MODERN TRADING IN THE AGE OF ARTIFICIAL INTELLIGENCE: EXAMINING MARKET MANIPULATION AND THE CASE FOR REGULATION

Vasheesh Bhageerutty

A thesis submitted to the Auckland University of Technology in fulfilment of the requirements for the degree of
Doctor of Philosophy (PhD) in Law

October 2026

Faculty of Business, Economics and Law
Auckland University of Technology
Auckland, New Zealand

Abstract

The rapid adoption of artificial intelligence (AI) powered trading systems has transformed modern stock exchange markets, including the New Zealand Exchange (NZX). AI-driven trading strategies offer significant benefits in liquidity, speed, and market efficiency. However, they also introduce new risks to market integrity, particularly AI-enabled market manipulation, emergent collusion, and distorted price discovery arising from system design, optimisation objectives, and opaque decision-making. Legal frameworks developed to address human-centred misconduct are increasingly strained by the scale, speed, and opacity of AI-driven trading. This thesis examines whether New Zealand law can effectively regulate AI-driven market manipulation on the NZX and, if not, how the legal framework must evolve. It analyses the adequacy of existing public law mechanisms, including the market manipulation provisions of the Financial Markets Conduct Act 2013, and assesses the extent to which concepts such as foreseeability, knowledge, and causation can be applied to AI-enabled trading systems.

The thesis also explores the role of private law remedies in promoting accountability and victim compensation. It assesses the applicability of negligence and statutory product liability under the Consumer Guarantees Act and Fair Trading Act, identifying doctrinal constraints, particularly in cases involving pure economic loss, distributed agency, and the absence of a direct consumer relationship. In response, it proposes two targeted reforms: a statutory product liability regime under the FMC Act imposing strict liability for defective AI financial products, and a new tort of AI Governance Liability grounded in a statutory duty of diligent governance.

Ultimately, the thesis proposes an integrated legal and governance framework that combines preventive regulation with effective enforcement and compensation mechanisms. Key recommendations include pre-market approval for high-risk AI trading systems, enhanced transparency and explainability obligations, compliance by design requirements, continuous monitoring with tested kill switches, and clear board level accountability. It concludes that while New Zealand's existing legal architecture provides a foundation for regulating AI-driven trading, targeted reforms are required to safeguard market integrity and maintain investor confidence in an increasingly automated financial ecosystem.

Attestation of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person (except where explicitly defined in the acknowledgements), nor used artificial intelligence tools or generative artificial intelligence tools (unless it is clearly stated, and referenced, along with the purpose of use), nor material which to a substantial extent has been submitted for the award of any other degree or diploma of a university or other institution of higher learning.

Signature:

Date: 28 October 2026

Acknowledgements

First and foremost, I would like to express my deepest and most sincere gratitude to my primary supervisor, Professor Allan Beever. His intellectual guidance, unwavering support, and rigorous academic standards have been fundamental to the successful completion of this thesis. His generosity in sharing his knowledge, his thoughtful feedback, and his insistence on clarity and precision have profoundly strengthened my legal reasoning and academic writing. Under his supervision, I have been continually challenged to think more deeply, write more carefully, and push beyond my own perceived limitations. I would also like to extend my sincere thanks to my secondary supervisor, Dr Moshood Abdussalam, whose guidance and support have been invaluable throughout this research journey.

I am grateful to the Auckland University of Technology for providing a supportive academic environment and the institutional resources necessary to undertake and complete this research. The University's support throughout this journey has been instrumental in my personal and professional growth, and I am thankful for the opportunities it has afforded me.

Above all, I owe an immeasurable debt of gratitude to my wife, Pooja. This thesis would not have been possible without her constant love, belief in me, and extraordinary support. She stood by me through countless sleepless nights, periods of self-doubt, and the long, demanding process of research and writing. Her patience, deep understanding, and innumerable sacrifices, including postponed plans, missed celebrations, and daily compromises, sustained me throughout this journey. Completing this thesis while navigating the challenges of life after COVID, in a new country far from family and familiar support networks, was not easy. During this time, my wife’s encouragement and resilience carried me through the most challenging moments, and her faith in me never wavered, even when mine did.

Table of Contents

Abstract ..... **i**
Attestation of Authorship ..... **ii**
Acknowledgements ..... **iii**
Chapter 1: Introduction **1**
1.1. RESEARCH BACKGROUND 3
1.2. PURPOSE 10
1.3. RESEARCH METHODS 11
1.4. STRENGTHS AND WEAKNESSES OF CHOSEN RESEARCH METHODS 15
1.5. SIGNIFICANCE OF STUDY 17
1.6. THESIS OUTLINE 20
Chapter 2: Decoding Algorithmic Trading: A Journey from Traditional Markets to the Age of AI **24**
2.1. CHAPTER OVERVIEW 24
2.2. ALGORITHMIC TRADING: PUTTING IT TOGETHER 25
2.2.1 Defining Market Fundamentals 25
a) Securities, Stocks or Shares 25
b) Virtual or Digital Assets 25
c) Trading Venues 26
d) Market Participants 27
e) Market Supervisors 27
2.2.2 Defining Algo-Trading (including HFT) 28
2.2.3 Defining AI (including ML) 35
2.3. MODERN TRADING: FROM FLOOR TO ALGORITHM TO AI 37
2.4. THE APPLICATION OF AI IN MODERN ALGO-TRADING 43
2.4.1. Forecasting Stock Market Trends 43
2.4.2. Leveraging Predictive Analytics in Trading Strategies 44
2.4.3. Sentiment Analysis 45
2.4.4. Risk Assessment and Management 47
2.5. CONCLUSION 48
Chapter 3: AI-Driven Algorithmic Trading and Emerging Risks **50**
3.1. THE RISK OF HARM: ALGORITHMIC TRADING AND POTENTIAL MARKET DETRIMENT 51

3.1.1 Cyber Risks: Risks of Operational Failure & Inherent Errors 51
3.1.2. Malicious Use by Humans 54
3.1.3. Trading Misconducts 55
a) Collusion 55
b) Insider Trading 59
c) Market Manipulation (Thesis Focus) 61
3.2 THE LIABILITY CONUNDRUM: ASSIGNING RESPONSIBILITY FOR AI-DRIVEN HARM 68
3.3 CONCLUSION 78
Chapter 4: Market Manipulation in New Zealand: The Regulatory Framework of the FMC Act 79
4.1. DEFINITION 79
4.2. PROHIBITION OF MANIPULATION 85
4.3. CONSEQUENCES OF BREACH 103
4.4. CONCLUSION 104
Chapter 5: Case Law Trends in Addressing Modern Market Manipulation 107
5.1. INTRODUCTION 107
5.2. NEW ZEALAND 108
5.3. INTERNATIONAL CONTEXT 114
5.4. LESSONS 128
5.5. CONCLUSION 128
Chapter 6: Negligence and Product Liability 130
6.1. INTRODUCTION 130
6.2. NEGLIGENCE 131
6.2.1. A Rights-Based Account of Negligence 132
6.2.2. Duty of Care 134
6.2.3. Breach of Duty and Standard of Care 140
6.2.4. Causation 142
6.2.5. Damage 145
6.2.6. Comparative Case Law 146
6.3 PRODUCT LIABILITY 147
6.3.1. The Consumer Guarantees Act 147
6.3.2 The Fair Trading Act 159
6.4. CHOOSING BETWEEN NEGLIGENCE AND PRODUCT LIABILITY 165
6.5. CONCLUSION 169
Chapter 7: Pathways Forward - Achieving an Effective Regulatory Framework 173
7.1. OVERVIEW 173

7.2. WHY A DEDICATED AI TRADING COMPLIANCE REGIME? 176
7.3. PRE-MARKET APPROVAL: CONFORMITY ASSESSMENTS AND COMPLIANCE AUDITS 181
7.4. COMPLIANCE BY DESIGN AND SPEED BUMPS OR CIRCUIT BREAKERS 189
7.5. TRAINING AND AWARENESS 192
7.6. HUMAN OVERSIGHT, GOVERNANCE AND RESPONSIBILITY FRAMEWORK 195
7.7. RISK MANAGEMENT FRAMEWORK 198
7.8. MARKET SURVEILLANCE AND ADDITIONAL MEASURES 202
7.9. AI GOVERNANCE LIABILITY: A STATUTORY CAUSE OF ACTION FOR SUPERVISORY FAILURES 207
7.10. NOVEL PRODUCT LIABILITY FRAMEWORK 214
7.10.1. Rationale 214
7.10.2. Comparative Inspiration: The EU Product Liability Directive 2024 216
7.10.3. Proposed FMC Act Amendments 219
7.10.4. Policy Justification and Legal Coherence 223
7.11. VICTIM COMPENSATION 225
7.12. CONCLUSION 231
Chapter 8: Overall Conclusion 235
Bibliography 245
A. Legislation, Codes and Regulations 245
B. Cases 246
C. Official Publications and Regulatory Reports 248
D. Books 250
E. Chapters in Edited Books 251
F. Journal Articles 252
G. Online and Institutional Resources 256

Chapter 1: Introduction

"*Capital markets are undergoing a fundamental shift worldwide where robots, specifically complex algorithms, are increasingly replacing human brokers*" – *Maureen O'Hara.*[1]

The advent of Artificial Intelligence ("AI") has transformed the financial services sector, including the stock exchange market, making it more innovative and efficient.[2] The fundamentals of the modern market have evolved tremendously from traditional stock exchanges in terms of market structure, traders' way of trading, and price-discovery mechanisms.[3] Once primarily driven by human intuition, expert know-how, and analytical skills, stock trading, which carried the risk of bias and emotional stress-led errors, is now dominated by AI-driven control such as algorithmic trading.[4] The stock market is witnessing a significant shift from traditional, human-driven trading methods towards increasingly sophisticated AI-powered algorithmic trading systems.[5] This research study decodes the radical transformation of stock exchanges, emerging legal risks associated with AI-driven trading, and how crafting a well-designed regulatory framework is key to maintaining market integrity and confidence.[6]

In New Zealand, the 2018 Thematic Review by the New Zealand Exchange ("NZX") serves as the genesis of an in-depth analysis of the development and impact of algorithmic trading on the NZX.[7] Given algorithmic trading has been creating a "noise factor", the thematic review went on to define algorithms as "computer programmes designed to replicate the set of steps a human dealer would take to execute an order in the market" and algorithmic trading ("algo-trading") as "trading in financial products where an algorithm automatically determines

---

[1] Maureen O'Hara "High frequency market microstructure" (2015) 116(2) Journal of Financial Economics 25 at 258 <https://doi.org/10.1016/j.jfineco.2015.01.003>.
[2] Tom CW Lin "The New Investor" (2013) 60 UCLA Law Review 678 at 700–709.
[3] Andrei Kirilenko and Andrew W Lo "Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents" (2013) 27 Journal of Economic Perspectives 51 at 53–57.
[4] O'Hara, above n 1, at 260–263.
[5] NZX *Algorithmic Trading: Thematic Review* (NZX, Wellington, December 2018) at 5–7.
[6] Karen Yeung "Algorithmic Regulation: A Critical Interrogation" (2018) 12 Regulation & Governance 505 at 510–517.
[7] NZX, above n 5, at 5–6.

the management of orders concerning timing, price, and quantity, based on predetermined parameters, rules, and conditions".[8] The setting out of these definitions has ushered a starting point to understand the various risks posed by AI-trading innovations and why regulation is vital.[9]

AI's technological advantages have become indispensable for driving progress in the stock exchange market.[10] Clearly, AI surpasses the inherent capabilities of human traders, rendering them unable to compete against it effectively.[11] The intersection between AI and human traders presents a complex legal landscape within the realm of algo-trading. This complexity has become a focal point of ongoing debate in the literature.[12] On the one hand, algo-trading is innovating rapidly; on the other hand, the regulatory framework remains quasi-stagnant, presenting significant challenges to the concept of liability for market misconduct.[13] Who is to be held responsible for AI-led trading abuses or market malpractices? There are several types of market misconduct, including collusion, insider trading and market manipulation. Notably, the primary challenge lies in effectively combating market manipulation due to its potential to jeopardise market integrity, undermine investor confidence and distort price discovery.[14]

How can we ensure that the exchange markets remain fair and orderly? This Chapter introduces the fascinating and complex world of algo-trading, briefly examining how the existing regulatory framework is currently addressing the evolving risks associated with AI-driven stock trading, including market manipulation and exploring potential avenues for its development to ensure adequate safeguards.[15] It refers to other leading researchers' critical findings and emphasises the need for New Zealand to develop a feasible regulatory framework to combat these challenges.

---

[8] NZX, above n 5, at 7.
[9] NZX, above n 5, at 9.
[10] Frank Pasquale *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press, Cambridge (MA), 2015) at 100–115.
[11] Lin, above n 2, at 703–706.
[12] Alessio Azzutti, Wolf-Georg Ringe and H Siegfried Stiehl "Machine Learning, Market Manipulation, and Collusion on Capital Markets: Why the 'Black Box' Matters" (2021) 43 University of Pennsylvania Journal of International Law 79 at 86–92.
[13] Yeung, above n 6, at 514–516.
[14] International Organization of Securities Commissions *Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency* (Final Report, October 2011) at 22–28.
[15] Andrew Tutt "An FDA for Algorithms" (2017) 69 Administrative Law Review 83 at 92–96.

The central argument developed through this thesis is that the challenge posed by AI-driven trading is not simply whether existing market manipulation law is capable of applying to conduct carried out through an algorithm. The more difficult question is whether New Zealand's existing legal framework can effectively prevent harmful conduct, attribute responsibility where trading decisions are increasingly autonomous or opaque, and provide meaningful redress where loss occurs. The thesis therefore examines the adequacy of the existing regulatory and private law framework before considering whether targeted reforms are required. In doing so, it connects the technological and regulatory challenges identified in the earlier chapters with the accountability, compensation and governance reforms developed in Chapter Seven.

This thesis situates itself within the existing scholarship on algorithmic trading, market manipulation, AI liability and financial regulation, while focusing specifically on how those areas intersect within New Zealand's legal and institutional framework. Rather than treating AI-driven market manipulation as an entirely new legal problem, the thesis examines the extent to which existing legal principles remain capable of responding to technologically mediated conduct and identifies the particular areas in which autonomy, opacity and the speed of algorithmic trading expose regulatory, accountability or remedial gaps.

1.1. RESEARCH BACKGROUND

According to Lee, Rui and Wang, large-scale electronic trading started in 1971 following the setting up of the first trading venue, the National Association of Securities Dealers Automated Quotations ("NASDAQ"). This milestone paved the way for the progressive use of computerised trading, leading to the emergence of other electronic trading venues known as "electronic communications networks" and, eventually, the introduction of algorithms to facilitate the electronic entry and execution of orders.[16] These algorithms operated based on a pre-determined set of rules.[17] Subsequently, McGowan highlights that the past decade has witnessed a dramatic change in how securities are traded on the capital markets, driven by the rise of sophisticated computer algorithms and high-speed networks to analyse

[16] Bong-Soo Lee, Oliver Meng Rui and Steven Shuye Wang "Information transmission between the NASDAQ and Asian second board markets" (2004) 28(7) Journal of Banking & Finance 1637 at at 1639–1642. <https://doi.org/10.1016/j.jbankfin.2003.05.001>.
[17] O'Hara, above n 1, at 259.

voluminous data, identify intricate patterns, and execute trades with remarkable precision in fractions of a second.[18] Algo-trading has virtually phased out the floor traders, brokers and smaller investors.[19] Menkveld identified that algo-trading makes trading faster and labour cheaper, resulting in efficient order execution and minimising price impacts.[20]

In 2018, the NZX, market operator and frontline regulator, focused on algo-trading in its thematic review, particularly recognising its influence on international markets and New Zealand.[21] It underlines that the primary or general forms of algo-trading have been deployed for on-market trading for many years. The review also highlights the increasing development and use of more sophisticated trading strategies on the NZX market, including high-frequency trading.[22] Alarmingly, it was noted that there are currently no specific rules governing algorithmic trading in New Zealand. As such, there are no prescribed requirements for the "development and testing of trading algorithms, operated by or through a separate participant identification of orders generated by algorithmic programmes, the types of, and parameters for, pre-trade filters and post-trade monitoring and the review of those controls".[23]

The transformation to algo-trading is further amplified by the integration of AI, particularly machine learning ("ML") and deep learning methodologies.[24] These changes have considerably impacted markets like Europe and the United States.[25] Since the last thematic review by the NZX on algo-trading was conducted in 2018, concrete data on the current prevalence of AI trading in the New Zealand Exchange market is limited. The absence of current publicly available data limits the conclusions that can properly be drawn about the prevalence of AI-driven trading on the NZX. This is a methodological limitation of the present research and is addressed further in section 1.4.

However, AI use in trading is likely growing with the automation trend, making it a natural evolution for the financial industry. AI offers a plethora of benefits in terms of enhanced

---

[18] Michael McGowan "The Rise of Computerized High Frequency Trading: Use and Controversy" (2010) 9 Duke Law & Technology Review 1 at 4–9.
[19] McGowan, above n 18, at 10–14; see also O'Hara, above n 1, at 260.
[20] Albert J Menkveld "The Economics of High-Frequency Trading: Taking Stock" (2016) 8(1) *Annual Review of Financial Economics* 1 at 8-14 <https://doi.org/10.1146/annurev-financial-121415-033010>.
[21] NZX, above n 5, at 5.
[22] NZX, above n 5, at 8–10.
[23] NZX, above n 5, at 11.
[24] Bryan H Choi "Software as a Profession" (2020) 33 *Harvard Journal of Law & Technology* 557 at 572–586.
[25] O'Hara, above n 1, at 261–263.

real-time data analysis over milliseconds, improved predictive capabilities, ultra-fast execution and the development of intricate trading strategies without the need for any human intervention. Similarly, it helps traders maximise fleeting market trends and fluctuations that generate more profit. These algorithms can operate 24/7, enabling continuous trade executions that provide greater flexibility and responsiveness across global markets operating in different time zones. Given their autonomous learning abilities, they can auto-learn from and auto-unlearn market experiences, refine these strategies, and devise their own risk management techniques to maintain an ongoing competitive edge over traditional traders or investors. Hence, it minimises loss and downside risk to a greater extent during highly volatile market conditions.[26]

While modern algorithmic trading has undoubtedly brought numerous advantages to the financial markets, it has simultaneously amplified several risk factors.[27] Market participants at large, including traders, confront various limitations in operating, monitoring, and regulating the modern market. Khashanah, Florescu, and Yang emphasised that advanced algo-trading provides new ways to manipulate the market,[28] increases market manipulation levels[29] and renders monitoring and detecting market manipulation much more challenging.[30] The IOSCO report further reveals how market manipulation may harm the market integrity and undermine public confidence in the securities markets by distorting prices, damaging the hedging functions, and creating an artificial appearance of market activity.[31]

Per Yadav, the most expressive example of systemic risk posed by algo-trading is extreme 'bad volatility' episodes known as flash crashes.[32] Often, algo-trading is linked to the

---

[26] Fatima Dakalbab and others "Artificial intelligence techniques in financial trading: A systematic literature review" (2024) 36(3) *Journal of King Saud University*: Computer *and Information Sciences* 102015 at 7–12 https://doi.org/10.1016/j.jksuci.2024.102015.

[27] Yesha Yadav "How Algorithmic Trading Undermines Efficiency in Capital Markets" (2015) 68 *Vanderbilt Law Review* 1607 at 1610–1614.

[28] Kiyohiko G Nishimura, 'Electronic Trading and Financial Markets' (International Paris-Europlace Financial Forum, Tokyo, 29 November 2010) at 4–8; Tom CW Lin, 'The New Market Manipulation' (2017) 66 *Emory Law Journal* 1253 at at 1260–1268.

[29] Charles R Korsmo, 'High-Frequency Trading: A Regulatory Strategy' (2013) 48 *University of Richmond Law Review* 523 at at 540–548.

[30] Joseph E Stiglitz "Tapping the Brakes: Are Less Active Markets Safer and Better for the Economy?" (paper, Federal Reserve Bank of Atlanta 2014 Financial Markets Conference, 15 April 2014) at 3–9; Matt Prewitt "High-Frequency Trading: Should Regulators Do More?" (2012) 19 *Michigan Telecommunications and Technology Law Review* 131 at 141–151.

[31] International Organization of Securities Commissions Technical Committee *Investigating and Prosecuting Market Manipulation* (May 2000) at 4–9.

[32] Yesha Yadav "Algorithmic Trading and Market Regulation" in Walter Mattli (ed) *Global Algorithmic Capital Markets: High Frequency Trading, Dark Pools, and Regulatory Challenges* (Oxford University Press, Oxford, 2019) 232 at 240–248; O'Hara, above n 1, at 263–266.

2010 Flash Crash, a United States trillion-dollar crash where Navinder Singh Sarao, the trader, used an automated algorithmic trading software to engage in a market manipulation tactic called 'spoofing'. The prices of different securities collapsed and rebounded to their original values very rapidly – all within 30 minutes.[33] A joint investigation by the US Securities and Exchange Commission ("SEC") and the US Commodity Futures Trading Commission ("CFTC") asserted that this market incident was triggered by a misbehaving execution algorithm and exacerbated by the action or inaction of the specialised algo-traders.[34] Navinder Singh Sarao was found guilty of fraud, spoofing, and market manipulation charges.[35]

While New Zealand has not experienced a similar 'flash crash' on the scale of the 2010 event, it is crucial to acknowledge that this does not equate to immunity. The smaller size of the New Zealand exchange market can offer some degree of insulation from certain types of systemic risks that might impact larger and more interconnected markets. However, vulnerabilities remain. A sudden global market downturn could have substantial ripple effects on the New Zealand exchange market, potentially causing significant disruptions. The NZX study noted that advanced algo-trading mechanisms like High-Frequency Trading ("HFT") substantially increased the "noise factor" in the New Zealand exchange market due to their high speed and interconnectivity.[36]

Azzutti and Ors identified two emerging market abuse risks posed by trading algorithms: market manipulation and tacit collusion. Under New Zealand's transaction-based market manipulation regime, liability does not depend upon proof of a separate specific intention to manipulate the market. The statutory inquiry focuses on whether the conduct has, or is likely to have, the prohibited effect and whether the person knew or ought reasonably to have known of that effect. A trader's purpose may nevertheless be relevant evidence when determining whether particular trading was manipulative.[37]

---

[33] Andrei A Kirilenko, Albert S Kyle, Mehrdad Samadi and Tugkan Tuzun "The Flash Crash: High-Frequency Trading in an Electronic Market" (2017) 72 *Journal of Finance* 967 at 970–976.
[34] US Commodity Futures Trading Commission and US Securities and Exchange Commission *Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues* (30 September 2010) at 1–6.
[35] *Sarao v Government of the United States of America* [2016] EWHC 2737 (Admin) at [11] – [24]; *United States v Sarao* No 1:15-cr-00075 (ND Ill, 28 January 2020) at [1] – [10].
[36] NZX, above n 5, at 11–12.
[37] Financial Markets Authority v Warminger [2017] NZHC 327 at [60] - [61] and [68].

Tacit collusion describes the phenomenon where firms recognise their interdependence in the market and, without any agreement, all start to adopt similar pricing and output strategies. While not inherently illegal, such behaviour can result in prices that exceed competitive levels, potentially harming market participants. Essentially, they questioned the adequacy of the existing regulatory frameworks, such as the traditional liability concepts, including "intent" and "causation," and enforcement mechanisms governing algo-trading to cope with the deepening risks of these market misconducts.[38]

The NZX Participant Rules deal with Market Participants' binding obligations (including those of brokers and other firms) in relation to trading on the NZX market. Primarily, they focus on market participants' disclosure and conduct obligations towards investors rather than on promoting and maintaining orderly markets. This emphasis presents a critical issue when deploying sophisticated AI-based algorithms.[39] The NZX has a monitoring framework for identifying and addressing 'anomalous' algorithmic trades that involves monitoring trades for unusual patterns or deviations from expected behaviour.[40] In cases where anomalies are detected and considered harmful to market integrity, the NZX may impose public censures and regulatory sanctions on the responsible parties.[41] Separately, the Financial Markets Authority ("FMA") actively enforces market manipulation regulations, although recent cases have not exclusively targeted algo-trading.[42] Furthermore, the NZX's Surveillance team monitors activities that may indicate market misconduct, including technical rule breaches and unlawful conduct such as market manipulation and insider dealing.

Part 5, subpart 3 of the Financial Markets Conduct Act 2013 ("FMC Act") prohibits information based and transaction-based market manipulation in New Zealand. Contraventions of ss 262 and 265 are Part 5 market provisions that may give rise to civil liability under pt 8 of the FMC Act. The available civil consequences include declarations of contravention, pecuniary penalties and compensatory orders. A compensatory order may be sought by the FMA or another person where an aggrieved person has suffered, or is likely to suffer, loss or

---

[38] Azzutti, Ringe and Stiehl, above n 12, at 92–99 and 130–137.
[39] NZX, above n 5, at 18–21.
[40] Alain P Chaboud and others "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market" (2014) 69(5) The Journal of Finance 2045 at at 2050–2058. <https://doi.org/10.1111/jofi.12186>.
[41] NZX, NZ RegCo, About NZ RegCo <https://www.nzx.com/regulation/nzregco/about-nzregco>
[42] Nam Ngo and others *Market Cleanliness in the New Zealand Equity Market* (Financial Markets Authority, Occasional Paper Series, October 2024) at at 5–14 < https://www.fma.govt.nz/assets/Research/Market-Cleanliness-Report-2024.pdf>.

damage because of the contravention. In sufficiently serious circumstances, market manipulation may also attract criminal liability. Accordingly, the statutory framework performs both enforcement and compensatory functions, although establishing entitlement to compensation may remain difficult in cases involving complex AI-driven trading and diffuse market loss.[43]

Depending on the circumstances, the FMA may use a range of enforcement responses, including warnings, civil proceedings and prosecution. Importantly, the FMC Act is not confined to punitive or deterrent remedies. As noted above, its civil liability regime also provides a statutory route through which compensation may be sought for loss caused by a contravention. However, the only manipulation cases recently tried in New Zealand are *FMA v Henry*[44] in 2014 and *FMA v Warminger*[45] in 2017. In these cases, the FMA issued civil proceedings in the High Court against Henry and Warminger for alleged market manipulation in breach of the provisions of the Securities Markets Act 1988 (now repealed and replaced by the FMC Act). The proceedings resulted in findings of contravention and the imposition of pecuniary penalties. These cases remain important to the present research because they demonstrate how New Zealand courts approach market effect, knowledge and trading purpose, concepts that become more difficult to apply where trading decisions are mediated by increasingly autonomous systems.

In response to the increasing prevalence of complex algo-trading strategies across EU markets, many member states such as Germany,[46] Italy,[47] and France[48] have recognised the risk of market disruptions and enacted significant legal provisions applicable to traders, trading venues and market intermediaries. The MiFID II introduced a risk-based approach to regulating algo-trading while promoting fair competition and investor protection. Their framework imposes strict requirements for investment firms using algorithmic trading. These include trading limits and thresholds to prevent errors and ensure market stability, business continuity arrangements to deal with any operational failure, notifications to authorities and trading

[43] Financial Markets Conduct Act 2013, ss 262-269, 385 and 485-495, particularly ss 494 and 495.
[44] *FMA v Henry* [2014] NZHC 1853 at [70] - [78].
[45] *FMA v Warminger* [2017] NZHC 327 at at [73]– [80].
[46] Martin Haferkorn and Kai Zimmermann "The German High-Frequency Trading Act" (2014) SSRN Working Paper No 2514334 at 4–9.
[47] Commissione Nazionale per le Società e la Borsa (CONSOB) *Algotrading and HFT: Mapping of the Regulatory Framework and Reporting Obligations to Consob* (CONSOB, 2023) at 6–13.
[48] Code monétaire et financier (France) arts L533-10-3 to L533-10-7 <https://www.legifrance.gouv.fr>.

venues about the nature of algorithmic trading strategies and details of their compliance risk controls.[49]

Some firms are subject to stricter obligations, such as continuous trading during specified hours and agreements with trading venues. Moreover, firms providing direct electronic access to a trading venue are responsible for managing risks, monitoring clients to prevent market abuse, maintaining proper records and ensuring compliance with the trading venue rules. The European Securities and Markets Authority is responsible for clarifying these rules through developing detailed technical standards.[50]

The EU's approach to regulating algo-trading highlights the importance of a forward-looking regulatory framework and serves as a good example for New Zealand. By closely observing the EU's regulatory developments, including the challenges and successes they achieve, New Zealand can identify and assess potential risks associated with the evolving sophistication of algo-trading within its own market. While New Zealand may not currently prioritise specific regulations for advanced AI-enabled algo-trading, the EU's experience will allow the regulators to adopt a more informed and measured response to adapt and refine their oversight frameworks as needed.

The comparative analysis is not undertaken on the assumption that overseas regimes can simply be transplanted into New Zealand. Rather, international approaches are used to identify regulatory techniques that respond to problems also identified in the New Zealand framework and to assess whether those techniques can be adapted to the size, structure and institutional characteristics of the New Zealand market. This analysis ultimately informs the preventative and governance-based reforms developed in Chapter Seven.

The central research question is: how should New Zealand's legal and regulatory framework respond to market manipulation and related harm arising from the increasing use of AI-driven algorithmic trading?

---

[49] Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on Markets in Financial Instruments [2014] OJ L173/349 (MiFID II), art 17.
[50] MiFID II, above n 49, art 17(3) - (5).

- How prevalent is AI-driven trading in the New Zealand financial market compared to the global trends?
- What specific risks are introduced or amplified by the use of AI in algorithmic trading, considering the unique characteristics of the NZX market?
- What manipulative strategies can be perpetrated using AI - spoofing, layering, or pump-and-dump schemes?
- How do current market manipulation laws, NZX rules and FMA monitoring frameworks address these risks and their critical limitations?
- What lessons can be learned from significant market events, such as the 2010 flash crash involving algorithmic trading and international efforts to prevent such incidents?
- What regulatory approaches can be considered to mitigate these risks?
- How can modern regulatory frameworks such as compliance by design be implemented to enhance the regulation of AI-driven market manipulation?
- To what extent can civil law remedies, including negligence and statutory product liability under the Consumer Guarantees Act and Fair Trading Act, provide redress for victims of AI-driven market manipulation or defective trading systems?
- What are the doctrinal and evidentiary challenges in applying traditional tort principles to AI-mediated harm, and how might these be addressed through legal reform?
- What combination of preventative governance, compliance by design, statutory product liability and AI Governance Liability would best address the regulatory, accountability and remedial gaps identified by this thesis?

1.2. PURPOSE

To address the above research questions, this thesis investigates the evolving landscape of market manipulation in New Zealand, focusing on challenges posed by the increasing use of AI in algo-trading. This research study:

- analyses the historical developments and inner workings of the securities market, mainly focusing on the evolution of trading technologies from basic rule-based systems to sophisticated ML models;
- examines the intersection of AI in algo-trading and its profound impact on market dynamics and modern trading malpractices in terms of insider dealing and market manipulation;

- analyses New Zealand's existing market manipulation framework and critically evaluates their adequacy in addressing AI-driven manipulation;
- examines both the statutory compensation mechanisms available under the FMC Act and the private law avenues available to persons suffering loss, particularly negligence and statutory product liability under the Consumer Guarantees Act 1993 and Fair Trading Act 1986, and critically evaluates the limitations of those mechanisms in the context of AI-driven trading; and
- develops an integrated reform framework that responds to the regulatory and remedial gaps identified in the preceding chapters, including preventative governance, pre-market approval for high-risk AI trading systems, compliance by design, continuing monitoring and accountability, statutory product liability under the FMC Act and a proposed tort of AI Governance Liability.

Taken together, these inquiries provide a doctrinal and policy analysis of the regulatory challenges posed by AI-driven market manipulation in New Zealand. The thesis does not proceed on the basis that existing law is wholly incapable of responding to AI. Rather, it identifies where existing public and private law mechanisms remain effective, where technological characteristics such as autonomy and opacity create practical or doctrinal difficulties, and where targeted reform is justified. The recommendations in Chapter Seven are therefore derived from the deficiencies identified through the preceding analysis rather than being presented as independent regulatory proposals.

1.3. RESEARCH METHODS

As previously discussed, this research aims to examine the regulation of modern algorithmic trading (including AI) and market manipulation in New Zealand. Given the nascent nature of this issue, the research will adopt a multi-faceted approach, combining legal doctrinal analysis, comparative legal analysis and case study analysis. This approach provides a comprehensive understanding of:

(i) the current legal landscape around market manipulation in New Zealand,

(ii) the regulatory developments made in other jurisdictions,

(iii) how market manipulation cases are dealt with to identify patterns and address AI-driven manipulation's potential challenges, and

(iv) the civil law doctrines, most notably negligence and product liability and the remedies for individuals or entities that suffer financial loss.

This study thoroughly analyses the relevant sources of law, including New Zealand legislation, namely the Financial Markets Conduct Act 2013 and other applicable laws and regulations, identifies provisions related to market manipulation and reveals how these could be relevant to address modern algorithmic trading. Moreover, it reviews pertinent case law in New Zealand and internationally, showing how the Courts have interpreted and applied market manipulation laws in the context of traditional manipulation techniques. This examination provides a baseline for understanding the potential application of these laws to AI-driven manipulation.

Furthermore, this study examines secondary authorities, including regulatory guidance, rules, and policies issued by relevant agencies such as the FMA, law reviews, government reports, academic commentary, legal dictionaries, research papers, journals and news articles, which are within the boundaries of the research question. This will help understand the current approach to market oversight and enforcement. It involves conventional legal research methods to understand, describe, and critically examine the regulatory regime governing market manipulation and algo-trading within the NZX market.

In addition to the civil liability and compensatory mechanisms available under the FMC Act, this research also investigates the role of private law in addressing harm caused by AI-driven trading systems. This includes a detailed doctrinal analysis of tort law, with a particular focus on negligence, and statutory remedies under the Consumer Guarantees Act 1993 and the Fair Trading Act 1986. The study examines the doctrinal constraints of these frameworks, especially in cases involving pure economic loss, distributed agency, and the absence of a direct consumer relationship. It evaluates the extent to which these private law mechanisms can provide redress to victims of algorithmic manipulation and identifies the evidentiary and conceptual challenges that arise in attributing liability in automated environments. This civil liability analysis is complemented by the development of proposed legal reforms, examined through a combination of doctrinal reasoning and policy analysis, assessing their compatibility with existing legal principles and their potential to fill the accountability gaps identified in the current framework.

The methods perform different but complementary functions. Doctrinal analysis establishes the content and limits of existing New Zealand law. Comparative analysis identifies alternative regulatory techniques and provides a basis for evaluating possible reform. Case study analysis tests the doctrinal analysis against concrete examples of technologically mediated market conduct. The thesis then uses normative legal analysis to determine whether the deficiencies identified through those methods justify reform and, if so, what form that reform should take.[51]

The research applies a hierarchy of sources. Primary New Zealand legal sources, particularly legislation and judicial decisions, are authoritative for determining the content of existing law. NZX Rules and official regulatory materials are used to establish the applicable institutional and regulatory framework. Scholarly literature is used to identify competing interpretations, situate the research within existing knowledge and critically evaluate possible reforms. International legal materials are used comparatively and are not treated as authority for the content of New Zealand law.

The thesis also uses comparative analysis because several overseas jurisdictions have greater regulatory experience with algorithmic trading and technologically mediated market misconduct. The European Union and the United States are examined because they operate developed electronic financial markets, have experienced significant issues involving algorithmic trading, and have developed regulatory responses relevant to the problems considered in this thesis. The comparative analysis identifies lessons, regulatory approaches and practices that may inform the development of New Zealand's framework for AI-driven market manipulation. However, those overseas approaches are not treated as directly transferable to New Zealand. Regulatory measures developed for larger and more liquid markets must be considered in light of the structure of the FMC Act, the institutional roles of the FMA and NZX, and the size and characteristics of the New Zealand market. Comparative law is therefore used to identify and evaluate regulatory options that may be adapted to the New Zealand context, rather than as a basis for automatically transplanting overseas rules.[52]

---

[51] Mark Van Hoecke "Legal Doctrine: Which Method(s) for What Kind of Discipline?" in Mark Van Hoecke (ed) *Methodologies of Legal Research* (Hart Publishing, Oxford, 2011) 1 at 4–19; Terry Hutchinson "The Doctrinal Method: Incorporating Interdisciplinary Methods in Reforming the Law" (2015) 8 *Erasmus Law Review* 130 at 132–138.

[52] Mathias Reimann "The Progress and Failure of Comparative Law in the Second Half of the Twentieth Century" (2002) 50 *American Journal of Comparative Law* 671 at 696–700.

The research design is as follows:

a) The introduction to the research topic starts with the evolution of the trading market, including developments in algo-trading and HFT, the growing use of emerging technologies, particularly AI, and its presence in the New Zealand Exchange market. It further analyses and refines the legal definition of algo-trading and related terms, intending to discuss whether AI-driven manipulation falls within the ambit of existing laws.
b) Introduction to a risk-based approach: it is essential to identify the principal risks of algo-trading, particularly collusion and market manipulation. This thesis primarily emphasises on market manipulation due to the broader spectrum of related disruptive practices and their potentially profound implications for the market.
c) Analysis of the anti-manipulation framework in New Zealand: This involved examining the Financial Market Conduct Act, NZX Participant Rules, regulatory guidance, case laws, and other sources to identify the regulatory gaps.
d) Exploring legal approaches: considering different legal tools and thoroughly discussing them. This thesis will emphasise a risk, supervisory framework, and civil and administrative remedies. This includes a critical evaluation of the limitations of negligence, the CGA, and the FTA in addressing AI-related harm, and the development of new forms to ensure that victims are not left without recourse.
e) Final Recommendations: the findings from the technological, public law, case law and private law analyses are brought together to identify the reforms justified by the deficiencies established in the preceding chapters. Chapter Seven evaluates an integrated framework comprising pre-market approval for high-risk AI trading systems, compliance by design, continuing monitoring and governance obligations, enhanced regulatory surveillance, statutory product liability under the FMC Act and the proposed tort of AI Governance Liability. The recommendations are evaluated against legal coherence, market integrity, effective accountability and compensation, proportionality, technological neutrality and institutional feasibility.

1.4. STRENGTHS AND WEAKNESSES OF CHOSEN RESEARCH METHODS

The combination of doctrinal, comparative and case-study analysis is well-suited to the research question for four reasons. First, the question is doctrinal. It concerns the interpretation and application of New Zealand statutory and common-law provisions to a class of conduct (AI-mediated market manipulation) that has not yet been litigated in this jurisdiction. Doctrinal analysis is therefore the appropriate primary method.[53] It enables systematic examination of the FMC Act, NZX Participant Rules, FMA enforcement guidance and the relevant common-law authorities, and produces conclusions that are directly usable by regulators, judges and practitioners.[54]

Secondly, the question is normative. It asks not only whether existing law applies, but how it should evolve. Comparative analysis allows the thesis to draw on jurisdictions (EU, UK, US, Australia, Canada) which have engaged with AI-mediated harm more extensively, without committing to wholesale transplantation of foreign regimes.[55] Thirdly, the question is empirically informed but not empirical. Existing case studies - the 2010 Flash Crash, *Sarao*, *Coscia*, *Da Vinci Invest* provide concrete examples against which doctrinal proposals can be stress-tested.[56]

Fourthly, the question is systemic. It engages with the architecture of market regulation, private law and governance simultaneously. A purely doctrinal method, narrowly conceived, would understate those interactions; a purely socio-legal method would not produce the operational legal proposals that this thesis advances.

The combination of methods is also intended to reduce the limitations inherent in relying on any one method. Doctrinal analysis provides legal precision but cannot establish how frequently AI systems are currently used on the NZX or how proprietary systems operate in practice. Comparative analysis broadens the range of regulatory options but creates a risk of inappropriate transplantation from markets with different institutional and economic characteristics. Case studies provide concrete examples against which legal rules can be tested, but prominent incidents may be exceptional and cannot establish the prevalence of algorithmic

---

[53] Mark Van Hoecke “Legal Doctrine: Which Method(s) for What Kind of Discipline?” in Mark Van Hoecke (ed) *Methodologies of Legal Research* (Hart Publishing, Oxford, 2011) 1 at 4–19.
[54] Geoffrey Samuel "Is Legal Knowledge Cumulative?" (2012) 32 Legal Studies 448 at 459–469.
[55] Mathias Reimann "The Progress and Failure of Comparative Law in the Second Half of the Twentieth Century" (2002) 50 Am J Comp L 671 at 696–700.
[56] Terry Hutchinson "The Doctrinal Method: Incorporating Interdisciplinary Methods in Reforming the Law" (2015) 8 Erasmus L Rev 130 at 132–138.

misconduct. The thesis therefore uses each method for the purpose to which it is best suited and does not treat any one method as capable of answering the research question independently.

The methodology also has limitations, of which four warrant explicit acknowledgement. First, the analysis is doctrinal and conceptual; it does not include original empirical investigation of NZX trading data or of the use of AI by NZX-registered participants.[57] Empirical work of that kind is constrained by data-access barriers and falls outside the scope of a single doctoral study, but it represents a clear avenue for future research.

Secondly, the technological substrate is moving quickly. The proposals are designed to be technology-neutral and governance-focused, but any doctrinal account will need periodic recalibration as machine-learning techniques (and especially large language models) develop.[58] Thirdly, the comparative dimension is selective rather than exhaustive. The thesis draws principally on EU, UK, US, Australian and Canadian sources, and does not engage in depth with, for example, Singaporean or Hong Kong regulatory developments, which would warrant their own treatment. Fourthly, the thesis is primarily domestic in orientation. Cross-border enforcement, conflicts of laws and extraterritorial issues are flagged but not resolved and are identified as topics for future research.

A further limitation arises from the normative character of the recommendations. There is no purely doctrinal answer to how the law should balance technological innovation, market integrity, compensation and regulatory cost. The reform proposals therefore necessarily involve evaluative judgement. To make that judgement transparent, Chapter Seven assesses the proposed reforms against identified criteria: legal coherence, prevention of market harm, effective accountability and compensation, proportionality, technological neutrality and institutional feasibility. This provides an explicit basis for moving from the doctrinal identification of regulatory gaps to recommendations for legal reform.[59]

---

[57] For an example of the empirical work that would complement this doctrinal study, see Kirilenko, Kyle, Samadi and Tuzun, above n 33, at 980–985.
[58] See generally Lyria Bennett Moses "How to Think about Law, Regulation and Technology: Problems with 'Technology' as a Regulatory Target" (2013) 5 Law, Innovation & Technology 1 at 7–13.
[59] Terry Hutchinson "The Doctrinal Method: Incorporating Interdisciplinary Methods in Reforming the Law" (2015) 8 *Erasmus Law Review* 130 at 132–138.

1.5. SIGNIFICANCE OF STUDY

"Artificial intelligence rose to become a key technology in the 2010s and is assumed to be the dominating technology in the 2020s." - Yves Hilpisch[60]

The primary motivation for this research is to address a gap in the existing scholarly literature, while recognising that this gap is one of focus and integration rather than an absence of scholarship on the underlying subjects. There is a substantial international scholarship on financial-market manipulation,[61] together with a rapidly developing scholarship on AI and legal liability across a range of areas, including tort,[62] product liability,[63] regulation[64] and corporate governance.[65] There is also established New Zealand scholarship on the FMC Act and securities market regulation more generally.[66] What is less developed is a sustained doctrinal and policy analysis that brings these strands together to examine whether and how the FMC Act, supported by complementary private law and governance mechanisms, can respond effectively to AI-driven market manipulation in the New Zealand exchange market. The thesis therefore positions itself within, and seeks to contribute to, existing scholarship rather than asserting a vacuum of prior work.

---

[60] Yves Hilpisch Artificial Intelligence in Finance: A Python-Based Guide (O'Reilly Media, Sebastopol (CA), 2020) at xv.

[61] Daniel R Fischel and David J Ross "Should the Law Prohibit 'Manipulation' in Financial Markets?" (1991) 105 Harvard Law Review 503 at 510–520; Merritt B Fox, Lawrence R Glosten and Gabriel V Rauterberg *The New Stock Market: Law, Economics, and Policy* (Columbia University Press, New York, 2019) at 209–246.

[62] Ryan Abbott *The Reasonable Robot: Artificial Intelligence and the Law* (Cambridge University Press, Cambridge, 2020) at 47–73; Andrew D Selbst "Negligence and AI's Human Users" (2020) 100 Boston University Law Review 1315 at 1341–1356; David C Vladeck "Machines Without Principals: Liability Rules and Artificial Intelligence" (2014) 89 Washington Law Review 117 at 124–146; Mark Lemley and Bryan Casey "Remedies for Robots" (2019) 86 University of Chicago Law Review 1311 at 1322–1329.

[63] Catherine M Sharkey "Product Liability in the Digital Age: Liability for Algorithmic Harms" in Kerry McCausland (ed) *Research Handbook on Law and Artificial Intelligence* (Edward Elgar, Cheltenham, 2024) 211 at 218–227; Christiane Wendehorst "Strict Liability for AI and Other Emerging Technologies" (2020) 11 Journal of European Tort Law 150 at 161–169; Bryan H Choi "Software as a Profession" (2020) 33 Harvard Journal of Law & Technology 557 at 572–586.

[64] Tutt, above n 15, at 92–96; Margot E Kaminski "Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability" (2019) 92 Southern California Law Review 1529 at 1540–1551; Yeung, above n 6, at 510–517.

[65] Luca Enriques and Dirk Zetzsche "Corporate Technologies and the Tech Nirvana Fallacy" (2020) 72 Hastings Law Journal 55 at 75–99.

[66] Susan Watson "Investor Protection in New Zealand: A Tale of Two Statutes" (2014) 20 New Zealand Business Law Quarterly 245 at 248–260; David Brown "Securities Law: New Zealand" in *Encyclopaedia of New Zealand Law* (online ed, LexisNexis); Andrew Beck "The Financial Markets Conduct Act and Market Manipulation" [2014] NZLJ 354 at 355–358.

The distinct contribution of this thesis lies in bringing together existing scholarship on market manipulation, algorithmic trading, AI liability, private law and governance and applying those bodies of knowledge to the specific New Zealand regulatory context. The contribution claimed is therefore one of integration, application and reform, rather than an assertion that the underlying subjects have not previously been examined. It contributes to the academic literature on the intersection of AI, finance, and regulation and provides valuable insights for researchers, policymakers, and practitioners working in this area. It is also likely to stimulate further research, leading to a deeper understanding of the complex challenges posed by AI in the securities markets. Furthermore, to some extent, it can equip market participants with some knowledge and skills needed to navigate the evolving landscape of AI-driven trading.

The thesis makes three principal contributions. First, it evaluates whether the existing FMC Act market manipulation regime can accommodate AI-driven conduct while identifying the practical difficulties created by opacity, autonomy, knowledge and attribution. Secondly, it extends the analysis beyond public enforcement by examining the availability and limitations of compensation under the FMC Act, negligence and statutory product liability. Thirdly, it develops a governance-based reform framework that links preventative regulation with responsibility and redress, including statutory product liability under the FMC Act and AI Governance Liability. These contributions correspond directly to the legal deficiencies identified in Chapters Four, Five and Six.

When amplified by AI, market manipulation poses a significant threat to market integrity, as evidenced by events like the 2010 flash crash. Particularly, trading algorithms can be ill-intentionally programmed or used to manipulate the market, with the perpetrators hiding behind algorithms' black-box nature and taking advantage of such opacity in trade execution by autonomous AI systems. This study addresses the deliberate manipulation facilitated by AI and the potential for algorithmic behaviours to exacerbate market volatility. By examining and mitigating these risks in New Zealand, the study contributes to safeguarding market stability, restoring and maintaining investors' confidence in this modern era of algo-trading and ensuring fair trading practices. Likewise, developed algorithmic trading shall serve as a force for positive change in the financial ecosystem rather than a threat.

The findings of this research study are vital for informing the development of effective regulatory responses. It analyses the limitations of the FMC Act, existing NZX rules and FMA monitoring practices in the face of AI-driven manipulation and algorithmic instability. For example, it identifies that the FMC Act does not address algo-trading in any of its provisions, including those on market manipulation. It is, thus, questionable whether, in the current regulatory context, the NZX as a market operator and the FMA as a regulator can effectively deal with risks associated with algo-trading-driven market manipulation. Similarly, no cases or precedents in New Zealand have established how the court could potentially approach a situation where trading algorithms are directly involved in manipulating the market. Critically, it provides recommendations for adapting these frameworks, drawing lessons from events like the flash crash, to prevent similar occurrences in the New Zealand exchange market.

This study also contributes to the ongoing debate about the appropriate balance between innovation and regulation. It recognises that while AI has the potential to enhance market efficiency and liquidity, it also introduces new forms of risk that traditional legal frameworks are not equipped to manage. By proposing reforms such as a statutory product liability regime and the tort of AI Governance Liability, the thesis offers a roadmap for modernising New Zealand's legal infrastructure in a way that supports innovation while protecting market integrity. These proposals are grounded in legal principle but are also informed by practical considerations, ensuring that they are capable of implementation within the existing institutional landscape.

While the focus is on New Zealand, this research study draws comparisons with other jurisdictions that have dealt with similar issues. The study identifies best practices and potential pitfalls by examining how different countries have adapted their regulatory frameworks to address AI-driven manipulation. This comparative analysis adds a valuable dimension, allowing New Zealand to learn from international experience, particularly from the EU member states and avoid reinventing the wheel. It can also position New Zealand as a leader in addressing this emerging challenge. For instance, this research study explores different alternatives, such as compliance by design frameworks and their relevance in the context of algorithmic trading and market stability. Another key benefit of this study is its proactive approach. Examining the potential for AI-driven manipulation before it becomes widespread allows regulators to anticipate and mitigate risks before they materialise. This is far more effective than reacting to a crisis after significant damage has been done. The study will help

identify early warning signs and develop preventative measures, minimising the potential for market disruption.

In doing so, the thesis contributes to the resilience of New Zealand's financial markets by equipping regulators, courts, and market participants with the conceptual and legal tools necessary to respond to the challenges of AI. It also provides a foundation for future empirical and interdisciplinary research, encouraging collaboration between legal scholars, technologists, and financial regulators. By framing AI not merely as a technological innovation but as a governance challenge, the study shifts the regulatory conversation towards accountability, transparency, and systemic risk management.

1.6. THESIS OUTLINE

This thesis comprises eight chapters, each further split into separate sub-chapters and parts, as established in the table of contents. As discussed above, Chapter One only introduces the background and lays the groundwork for this thesis. The research background refers to the writings of leading researchers in this area, noting that there are limited sources on this innovative area of algorithmic trading. It also explains the purpose of this thesis and which specific issues are addressed, including the legal and regulatory challenges around market manipulation in the modern market and proposed regulatory solutions.

Moreover, it outlines the research methods employed throughout this thesis. It establishes the primary sources as legislation, case-law precedents, and studies and publications by regulatory bodies such as the NZX and NZX RegCo within New Zealand. References to the international context are also made as secondary sources. Additionally, the research design establishes the underlying strategy adopted, which includes a brief introduction to the topic, a risk assessment to identify the different risks, a focused examination of market manipulation, an analysis of the current anti-manipulation framework in New Zealand for loopholes and recommendations for an improved or regulatory framework. The final part of this Chapter sets out the significance of the study, explaining the importance of the research and its aims and objectives.

Next, Chapter Two defines algorithmic trading, AI, ML and associated key terms such as high-frequency trading, securities or stocks, virtual assets, trading venues, market participants,

market supervisors, and different trading strategies. Overall, it provides a customised definition of algorithmic trading that is tailored explicitly to the New Zealand context. It also provides a brief historical journey exploring the origins of the stock exchange markets, where traders first started buying and selling goods or shares and a comprehensive overview of how algorithms influence the financial landscape in modern exchanges. Subsequently, it dives into the development of trading practices and the recent rise of algorithmic trading within the NZX market and globally. It explores the changing landscape for investors and brokers, including the execution of trades at high speeds powered by innovative technologies such as AI, HFT and ML.

Chapter Three explores the transformative impact of AI and ML on the exchange market, highlighting the benefits and risks of these technologies. This Chapter focuses on the main advantages of sophisticated AI, such as reducing labour and costs for trading firms, enhanced market efficiency, volume and speed of order entry and execution for enhanced liquidity market participants and excellent investor protection. It also addresses the negative externalities encompassing risks of operational failures and inherent errors, malicious use by humans, autonomous decision-making, the "black box" problem and disruptive trading practices like collusion, insider trading and market manipulation.

Furthermore, it discusses adopting a risk-based approach to outline all risks associated with algorithmic trading and demonstrates its specialised focus on the issue of market manipulation. It explains market manipulation in more straightforward terms and introduces the two types of market manipulation: information-based and transaction-based. It also details other forms of market manipulation, namely deceptive trading activity and electronic front-running. This Chapter also refers to remarkable past market incidents in New Zealand and other relevant and international jurisdictions such as the US, EU, and UK.

Chapter Four examines the current laws in New Zealand designed to prevent market manipulation and how these might struggle to keep pace with algorithmic trading. It outlines the existing anti-market manipulation framework in place, covering the legal definition of market manipulation and which specific provisions of the law and NZX rules are applicable in this context. Moreover, it analyses the different classifications of market manipulation but focuses on the ones applicable to the New Zealand context. In doing so, reference is also made to international case precedents, mainly the US and a few domestic cases.

Chapter Five analyses how courts in New Zealand and across international jurisdictions have interpreted and applied market manipulation laws over time, despite the rapid evolution of trading technologies. By examining established precedents, the chapter provides a roadmap for how the judiciary is likely to treat AI-powered trading, particularly by reinforcing the view of AI as a sophisticated "tool" rather than a legal entity. This perspective suggests that courts will adapt existing statutory tests such as those found in the FMC Act to focus on the conduct of the human operator or the firm that deploys the technology, ensuring that the use of advanced algorithms does not shield participants from liability.

The Chapter Five further explores how legal arguments are shifting to accommodate the "black box" nature of AI while maintaining the integrity of market conduct rules. By drawing on domestic cases and international examples it illustrates a judicial trend toward assessing the objective "market effect." This analysis serves to guide future regulation by demonstrating that while the medium of manipulation has become more complex, the fundamental legal inquiry remains centred on whether the technology was used to create a false or misleading appearance of trading or to engineer an artificial price.

Chapter Six shifts the focus from public enforcement to private law remedies. It evaluates the potential for tort law to provide compensation to victims of AI-driven market manipulation, particularly through negligence and product liability claims. The chapter identifies significant doctrinal constraints in both areas, including the limitations of negligence in cases involving pure economic loss and the restricted scope of the CGA and FTA. It also explores the difficulties of attributing liability in automated environments, including the relevance of vicarious liability and agency principles. The chapter concludes that while private law has a role to play, it is insufficient on its own to address the full spectrum of harms arising from AI-driven trading.

Chapter Seven builds on these findings by proposing a comprehensive regulatory response. It argues for a shift from reactive enforcement to proactive governance, recognising that the speed and opacity of AI systems demand new tools and institutional capacities. The chapter introduces a compliance regime centred on pre-market approval, algorithmic registration, and continuous oversight. It also proposes two legal innovations: a statutory product liability regime under the FMC Act to address harms caused by defective AI financial products, and the tort of AI Governance Liability to hold firms accountable for systemic supervisory failures.

These proposals are designed to complement existing legal mechanisms and form part of a broader governance-based framework that integrates public regulation, private law, and institutional accountability.

Finally, Chapter Eight consolidates the findings of the thesis and articulates the overarching conclusions. It reaffirms that while AI does not require a wholesale reinvention of legal principles, it does necessitate a recalibration of regulatory structures to ensure that responsibility remains human-centred and that legal remedies remain effective. The chapter underscores the importance of embedding governance into the design and operation of AI systems, aligning incentives through insurance and compensation mechanisms, and maintaining a balance between innovation and integrity. It concludes that New Zealand already possesses important regulatory and private law foundations, but that those foundations require targeted supplementation to address the preventative, attribution and remedial challenges created by increasingly sophisticated AI trading systems.

Chapter 2: Decoding Algorithmic Trading: A Journey from Traditional Markets to the Age of AI

"The stock market is a complex system that no one fully understands, and nobody can predict with certainty where it will go next." - Burton Malkiel

2.1. CHAPTER OVERVIEW

This chapter provides the technological and conceptual foundation upon which the rest of the thesis is built, ensuring that the legal analysis in subsequent chapters is grounded in a clear understanding of how modern financial markets operate. By introducing the core mechanics of trading and tracing the evolution from traditional market structures to algorithmic and AI-driven environments, the chapter establishes the factual baseline against which later regulatory and doctrinal arguments are tested.

Section 2.2 analyses the key terms and concepts most relevant to the modern exchange market and develops working definitions of algorithmic trading and AI. This conceptual clarity is essential because, as later chapters demonstrate, the regulatory and doctrinal challenges posed by AI-driven trading flow directly from its distinctive technological features. Section 2.3 situates algorithmic trading and AI within the broader historical development of financial markets, showing that they are not isolated phenomena but the latest stage in a long process of automation. Section 2.4 examines how AI is currently applied within modern trading environments, identifying characteristics such as speed, opacity, autonomy, and adaptive behaviour that distinguish AI-driven trading from earlier forms of automation.

Together, these sections demonstrate why AI-driven trading poses regulatory challenges that the existing legal framework was not designed to address. The technological account developed here informs the analysis of AI-specific risks in Chapter 3, the assessment of the FMC Act's market manipulation regime in Chapter 4, the case-law analysis in Chapter 5, the doctrinal limits of private law explored in Chapter 6, and ultimately the governance-based reforms proposed in Chapter 7, including the new tort of AI Governance Liability.

Section 2.5 concludes by drawing these threads together and previewing the regulatory questions taken up in the remainder of the thesis.

2.2. ALGORITHMIC TRADING: PUTTING IT TOGETHER

Algo-trading is built on the foundation of traditional financial markets and trading practices. It does not replace the underlying trading principles; it is, in fact, an extension of them. Just as understanding the meaning of a sentence requires knowledge of its words, comprehending the concept of algo-trading necessitates a solid grasp of the key market terms and concepts. More to the point, these terms are the essential components that assist in unlocking the complexities of algo-trading and navigating the world of financial markets to identify their inherent risks and potential remedies. The key terms are as follows: securities, virtual assets, market participants, market supervisors, and trading venues. Detailed definitions of these key terms and algo-trading will follow in the subsequent section.

2.2.1 Defining Market Fundamentals

a) Securities, Stocks or Shares

In essence, a security, stock or share is a tradeable financial instrument or futures contract forming a contractual agreement allowing the holders to claim the cashflows and control of a particular business.[67] The buyers, for example, will initially buy equity or debts in exchange for money, which will, in turn, obtain the promise of some control rights and entitlements of residual earnings over a given company. The modern market is now trading new commodities, such as digital or virtual assets, such as cryptocurrencies and other digital currencies, adding an innovation layer.

b) Virtual or Digital Assets

The Financial Action Task Force ("FATF") defines a virtual asset as a digital representation of value that can be digitally traded or transferred and used for payment and investment purposes. This definition excludes the digital representation of currencies,

[67] Merritt B Fox, Lawrence R Glosten and Gabriel V Rauterberg The New Stock Market: Law, Economics, and Policy (Columbia University Press, New York, 2019) at 24–29.

securities, and other financial assets already covered elsewhere by the relevant laws.[68] Notably, cryptocurrencies such as Bitcoin, Litecoin, Ethereum, etc., are virtual or digital currencies wherein transactions are secured using computerised databases or a well-defined and decentralised network or system. This pool of technology may include cryptography, public key, private key, blockchain, or mathematical function. A blockchain network is a public financial transaction database to track orders, payments, accounts, and production.[69]

c) Trading Venues

Trading venues are primarily digital marketplaces where traders buy and sell, but they are not limited to shares, virtual assets, or digital currencies. The US Exchange Act, for example, defines an exchange, commonly classified as a Self-Regulatory Organisation ("SRO"), as "*any rule-based marketplace for transacting securities that register as an exchange."* [70] Fox, Glosten, and Rauterberg (2020) classify modern trading venues into three categories, namely, stock exchanges, alternative trading systems ("ATS"), and non-ATS off-exchange trade, which is mainly internationalisation.[71] Any trading transaction not conducted on national stock exchanges, that is, off-exchange trade, is done on ATS and non-ATS.[72]

ATS is a non-self-regulatory exchange-like trading venue that facilitates stock trading between buyers and sellers, typically operated by broker-dealers.[73] ATS can take the form of electronic communication networks ("ECNs") and "dark pools". On the other hand, non-ATS is an internalised practice where a retail broker-dealer executes a customer's order on its own books rather than sending it to an exchange or another trading venue. The national stock exchanges remain among the most vital trading venues. SROs are quasi-public institutions delegated with regulatory functions, partly to regulate, establish rules and governance frameworks, and supervise trading activities of market participants within their respective exchanges. For instance, the NASDAQ, NYSE, LSE, NZX, ASX, and other well-known stock exchanges.

---

[68] Financial Action Task Force Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers (FATF, Paris, October 2021) at [13]–[16].
[69] Himani Sardar and Chavi Sood "Cryptocurrency Regulation: Need of the Hour" (2024) 2(1) GLS KALP: Journal of Multidisciplinary Studies 23 at 25–28.
[70] Securities Exchange Act of 1934, 15 USC § 78c(a)(1).
[71] Fox, Glosten and Rauterberg, above n 67, at 29–32.
[72] Fox, Glosten and Rauterberg, above n 67, at 33–37.
[73] Fox, Glosten and Rauterberg, above n 67, at 38–41.

d) Market Participants

The New Zealand Exchange ecosystem comprises diverse market participants who play crucial roles in facilitating securities trading. A Market Participant is a business accredited by NZX that participates in NZX markets[74] operating within the Cash Market, Derivatives Market, Fonterra Shareholders Market, and Primary Market. NZX also accredits financial advisors within these businesses as NZX Advisors.

At the core lie Traders and Trading Firms, who buy and sell securities on behalf of clients or for their accounts. Brokers act as intermediaries, connecting these Traders with the NZX market and executing their orders. Investors, the ultimate source of capital, participate either indirectly through investment vehicles managed by Traders or directly through Direct Market Access ("DMA") accounts, allowing them to submit orders electronically to the NZX. The NZX serves as the central platform, matching buy and sell orders, ensuring efficient price discovery, and facilitating secure settlement of transactions. This interplay between these participants fosters a dynamic marketplace that channels investment capital into the New Zealand economy.

e) Market Supervisors

As briefly discussed in Chapter 1, the FMA is the primary supervisor responsible for endorsing fairness, efficiency, and transparency within the financial markets. As part of its obligations, the FMA also supervises the NZX to ensure it complies with its duties as the only licensed market operator and frontline regulator vested with the power to monitor and regulate the trading of stocks, shares, bonds, and derivatives on the NZX markets. As part of the NZX, the NZX Regulation team ("NZ RegCo") is the surveillance and regulatory agency responsible for monitoring algorithmic trading by conducting investigations and referring likely market manipulation cases to the FMA where there is reasonable suspicion.

[74] New Zealand Exchange "Market Participants" <https://www.nzx.com/services/market-participants>.

2.2.2 Defining Algo-Trading (including HFT)

Like many other jurisdictions worldwide, New Zealand laws have not yet assigned a definitive meaning to algo-trading. As in any field of study, a definition helps to understand and identify the parameters of the subject area. Given the complex and multifaceted nature of algo-trading, Azzutti warned that without a precise definition, there could be regulatory difficulties to address the evolving market manipulation and systemic and black-box risks.[75] Opting for a concrete definition contributes to a more effective assessment of the issues at hand and, thus, improves propositions. Moreover, a well-defined framework facilitates better governance and accountability, enabling regulators to implement measures to mitigate potential abuses and enhance market integrity and investor protection.[76]

Examining the perspectives of regulators or organisations worldwide that have attempted to define algo-trading would provide important insights. To begin with, the Technical Committee of the International Organisation of Securities Commission ("IOSCO") has referred to algo-trading "as using a basic algorithm, set of rules, or computerised instructions to feed portions of an order into the market at pre-set intervals to minimise the market impact costs".[77] While algorithms are the heart of algo-trading, the term "basic" is not justified as algo-trading can also involve complex and sophisticated algorithms far beyond the "basic" set of rules. Moreover, it captures a specific trading strategy where orders are entered at pre-set intervals. Sophisticated trading algorithms might have the ability to react to real-time data. Minimising market impact costs is not the sole or the primary goal of algo-trading. For these reasons, this definition is too narrow and fails to capture the diverse range of algo-trading activities.

In a staff report, the US Securities and Exchange Commission ("SEC") highlighted that since algo-trading is about using algorithms in trading, it is crucial to start by defining the term 'algorithm'. Accordingly, it describes an algorithm as a "finite, deterministic, and effective

---

[75] Alessio Azzutti "AI Governance in Algorithmic Trading: Some Regulatory Insights from the EU AI Act" [2024] SSRN Electronic Journal at 4–8.
[76] Joseph Lee and Lukas Schu "Algorithmic Trading Regulation: The Frameworks for Human Supervision and Direct Market Interventions" [2021] SSRN Electronic Journal at 6–10.
[77] International Organization of Securities Commissions Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency (Final Report, October 2011) at 10.

problem-solving method suitable for implementation as a computer program."[78] Furthermore, the report explains that trading algorithms analyse a vast collection of market information to assess the market, decide when, where and how to trade and execute the resulting trading decisions in the marketplace.[79] The SEC's interpretation of algo-trading appears to be broader than the IOSCO definition. However, the description remains primarily functional and does not itself provide a comprehensive regulatory definition distinguishing different forms or levels of algorithmic automation.

Another study by Gomber, Arndt and Lutat defines algo-trading as "a tool for professional traders that may observe market parameters or other information in real-time and automatically generates trading decisions without human intervention."[80] While this definition highlights the professional application of algo-trading, it is somewhat restrictive. Algorithmic trading is not solely the domain of professional traders; retail investors also utilise these tools to enhance their trading strategies. Furthermore, using the term "may" in the definition could be seen as weakening its precision. Observing market parameters is a fundamental aspect of algo-trading, and thus, the definition should reflect this core function more assertively.

The Hong Kong Securities and Futures Commission's Code of Conduct recognises algo-trading as a form of electronic trading, computer-generated trading activities created by a pre-determined set of rules to deliver specific execution outcomes. This definition of algo-trading covers a pre-determined set of rules as being a key aspect of algo-trading. However, it does not capture the complexity or sophistication of these rules. Algo-trading can range from simple rule-based strategies to complex models where AI and ML are used. Moreover, it fails to capture an essential characteristic of algo-trading: speed, which is the potential for rapid execution and high trading volumes.[81]

In August 2022, the Australian Securities and Investments Commission ("ASIC") issued a Regulatory Guide ("RG241") on Electronic Trading, which defines "algorithmic

---

[78] United States Securities and Exchange Commission *Staff Report on Algorithmic Trading in US Capital Markets* (August 2020) at 5.
[79] SEC, above n 78, at 6.
[80] Peter Gomber and others "High-Frequency Trading" [2011] SSRN Electronic Journal at 14 <https://doi.org/10.2139/ssrn.1858626>.
[81] Securities and Futures Commission *Code of Conduct for Persons Licensed by or Registered with the Securities and Futures Commission* (March 2023) at 50.

programs" as "automated strategies which use programmable logic, system-generated (rather than human-generated) messages based on a set of predetermined parameters, logic rules and conditions". In the same guide, under the section "Key Terms", ASIC defines "algorithm or algorithmic trading" as "electronic trading activity where pre-determined parameters, rules and conditions deliver specific outcomes".[82] It further specifies that such programmes may include algorithmic trading.[83] These definitions identify the essential features of algorithmic trading - namely, automation and rule-based execution.

However, they do not distinguish between basic rule-based systems and more complex implementations involving machine learning. They do not explicitly reference execution speed, often a defining characteristic in practice. While such omissions may not undermine the definitional adequacy per se, they suggest a level of generality that may limit the definitions' utility in contexts requiring finer conceptual distinctions. While the term "pre-determined" indicates rule-based operation, the definition does not expressly distinguish conventional automated execution from more adaptive forms of algorithmic trading.

The European Union ("EU") defines algo-trading in the Markets in Financial Instruments Directive II ("MiFID II") as "trading in financial instruments where a computer algorithm automatically determines individual parameters of orders such as whether to initiate the order, the timing, price, or quantity of the order or how to manage the order after its submission, with limited or no human intervention and does not include any system that is only used to route orders to one or more trading venues or for the processing of orders involving no determination of any trading parameters or for the confirmation of orders or the post-trade processing of executed transactions".[84]

While the EU provides a structured definition of algo-trading, it notably omits a clear explanation of the foundational term "algorithm". Ebers highlights that algorithms are not novel concepts; they have existed long before the advent of computers and are integral to all modern software systems. Drawing on Kitchin's work, he defines an algorithm as a "series of

---

[82] Australian Securities and Investments Commission *Regulatory Guide 241: Electronic Trading* (August 2022) at [241.19].
[83] ASIC, above n 82, at 7.
[84] Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments [2014] OJ L173/349, art 4(1)(39).

well-structured steps to process data or instructions to generate an output".[85] Although this technical definition helps understand how algorithms function, it may fall short in regulatory contexts where broader concerns such as accountability, transparency, and ethical responsibility are at stake. In the context of algorithmic trading, where automated decisions can have significant financial and societal impacts, a more nuanced understanding of what constitutes an algorithm is essential. Clear and context-sensitive definitions are, therefore, pivotal in developing effective regulatory strategies that can address not only the mechanics of algorithmic systems but also their implications for governance and oversight.[86]

Furthermore, the EU definition emphasises algo-trading as an automated and computer-based decision process for executing and managing orders. Mandes explains that algorithmic execution may involve different approaches to the timing and division of orders: the execution of an order at once entirely using a market order with a high trading cost and splitting and scheduling the order at a constant execution rate over the entire trading period.[87] The definition of algo-trading by the EU also reflects the approach adopted by the Federal Financial Supervisory Authority (BaFin), the regulatory authority of Germany under its Securities Trading Act (*Wertpapierhandelsgesetz*). Section 80 of this Act defines algo-trading "as a computer program that automatically determines individual order parameters such as price, time for buying or selling, or quantity of order".[88] Likewise, the UK's Financial Conduct Authority ("FCA") has adopted the EU's definition of algo-trading in the Glossary of the FCA Handbook, which is a set of rules and guidance applicable to all regulated firms under the UK Financial Services and Markets Act 2000.

Turning to the New Zealand context, the 2018 NZX Thematic Review assessed the market and compliance trends focused on algo-trading due to its increasing popularity. Accordingly, the review defines algo-trading as the use of algorithms to autonomously mimic a detailed series of human dealer steps to execute a trading order, and this may involve limited or no human intervention. Furthermore, the review identifies algorithms as computer

---

[85] Rob Kitchin "Thinking critically about and researching algorithms" (2016) 20(1) *Information, Communication & Society* 14 at 17–18; Martin Ebers "Regulating AI and Robotics: Ethical and Legal Challenges" in Martin Ebers and Susana Navas (eds) *Algorithms and Law* (Cambridge University Press, Cambridge, 2020) 37 at 41–43.

[86] Kitchin, above n 85, at 17–18.

[87] Alexandru Mandes "Algorithmic and High Frequency Trading Strategies: A Literature Review" (2016) MAGKS Joint Discussion Paper Series No 2 at 3.

[88] Wertpapierhandelsgesetz (German Securities Trading Act) 26 July 1994 (Germany), s 80(2).

programmes which involve tailormade or stepwise processes and techniques to solve a particular problem or perform a job.[89] The definition of algo-trading, as covered in the NZX review, is problematic as it implies that a trading algorithm simply replicates what a human dealer does. This is not the case for all, or even most, algorithmic trading strategies. While some algorithms might be designed to mimic human behaviours, many can surpass human capabilities in many ways, particularly with real-time trading and monitoring. Moreover, algo-trading is not just executing a trading order. Algo-trading can adjust orders based on market conditions and engage in complex portfolio management strategies. This limitation is important to the thesis because a definition centred on the replication of human dealer behaviour may not adequately capture the autonomy and adaptive characteristics of AI-driven trading considered in Chapters 3 to 7.

Compared to the other definitions discussed previously, the EU definition of algo-trading is a good starting point.[90] It offers a more comprehensive and nuanced approach and captures the key features of algo-trading, including automation, decision-making, order management and limited human intervention. However, this definition would benefit from being improved by (i) avoiding the use of the term "algorithm", which does not help explain algorithmic trading, and (ii) being more explicit on speed. These proposed amendments are particularly beneficial for regulatory and practical purposes as they encompass the essential aspects of algorithmic trading that are critical for understanding its impact on financial markets.

Firstly, the term "algorithm" may not effectively convey the operational mechanics of algorithmic trading to a broader audience. Instead, focusing on the automated nature of trading processes and decision-making can provide a clearer understanding. Secondly, the definition should explicitly emphasise the importance of speed in algorithmic trading. Speed is a critical factor that distinguishes algorithmic trading from traditional trading methods, as it enables the rapid execution of trades and the exploitation of fleeting market opportunities. By highlighting these aspects, the definition can better serve academic and practical regulatory purposes, offering a more precise and accessible explanation of algorithmic trading.

---

[89] NZX Regulation "NZX Thematic Review 2018" Algorithmic Trading (2019) at 5.
[90] Directive 2014/65/EU, above n 84, art 4(1)(39).

In this respect, the EU definition can be further refined for clarity and precision as follows: Trading in financial instruments where a computer programme automatically and often at high speeds, determines individual parameters of orders, such as whether to initiate the order, the timing, price, or quantity of the order or how to manage the order after its submission, with limited or no human intervention and does not include any system that is only used to route orders to one or more trading venues or for the processing of orders involving no determination of any trading parameters or for the confirmation of orders or the post-trade processing of executed transactions. This working definition is adopted for the remainder of the thesis because it captures the features most relevant to the subsequent legal analysis: automated decision-making, speed, limited human intervention and the capacity of computerised systems to determine trading parameters.

The 2018 NZX review also recognises HFT as the most common trading strategy in New Zealand.[91] However, it does not define HFT. Spindler and the Technical Committee of IOSCO describe HFT as the automated trading of shares and securities on stock markets, which operates at a much faster execution speed than other algo-trading strategies.[92] This definition is a decent starting point. However, it misses a key characteristic of HFT: the use of sophisticated technology, including co-location and direct market access. Fox, Glosten and Rauterberg explain that HFT uses fast-speed communications to continuously update information concerning the transactions occurring on each stock it trades, including details about changes in the buy and sell limit orders posted by other traders on every trading venue. In turn, the HFT automatically sends this information to a computer that uses algorithms to change the limit prices and quantities of the stocks as they trade in each trading venue.[93]

Following the EU's proposals to impose additional regulations on high-frequency traders, Germany enacted its High-Frequency Trading Act in May 2013.[94] It refers to high-frequency trading as the buying and selling of financial instruments on one's own account, either as a direct or indirect participant in an organised domestic market or multi-lateral trading

---

[91] NZX, above n 89, at 6–8.
[92] Gerald Spindler "Control of Algorithms in Financial Markets: The Example of High-Frequency Trading" in Martin Ebers and Susana Navas (eds) Algorithms and Law (Cambridge University Press, Cambridge, 2020) 207 at 210–214.
[93] Charles R Korsmo "High-Frequency Trading: A Regulatory Strategy" (2014) 48 University of Richmond Law Review 523 at 540–548.
[94] Martin Haferkorn and Kai Zimmermann "The German High-Frequency Trading Act" [2014] SSRN Electronic Journal at 4.

facility. This trading strategy uses infrastructures designed to minimise latency. The decision-making process regarding the initiation, entry, transmission, or execution of an order occurs without human intervention in individual transactions or orders. Additionally, this strategy is marked by a high intraday message rate, including orders, quotes, or cancellations, when not provided as a service to others."[95]

The US staff report by the Securities and Exchange Commission ("SEC") on Algorithmic Trading Capital Markets commented that there is no statutory or regulatory definition of the term 'HFT' in the US. Nevertheless, it describes HFT as short-term strategies with the ability to place large orders and generate multiple trades daily. The report further refers to the 2010 Concept Release to highlight the characteristics of HFT, which include the use of high-speed and sophisticated technology infrastructure (computer programs, co-location services and individual data feeds offered by exchanges) to quickly and accurately process massive volumes of data, implement trading decisions based on this data and promptly submit numerous orders, which can be cancelled shortly after submission.[96]

The Australian Securities and Investment Commission ("ASIC"), in its 2014 Market Supervision report, emphasises that HFT is not a technical term. The definitions can vary because the term HFT can include a range of trading whereby numerous buy and sell orders are generated by high-speed computer programs on markets such as ASX, ASX 24 and Chi-X. Nevertheless, it also recognises that though HFT is often twinned with algorithmic trading, not all forms of algo-trading can necessarily be considered high-frequency.[97] However, in 2018, the ASIC re-iterated IOSCO's definition of HFT as a form of automated proprietary trading which uses high-speed algorithmic technology to pursue short-term, low-risk profits.[98] Accordingly, the ASIC based itself on IOSCO's defining principles to identify high-frequency traders.

---

[95] Haferkorn and Zimmermann, above n 94, at 4.
[96] United States Securities and Exchange Commission *Staff Report on Algorithmic Trading in US Capital Markets* (August 2020) at 6; *Concept Release on Equity Market Structure* 75 Fed Reg 3594 (21 January 2010) at 3606–3608.
[97] Australian Securities and Investments Commission *ASIC Supervision of Markets and Participants: July to December 2014* (Report 425, 2015) at 8–10.
[98] Australian Securities and Investments Commission *High-Frequency Trading in Australian Equities and the Australian–US Dollar Cross Rate* (Report 597, November 2018) at 5–9.

The European Securities and Markets Authority ("ESMA") does not have a single, concise definition of HFT. Instead, it relies on Article 4(1)(40) of the MiFID II to describe HFT as an algo-trading technique that can be identified through a combination of factors as follows: "(a) infrastructure intended to minimise network and other types of latencies, including at least one of the following facilities for algorithmic order entry: co-location, proximity hosting or high-speed direct electronic access; (b) system-determination of order initiation, generation, routing or execution without human intervention for individual trades or orders; and (c) high message intraday rates which constitute orders, quotes or cancellations."[99]

Considering the above, it becomes clear that a concise, universally agreed-upon definition is elusive. The complexity and rapidly evolving nature of HFT make it more practical to describe its salient characteristics. Therefore, this thesis adopts the approach of defining HFT through characteristics described by the German HFT Act, ASIC, SEC and ESMA. The key elements are infrastructure designed to minimise latency (including co-location, proximity hosting, or high-speed direct electronic access), system-determined order initiation, generation, routing, or execution without human intervention for individual trades, and high intraday message rates comprising orders, quotes, or cancellations. These combined features offer a more comprehensive and adaptable way to understand HFT.

2.2.3 Defining AI (including ML)

As discussed in previous paragraphs, algo-trading does not just rely on pre-determined parameters and rules. As algo-trading evolves, it leverages advanced technologies such as AI to analyse vast datasets, identify patterns, and make trading decisions with a level of sophistication that surpasses traditional algorithmic approaches. This evolution will be detailed in the next section. Therefore, understanding the core concepts of AI, including ML, is crucial. Without this understanding, it will be difficult to fully understand Chapter 3, which discusses AI-powered trading algorithms' capabilities, risks, and regulatory challenges.

The Financial Stability Board ("FSB 2017") referred to AI as the process of applying cognitive tools to perform duties traditionally involving human sophistication.[100] Turner

[99] European Securities and Markets Authority Economic Report: *High-Frequency Trading Activity in EU Equity Markets* (2014) at 5.
[100] Financial Stability Board *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications* (2017) at 3.

defines AI as the "ability of a non-natural entity to make choices by an evaluative process". He explains that he used the term 'non-natural' rather than 'man-made' to refer to the "artificial" part of the definition due to the emergent feature of AI to design and create other AI. Secondly, the definition refers to making choices due to the ability of AI to act independently, that is, to self-govern and make autonomous decisions. The final aspect of the definition is based on an evaluative process as the AI operates by evaluating principles from information or data fed against each other before concluding.[101]

The Artificial Intelligence Researchers Association in New Zealand have prepared a whitepaper outlining its strategic approach to AI development and governance. In this document, the Association references the definition of AI provided by the European Commission ("the Commission") in its whitepaper "On Artificial Intelligence – a European approach to excellence and trust" as follows: "a collection of technologies that combine data, algorithms and computing power".[102] Building on this foundation, the AI Researchers Association defines AI systems as "performing tasks that would require intelligence if a human did them". In other words, an AI system can perform tasks that typically require cognitive functions such as reasoning, learning, perception, or decision-making when carried out by humans. This reframing aims to distinguish AI from conventional automation by emphasising the simulation of human-like cognitive processes rather than the mere execution of tasks. The European Commission defined the AI system based on the OECD's definition of the AI system. The Commission emphasised that the definition of an AI system must be flexible and not technology-specific to accommodate future technological developments. Therefore, the Commission complemented the definition of an AI system with a list of the applicable latest techniques and approaches used for its development.[103]

The European Commission's 2021 proposal for the Artificial Intelligence Act defined an AI system as: "a software developed with one or more of the techniques and approaches listed in Annex I. For a given set of human-defined objectives, the AI system can generate outputs such as content, predictions, recommendations, or decisions influencing the environments they interact with." Annex I includes supervised and unsupervised machine

---

[101] Jacob Turner *Robot Rules: Regulating Artificial Intelligence* (Palgrave Macmillan, UK, 2019) at 16-17.
[102] European Commission, *White Paper on Artificial Intelligence: A European approach to Excellence and Trust. European Commission* (February 2020).
[103] European Commission *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)* COM(2021) 206 final, art 3.

learning techniques such as deep learning, logic and knowledge-based and statistical approaches. Article 4 empowers the European Commission to review, update or complement this list by adopting delegated acts under Article 73.[104]

Machine learning, within the realm of AI, refers to giving a machine the ability "to learn whenever it changes its structure, program or data (based on its inputs or in response to external information) in such a manner that its expected future performance improves."[105] Deep Learning ("DL") is a subset of ML which uses artificial neural networks with feature learning.[106] Supervised machine learning algorithms are the most used. With this model, a data scientist guides and teaches the algorithm what conclusions it should make. On the other hand, unsupervised machine learning uses a more independent approach in which a computer learns to identify complex processes and patterns without a human providing close, constant guidance. Unsupervised machine learning involves training based on data that does not have labels or a specific, defined output.

2.3. MODERN TRADING: FROM FLOOR TO ALGORITHM TO AI

This section establishes the historical context for trading. It showcases the gradual automation of market functions, starting with human intermediaries and culminating in algorithmic trading, a significant iteration in the continuous technological evolution. This part is considered crucial for establishing the inevitability of technological change in financial markets, making the case that regulation must be equally dynamic. This historical perspective shows the evolution from human-driven trading to modern algo-trading, which allows the thesis to argue later that existing regulations designed for a human-centric market are ill-equipped to a certain extent to address emerging challenges.

The stock exchange market originated in the early modern period when European countries initiated collaboration by establishing joint-stock companies to conduct financial trades and explore capital ventures.[107] Unlike contemporary brokers, traders and merchants

---

[104] European Commission *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)* COM (2021) 206 final, art 3.
[105] Nils J Nilsson *Introduction to Machine Learning* (Stanford University, 1998) at 1.
[106] Artificial Intelligence Researchers Association *Aotearoa New Zealand Artificial Intelligence: A Strategic Approach* (November 2021) at 3.
[107] Mark Smith *A History of the Global Stock Market: From Ancient Rome to Silicon Valley* (University of Chicago Press, 2004) at 12.

would gather in public places to trade company shares to the public. Around the mid-16th century and following the internationalisation of share trading across Europe, the first permanent stock exchange market was established in Antwerp City, Belgium, where business financiers traded in promissory notes and bonds. In 1602, the Amsterdam Stock Exchange was the prominent stock exchange for trading shares and stocks.[108] In those times, real 'stock' or tangible share certificates did not exist, nor would the shares' ownership technically change hands upon the trading of shares, as is the case in contemporary stock markets.[109] Nevertheless, these types of markets laid down the infrastructural and institutional foundation to form formal stock exchanges.[110]

In the 17th and 18th centuries, countries like Belgium, Spain, Portugal, Sweden, France, England, and even the US established their respective stock exchanges due to the increasing ease of share trading and shares going public. Around 1571, and in 1693, London founded its stock exchange, 'The London Stock Exchange' ("the LSE").[111] In the late 19th and 20th centuries, the stock exchanges became more proactive and developed. In 1971, the National Association of Securities Dealers Automated Quotations ("NASDAQ") was formed, a first-ever direct competitor of the NYSE nearly two centuries later. NASDAQ was among the first innovators to shift from traditional stock trading to computerised and electronic trading.[112]

The technological innovation led to the merger of NYSE-Euronext, the world's first transatlantic stock exchange. These developments substantially reduced the role of traditional floor trading and enabled investors to participate in increasingly electronic markets. This further triggered the development of other electronic trading venues known as "electronic communications networks". Eventually, algorithms to facilitate the electronic entry, execution of orders and better average prices are referred to as 'algo-trading'. Algo-trading has changed how we trade, dictating how liquidity, price discovery mechanisms, and modification of market designs now occur in these modern markets at a breakneck pace. Today, market participants can trade securities or digital or virtual assets over multiple trading venues: NASDAQ, NYSE, LSE, NZX, and ASX. The modern market can potentially enable trade over seventy-five

---

[108] Rjumohan *Stock Markets: An Overview and a Literature Review* (Munich Personal RePEc Archive Paper No 101855, 2019) at 6–10.
[109] Smith, above n 107, at 12.
[110] Rjumohan, above n 108, at 11–14.
[111] London Stock Exchange < https://www.londonstockexchange.com/discover/our-regions/united-kingdom >.
[112] Rjumohan, above n 108, at 19.

competing trading venues consisting of twelve exchanges and over thirty dark pools trading platforms.[113]

The New Zealand Exchange Market was first established in Dunedin in 1866, primarily trading shares in the gold mines amid the South Island's gold rush.[114] This setup engineered the regional development of multiple stock exchanges in Auckland (the 1870s), Wellington (1882), and Christchurch (1900), each operating independently. Since then, the New Zealand Exchange has undergone multiple stages of development. For instance, the New Zealand Stock Exchange Association regrouped all the regional exchanges in 1915. In 1983, the New Zealand Stock Exchange was formed, whereby qualified brokers' members gained national status instead of regional status. In 1991, New Zealand kept pace with US technological innovation, implemented computerised share trading, and phased out regional and traditional trading floors. Ultimately, in 2003, the New Zealand Exchange was rebranded to "NZX", and the overall governance and oversight of the exchange market were delegated to its members, with the brokers having the role of market participants.[115]

Per the NZX Thematic Review (2018),[116] algo-trading has accounted for a substantial proportion of trading activity on the New Zealand Stock Exchange by 2018. For instance, HFT and general algorithmic execution increased on-market trades (measured by the number of trades) by 47%, on-market volumes (number of securities traded) by 16.5%, and on-market turnover (dollar value) by 26% as of Q3 2018.[117] These trends are equivalent to algo-trading and HFT in the Australian market. Algo-trading has augmented the value of on-market trading on the NZX to 53.1% for the trading year as of November 2018. It records that the NZX market accounts for 10% of the usage of general execution trading algorithms compared to more than 90% of high-frequency trading algorithms being deployed within the equity securities on the NZX Main Board. The NZX anticipated this trend would continue increasing as market participants gained increased access to algorithmic solutions that are more efficient and cost-effective in meeting clients' demands. HFT dominates NZX, consisting of more than 91.2% of

[113] Fox, Glosten and Rauterberg, above n 67, at 3.
[114] New Zealand's Stock Exchange <www.nzx.com/about-nzx/organization-structure>.
[115] New Zealand Exchange, above n 114.
[116] NZX, above n 114, at 2.
[117] NZX, above n 89, at 13.

the overall on-market algorithmic trades, 94.9% of on-market algorithmic volume, and 95.7% of on-market algorithmic trades.[118]

Table 1 – shows the increasing usage of algorithmic trades between the period 2015 to 2018.

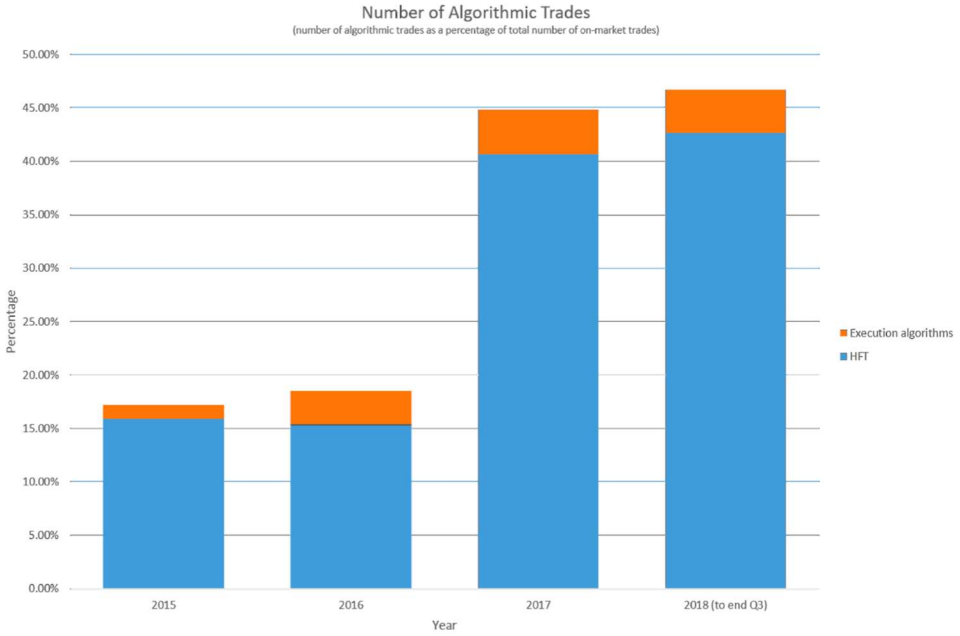
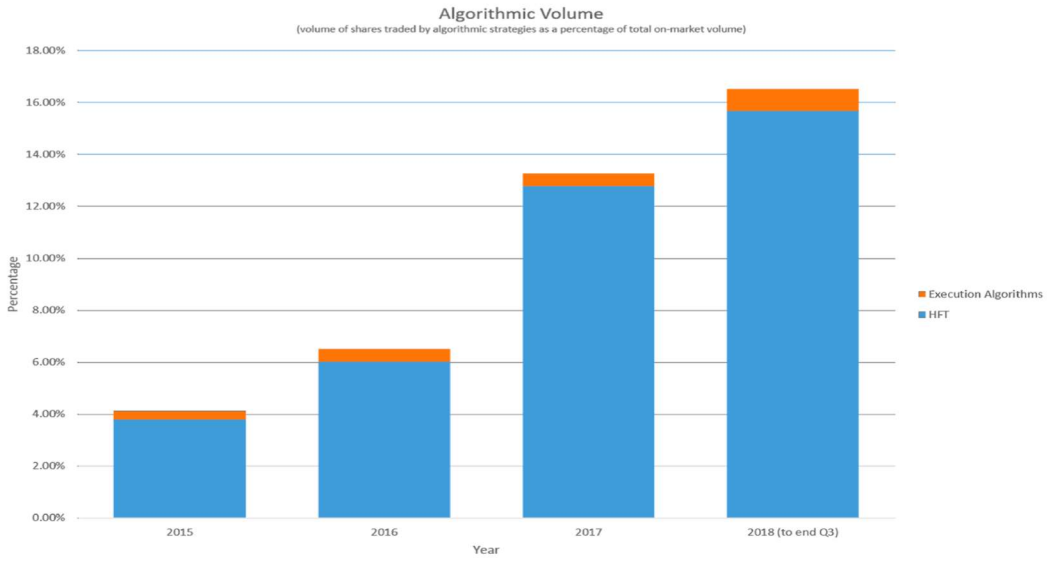
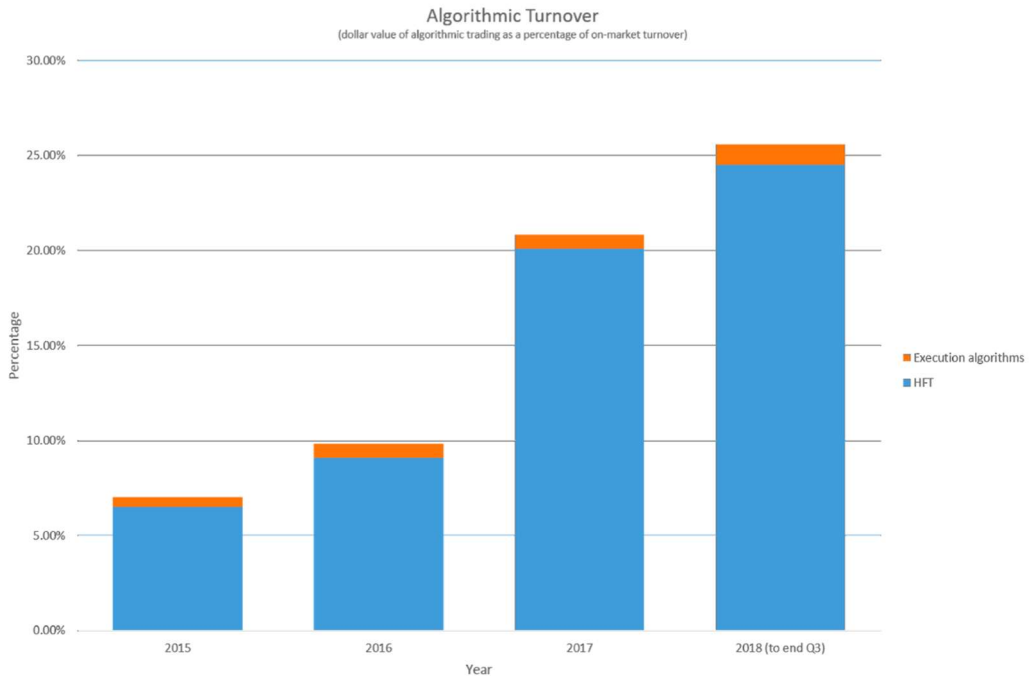


Table 2 – shows the increasing trend of the volume of shares traded by algorithmic trading between the period 2015 to 2017.



[118] NZX, above n 89, at 14.

Table 3 shows the increasing value of algorithmic trading between 2015 and 2018.



The review identified a diverse range of algo-trading and HFT strategies employed within the NZX market. The participate algorithmic strategy (also called a "go-along", "percent-of-volume" (POV) or "volume-in-line" algo) is one of the most common types of trade algorithms which allows the execution of client instructions in line with a set participation ratio. The participation ratio would be in the form of a certain percentage of the day's traded volume. The user could set this ratio as a desired percentage using parameters relevant to their risk appetite. In other words, the user specifies the fraction of the volume traded in the market. As a result, the user could choose between aggressive and passive instructions on an order. These instructions operate differently; for example, the "aggressive instruction" on a sell order will try to bid rather than sit on offers. Conversely, for the same order, a passive instruction will prefer to sit on the offer side of the market as much as possible.

The Volume Weighted Average Price ("VWAP") is another common algorithmic strategy whereby the algorithm tries to buy and sell stocks close to the average price of a specific period, considering the volume traded at each price level. It reviews a particular number of trading days to build an average volume profile for each half-hour. The algorithm then breaks the order into smaller pieces and submits them to the market to match the volume profile. The user can define parameters to indicate how aggressive they want the strategy to be.

Basically, this strategy attempts to minimise tracking errors while maximising performance.[119] On the other hand, the Time Weighted Average Price ("TWAP") enables a user to achieve an average price or better price by executing orders over specific time intervals. This method splits the trading day into several time increments, and like the other strategies, it allows the users to set their preferred parameters of aggressiveness.[120]

The implementation shortfall strategy attempts to minimise costs when buying or selling stocks. It reduces the difference between the price when the decision was made and when the order was executed. It does so by looking at the current market price and using this as a reference point to adjust the strategy of how much to buy and sell before implementing the order. Ultimately, this strategy results in more informed decisions and lower execution costs. In contrast, the Arbitrage strategy has emerged from overseas and has been designed to ensure that High-Frequency Traders receive a rebate for providing liquidity to investors. In other words, this allows them to make profits by taking advantage of price differences between securities traded on multiple exchanges at the expense of other market participants. Several companies listed on the NZX Main Board also have listings on the Australian Exchange ("ASX"), creating opportunities for arbitrage between these two exchanges.[121]

While the 2018 NZX thematic review offers valuable insights into algo-trading activity in New Zealand, publicly available and comprehensive contemporary data on the extent of algorithmic and AI-driven trading on the NZX remains limited. Accordingly, this thesis does not assume that the 2018 figures accurately represent current levels of algorithmic trading. Rather, the review is used as evidence that algorithmic trading had already become a significant feature of the New Zealand market by 2018. This limitation is relevant to the thesis methodology and means that conclusions concerning the present use of AI-driven trading in New Zealand must be approached cautiously.

This evolution naturally leads to the increasing integration of AI within trading strategies. AI forms part of the broader technological transformation commonly associated with the Fourth Industrial Revolution, which advances in big data, computational powers and global interdependence have facilitated. Earlier, only financially robust corporations could leverage

---

[119] NZX, above n 89, at 15.
[120] NZX, above n 89, at 16.
[121] NZX, above n 89, at 17.

its potential. However, the barrier to entry has now substantially lowered, making AI accessible to a much broader range of users. AI is increasingly used in financial markets to analyse large datasets, identify patterns, support trading decisions and automate aspects of portfolio and risk management. These developments illustrate the movement from conventional rule-based algorithmic trading towards systems capable of more adaptive and autonomous decision-making.[122]

2.4. THE APPLICATION OF AI IN MODERN ALGO-TRADING

2.4.1. Forecasting Stock Market Trends

AI predicts stock markets through two main approaches: ML and DL. ML involves training computers to recognise patterns in stock market data without being directly programmed for every possibility. It uses various algorithms such as Support Vector Machines ("SVM"), which work by finding optimal boundaries to classify different stock movements for prediction, and Artificial Neural Networks ("ANN"), which are structured like simplified brains with interconnected nodes that learn from data, and Decision Trees ("DT") which make predictions based on a series of rules derived from the data. These algorithms analyse historical stock prices, trade volume, and technical indicators. Technical indicators are signals generated from price and volume data. Examples include Moving Averages, which smooth out price fluctuations over a specific period, and the Relative Strength Index, which indicates whether a stock is trading at an extreme high or low. The goal is to identify relationships within this data to forecast future price behaviour.[123]

DL is a more advanced form of ML that utilises more complex ANN with multiple layers. These deep networks effectively uncover intricate and non-obvious patterns in enormous datasets. Two notable types of DL networks used in stock prediction are Long Short-Term Memory ("LSTM") networks and Convolutional Neural Networks ("CNN"). LSTM networks are designed to remember information over extended periods, making them suitable for analysing sequences of data like stock prices evolving over time and understanding how past events influence future prices. CNN identifies patterns in complex data with multiple

---

[122] Syed Hasan Jafar, Hani El-Chaarani and Sairam Moturi *Artificial Intelligence for Capital Markets* (CRC Press, 2023) at 2.
[123] Rui Ren, Desheng Dash Wu and Tianxiang Liu "Forecasting Stock Market Movement Direction Using Sentiment Analysis and Support Vector Machine" (2019) 13(1) IEEE Systems Journal 760 at 762–765.

dimensions, allowing for the integration of diverse information, such as stock prices, alongside news sentiment or macroeconomic data.[124]

ML and DL techniques can process large volumes of stock-market data and identify complex relationships that may not be captured by conventional linear models. Their predictive performance, however, depends on matters including the quality of the data, model selection, market conditions and the variables incorporated into the model. Research has therefore increasingly considered ensemble approaches and the integration of alternative data, including financial news, sentiment indicators and macroeconomic information, to improve predictive performance. These techniques nevertheless remain subject to limitations because financial markets are influenced by changing economic, regulatory and behavioural conditions that cannot necessarily be inferred from historical data alone.[125] Importantly, as AI systems become more embedded in trading environments, their influence on market dynamics must also be considered. The presence of AI itself may alter market behaviour and contribute to feedback effects that challenge assumptions based solely on historical market relationships. This feature becomes particularly important in Chapter 3, which considers the risks created by increasingly autonomous and interconnected trading systems.

2.4.2. Leveraging Predictive Analytics in Trading Strategies

AI's role in stock trading is significantly enhanced by its ability to act as a highly sophisticated forecaster, utilising predictive analytics. These AI systems employ ML and DL algorithms to rapidly process and analyse vast amounts of financial data, including historical stock prices, current trading volumes, and various other influential factors. AI can examine financial reports to understand a company's health, analyse market news for real-time developments, and even perform social media sentiment analysis to gauge public opinion. By identifying intricate patterns and correlations within this diverse data, AI aims to accurately forecast future market trends and the potential direction of stock prices.[126]

AI's predictions are crucial in developing and refining trading strategies. Instead of relying on traditional methods or human intuition alone, traders can leverage AI's insights to

[124] Ren, Wu and Liu, above n 123, at 765–767.
[125] Ren, Wu and Liu, above n 123, at 767–770.
[126] Oyewale Oyedokun, Somto Emmanuel Ewim and Oluwaseun Peter Oyeyemi "Leveraging Advanced Financial Analytics for Predictive Risk Management and Strategic Decision-Making in Global Markets" (2024) 2(2) Global Journal of Research in Multidisciplinary Studies 16 at 18–22.

make more informed decisions about when to enter or exit trades. For instance, if AI predicts an upward trend for a specific stock with a high degree of confidence, a trader might initiate a buy order. Conversely, predicting a downward trend could prompt a sell order or even a short position. AI can also help identify optimal entry and exit points based on predicted price movements and volatility. Furthermore, AI contributes to risk management by analysing extensive data to assess potential risks and even predict market anomalies or downturns, allowing for pre-emptive actions to safeguard investments, sometimes using strategies involving Inverse Exchange-Traded Funds to hedge against market corrections.[127]

The power of AI in trading is further amplified through automated trading systems that can automatically execute trades based on pre-defined strategies and the AI's real-time predictions. These trading bots can operate continuously, reacting to market changes much faster than humans and eliminating the influence of emotional biases on trading decisions. AI algorithms can be based on both Technical Analysis ("TA"), which involves analysing price charts and indicators and Fundamental Analysis ("FA"), which consists of evaluating a company's financial health and intrinsic value. Sophisticated AI trading bots can analyse correlations between stocks and industry indices, assess a company's financial stability through proprietary algorithms, and optimize trade execution for better pricing and reduced costs. This automation and analytical depth empower traders to implement complex strategies and potentially achieve greater efficiency and profitability in the dynamic world of stock trading.[128]

2.4.3. Sentiment Analysis

AI's initial role in sentiment analysis for stock trading involves the automated collection of vast amounts of textual data from the internet. Sophisticated web crawlers are deployed to continuously monitor and extract information from a wide array of online sources relevant to the financial markets. These sources encompass formal news outlets, financial reporting platforms, and more dynamic and less structured environments such as social media feeds (like Twitter), investor forums, and blog comments. This data's sheer volume and diversity provide a rich landscape for understanding the multifaceted opinions and emotions that drive investor behaviour.[129]

---

[127] Oyedokun, Ewim and Oyeyemi, above n 126, at 22–25.
[128] Oyedokun, Ewim and Oyeyemi, above n 126, at 25–27.
[129] Yue Qiu, Zhewei Song and Zhensong Chen "Short-Term Stock Trends Prediction Based on Sentiment Analysis and Machine Learning" (2022) 26(5) Soft Computing 2209 at 2211–2215.

Once this textual data is gathered, AI employs a suite of Natural Language Processing ("NLP") techniques to dissect and interpret its content. This process includes breaking down text into manageable units like sentences and words, identifying the semantic meaning of these units, and detecting the presence of sentiment-laden language. AI algorithms often utilise pre-existing sentiment lexicons and ontology bases, which are curated collections of words and phrases tagged with their associated emotional polarity (positive, negative, or neutral) and intensity. By mapping the words and phrases found in the collected text against these lexicons, AI can assign sentiment scores to individual pieces of content.[130]

Aggregating these individual sentiment scores over a specific period, such as a trading day, allows AI to generate a comprehensive sentiment index. This index serves as a quantitative measure of the overall market mood or the prevailing sentiment towards a particular stock or sector. Advanced AI models can also weigh different sources of information based on their perceived reliability or influence. For instance, the sentiment expressed by verified financial analysts might be given more credence than anonymous opinions on social media. Furthermore, AI can analyse the context in which sentiment is expressed, helping to disambiguate sarcasm or nuanced language that simpler sentiment analysis techniques might misinterpret.[131]

Finally, the derived sentiment index becomes a valuable input for AI-driven trading strategies. Integrated into ML models, including Logistic Regression ("LR"), sentiment data is used alongside traditional financial indicators to forecast future stock price movements and market trends. LR is a statistical model that, despite its name, is primarily used for binary classification tasks – in this context, predicting whether a stock price will go up or down. It estimates the probability of a binary outcome, such as a positive or negative price change, based on a set of input variables, such as the sentiment index and other market data. The model outputs a probability score between 0 and 1, which can then be used to make a classification decision based on a chosen threshold. For example, a probability above 0.5 indicates a likely price increase. By incorporating sentiment analysis as an input, LR models in AI trading aim to capture the influence of investor sentiment on these probabilities, potentially leading to more accurate predictions and informed trading decisions.[132]

---

[130] Qiu, Song and Chen, above n 129, at 2215–2218.
[131] Qiu, Song and Chen, above n 129, at 2218–2221.
[132] Qiu, Song and Chen, above n 129, at 2221–2224.

2.4.4. Risk Assessment and Management

AI plays a multifaceted role in risk management within stock trading, leveraging its analytical capabilities to identify, assess, and potentially mitigate various financial risks. One key application lies in portfolio construction and risk assessment, where AI can evaluate "Value-at-Risk (VaR)" and "Expected Returns" for suggested portfolios. They can analyse historical stock data, volatility patterns, and correlations between different assets to construct portfolios that align with an investor's stated risk tolerance. By calculating metrics like VaR, AI can estimate the potential maximum loss over a specific time horizon at a given confidence level, providing investors with a clearer understanding of the downside risks associated with a particular portfolio.[133]

Furthermore, AI contributes to risk management by identifying and analysing key risk factors influencing stock trading, including investor emotions, firm performance measurements, geopolitical ramifications, behavioural economics, and economic indicators. AI is adept at processing vast amounts of data from these diverse sources to identify subtle, non-linear relationships and potential risk indicators that human analysts might overlook. By continuously monitoring these factors, AI can provide early warnings of potential market volatility or increased risk exposure. By using hierarchical clustering to group traders with similar risk appetites and preferences, AI can tailor stock recommendations to individual risk tolerances, ensuring that investors are not presented with opportunities that carry a level of risk they are uncomfortable with. Recommending stocks and portfolios that are consistent with an investor's risk profile is a crucial aspect of risk management, preventing them from making impulsive or unsuitable investment decisions that could lead to significant losses.[134]

In essence, AI enhances risk management in stock trading by providing sophisticated tools for portfolio analysis, identifying complex risk factors, and personalising investment recommendations based on individual risk tolerance. By quantifying potential risks through metrics like VaR and aligning investment strategies with investor profiles, AI empowers

---

[133] Alok Bhardwaj and others "Risk Assessment and Management in Stock Trading Using Artificial Intelligence" in *2024 3rd International Conference on Sentiment Analysis and Deep Learning (ICSADL)* (IEEE, 2024) 138 at 140–143.
[134] Bhardwaj and others, above n 133, at 143–145.

traders and investors to make more informed decisions and navigate the inherent uncertainties of the stock market with greater awareness and control.[135]

2.5. CONCLUSION

In conclusion, this chapter has laid a foundational understanding of the contemporary stock trading landscape, highlighting the critical shift from traditional human-driven practices to the increasing dominance of algorithmic and AI-powered systems. By defining the core components of the stock exchange market and meticulously defining algo-trading and HFT, it established the technological bedrock upon which modern finance now operates. The provided definition of AI by the European Commission, encompassing machine learning, deep learning, and various statistical approaches, underscores the broad and evolving nature of the intelligent technologies now deeply embedded within trading environments.

The integration of AI into modern trading represents a profound evolution that extends far beyond mere automation. The capacity of AI, particularly through ML and DL, to independently learn from vast datasets, adapt to dynamic market conditions, and generate outputs without direct human intervention signifies a paradigm shift in financial operations. This independence allows AI systems to identify subtle patterns and anomalies that might be imperceptible to human traders or traditional algorithmic systems, leading to potentially more accurate forecasts and more effective trading strategies. The inclusion of sentiment analysis as a key application further highlights AI's ability to incorporate qualitative, human-centric data into quantitative trading models, bridging the gap between market psychology and financial decision-making.

The proliferation of AI in diverse markets, as evidenced by international trends and the specific case of the NZX, underscores the universality and adaptability of these technologies. The various algorithmic strategies now deployed, many enhanced by AI, demonstrate the practical implementation of these advanced analytical capabilities. Whether HFT capitalises on minute price discrepancies or more strategic, longer-term approaches informed by predictive analytics, AI is the intelligent engine driving these automated processes. This widespread adoption suggests a growing confidence in the ability of AI to navigate the

[135] Bhardwaj and others, above n 133, at 145.

complexities of financial markets and provide a competitive edge in an increasingly data-rich and fast-paced environment.

Ultimately, the application of AI across forecasting, strategic trading, sentiment analysis and risk management represents a significant development in the operation of modern financial markets. AI is not simply a tool for executing trades; increasingly autonomous and adaptive systems may participate in decision-making throughout the trading process. The characteristics identified in this chapter, particularly speed, autonomy, adaptability, opacity and reduced human intervention, provide the technological foundation for the legal analysis that follows. Chapter 3 examines the risks that these characteristics create, Chapters 4 and 5 assess whether New Zealand's existing market manipulation framework can adequately respond to those risks, and Chapter 6 considers the limits of existing private law remedies. Those findings ultimately inform the governance, liability and regulatory reforms proposed in Chapter 7. In this way, the technological developments examined in this chapter provide the factual basis for the thesis's central question: whether New Zealand's existing legal framework can adequately respond to the distinctive risks created by AI-driven trading.

Chapter 3: AI-Driven Algorithmic Trading and Emerging Risks

Building on the technological account developed in Chapter 2, Chapter Three adopts a risk-based methodology to identify and evaluate the inherent risks of AI-driven algorithmic trading and to establish the regulatory problem that the remainder of the thesis seeks to address. The chapter is presented in two parts. The first part, Section 3.1, examines the risk of harm, beginning with cyber risks such as operational failures and inherent errors arising from AI and machine learning (Section 3.1.1), followed by the potential for malicious use by humans to create manipulative algorithms (Section 3.1.2). Trading misconducts, including collusion and insider trading, are then analysed (Sections 3.1.3.A and B) before the discussion progressively narrows toward market manipulation (Section 3.1.3.C), which constitutes the central analytical focus of this thesis.

The second part, Section 3.2, examines the liability conundrum generated by AI's autonomous "black box" operation, where regulatory insight is limited and the attribution of responsibility for AI-driven market manipulation becomes doctrinally complex. To contextualise these risks, the chapter draws upon historical market events in New Zealand and major international jurisdictions, highlighting behaviours that undermine market efficiency and public trust. The chapter culminates in Section 3.3, which integrates the identified risks and the associated regulatory and legal challenges. Although a range of risks are surveyed, the chapter ultimately focuses on market manipulation, and the analysis developed here directly informs the assessment of New Zealand's market manipulation regime in Chapter 4, the case-law analysis in Chapter 5, the doctrinal limits of private law examined in Chapter 6, and the integrated regulatory framework, including the proposed AI Governance Liability tort and bespoke product liability regime, developed in Chapter 7.

3.1. THE RISK OF HARM: ALGORITHMIC TRADING AND POTENTIAL MARKET DETRIMENT

3.1.1 Cyber Risks: Risks of Operational Failure & Inherent Errors

"*Artificial Intelligence ("AI") is an infant at best. Once it becomes a teenager and believes it is smarter than its parents, will AI rebel?" – Dave Waters.*

Advances in AI and machine learning are changing algorithmic trading in financial markets. Speed is a defining characteristic of algorithmic trading. Algorithms can react to market information and submit or cancel orders faster than human operators can intervene. This speed can magnify the consequences of erroneous instructions, coding defects or unexpected interactions between trading systems.[136] Complexity can also make it difficult for developers, operators and regulators to identify the source of an error before it affects trading or market liquidity.[137]

Moreover, algo-trading operates within a complex web of interconnectedness. Even smaller players and instruments can pose significant threats, regardless of their individual size or value, including the risk of a domino effect, where a single misstep by one participant can trigger a cascade of failures, creating vicious cycles of volatility that ripple throughout the entire system. The deep interconnectedness of the system also allows price fluctuations to propagate rapidly across multiple exchanges and asset types, potentially causing widespread financial disruption. This concern is echoed by the Financial Stability Board ("FSB"), a global body responsible for monitoring financial stability. The FSB, particularly in its first report on AI in financial services, expresses anxieties about the unforeseen connections that might arise due to the use of AI and ML. It highlights the potential risk associated with various institutions utilising previously unrelated data sources, potentially leading to unforeseen and unpredictable interdependencies within the financial network.[138]

While algorithmic trading offers undeniable advantages, it can also introduce potential risks for market quality. The growing automation, speed, interconnectedness, interdependence,

---

[136] Tom CW Lin 'The new market manipulation' (2017) 66 *Emory Law Journal* 1253 at 1275–1281.
[137] Lin, above n 136, at 1283–1286.
[138] Financial Stability Board *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications* (FSB, 1 November 2017) at 4–6.

and removal of intermediaries that make algo-trading attractive might unintentionally undermine traders' role in ensuring efficient and equitable market function. Algorithmic trading can make the trading process more susceptible to glitches and misinterpretations originating from the AI machines themselves. This vulnerability can create operational risks where malfunctions can disrupt trades and inherent errors programmed into the algorithms, both of which can be categorised under the umbrella of cyber risk.[139]

Operational risk arises because algorithmic trading depends on software, data, communications infrastructure and automated execution. Failures may result from coding errors, defective configuration, inaccurate inputs, connectivity problems or inadequate testing and monitoring.[140] The speed of automated execution can allow an error to generate substantial order flow before human intervention occurs. These risks do not arise in every algorithmic strategy, but they require effective pre-deployment testing, real-time monitoring and operational controls.[141]

Algorithmic trading introduces additional cyber risks, especially when dealing with complex markets that generate vast amounts of intricate data. These algorithms, particularly those designed for investment decisions, often struggle to handle qualitative and contextual information. This creates a bias where different algorithms, programmed by various traders, respond similarly to the same data points. As a result, algo-trading can rapidly amplify market trends. While this can be beneficial when the algorithms correctly interpret the data, it can be disastrous if they all miss or misinterpret crucial information. These errors can significantly impact price accuracy, especially for simpler algorithms that rely on limited data and are less adaptable to changing information.[142]

The negative consequences of malfunctioning algorithms extend beyond the trader who deployed them. These errors can strain the infrastructure of trading venues, reduce available liquidity, and quickly spread to other market participants, potentially impacting multiple markets. This is because even though modern markets, especially those in the EU, are

---

[139] Yesha Yadav "How Algorithmic Trading Undermines Efficiency in Capital Markets" (2015) 68 Vanderbilt Law Review 1607 at 1639–1645.
[140] Yadav, above n 139, at 1645–1650.
[141] Yadav, above n 139, at 1652–1656.
[142] European Securities and Markets Authority *Guidelines on Systems and Controls in a Highly Automated Trading Environment for Trading Platforms, Investment Firms and Competent Authorities* (Consultation Paper, ESMA/2011/224, July 2011) at [44] – [52].

fragmented, they are also highly interconnected through the activity of algo-traders. This interconnectedness, combined with the speed of algorithmic trading, amplifies the impact of misbehaving algorithms. It can fuel instability across different markets and create self-reinforcing trends that can morph into price bubbles.[143]

Algorithmic trading, including high-frequency trading, has been associated with episodes of acute market volatility. The 6 May 2010 Flash Crash illustrates how automated trading can contribute to market instability under stressed market conditions. During the event, prices of numerous US securities declined sharply within minutes before substantially recovering. The joint CFTC and SEC investigation identified a large automated sell programme, declining market liquidity and interactions among high-frequency traders as contributing factors. It did not conclude that the event was triggered by a malfunctioning algorithm. Rather, the event demonstrates how automated execution and interactions between algorithmic trading strategies can rapidly amplify market movements and reduce liquidity.[144]

Knight Capital's 2012 software failure provides a separate example of operational risk arising from defective deployment and controls in an automated trading system. It should be treated as an operational failure involving conventional trading software, rather than as evidence of autonomous AI behaviour. For smaller exchanges, even a minor internet outage can be disastrous. Since trading operates in a relentless environment, a robust technological infrastructure is critical for maintaining operational stability.[145]

Traditional financial models may seek to identify causal relationships between variables. In contrast, many ML methods are data-driven and identify patterns and statistical correlations within datasets. This shift from causation to correlation represents a potential paradigm shift in financial theory. While ML can provide fast and cost-effective data analysis, relying solely on correlations without considering causation could lead to errors, biased outcomes, and unintended consequences in financial decision-making. This shift towards ML-

[143] Steffen Kern and Giuseppe Loiacono "High-Frequency Trading and Circuit Breakers in the EU" in Walter Mattli (ed) *Global Algorithmic Capital Markets: High-Frequency Trading, Dark Pools, and Regulatory Challenges* (Oxford University Press, Oxford, 2018) 308 at 317–321.
[144] Kirilenko, Kyle, Samadi and Tuzun, above n 33, at 970–976.
[145] Azzutti, Ringe and Stiehl, above n 12, at 116.

driven algorithmic trading could trigger a significant overhaul of financial theory and the statistical tools used to navigate uncertainty in financial decision-making.[146]

3.1.2. Malicious Use by Humans

The growing adoption of AI in algorithmic trading has ignited concerns regarding its potential impact on market stability. However, not all algorithmic market disruptions are accidental. Sometimes, the algorithms are deliberately designed and used for illegal purposes, constituting "by-design" algorithmic market abuse, which introduces a new layer of complexity. These manipulative algorithms can be programmed from the start or even trained to exploit market weaknesses while still appearing profitable. For instance, human experts can train AI traders through historical data or simulations to "discover" manipulative strategies that maximise profits.[147]

One significant challenge with algorithmic decision-making in finance is the potential for bias. Notably, programmers can employ biased data or information to train the trading models based on race or gender, raising ethical concerns and exposing users to potential discrimination. A hypothetical example would be an AI trading algorithm trained on historical loan data reflecting past societal biases that could falsely identify race or gender as indicators of loan default risk, even if these correlations are spurious. As a result, when assessing creditworthiness for trading on margin or other financial products, the algorithm might unfairly disadvantage individuals from specific racial or gender groups with less favourable terms based on flawed historical patterns rather than their actual financial reliability, thereby perpetuating discriminatory outcomes in financial markets.[148]

A further risk arises where algorithms are deliberately used to create false or misleading market signals. The case of Navinder Sarao illustrates deliberate human use of automated trading techniques for this purpose. Sarao used manual and automated spoofing strategies, including placing and cancelling large orders, to create a false impression of supply in the E-mini S&P 500 futures market. His conduct demonstrates how automated trading tools can be used by human actors to implement manipulative strategies at speed and scale. However, the

---

[146] Azzutti, Ringe and Stiehl, above n 12, at 116–117.
[147] Azzutti, Ringe and Stiehl, above n 12, at 117–118.
[148] Federico Cecconi (ed) *AI in the Financial Markets: New Algorithms and Solutions* (Springer, Cham, 2023) at 92–95.

Sarao case did not involve machine learning, corrupted training data or an AI system autonomously learning to manipulate the market. It should therefore be treated as an example of deliberate human misuse of algorithmic trading rather than AI-driven manipulation.[149]

The 2014 SEC proceeding against Athena Capital Research provides another example of deliberate human use of algorithmic high-frequency trading. Athena used a sophisticated algorithm known as "Gravy" to engage in a strategy designed to affect the closing prices of NASDAQ-listed securities, commonly described as "marking the close". The algorithm placed large numbers of orders during the final seconds of trading and enabled Athena to dominate the available liquidity in certain securities, affecting their closing prices.[150] The case illustrates how conventional algorithmic trading systems can be deliberately designed and deployed to implement manipulative strategies at high speed. It should not, however, be characterised as an example of autonomous AI or machine-learning misconduct. Athena settled the SEC administrative proceeding without admitting or denying the findings, except as to the SEC's jurisdiction, and agreed to pay a civil monetary penalty of USD 1 million.[151]

3.1.3. Trading Misconducts

a) Collusion

At the heart of functional capital markets lies a critical objective: the efficient distribution of financial resources and the promotion of responsible risk-sharing amongst all market participants. However, this reputation can be tarnished by certain market behaviours that threaten efficiency and public trust, including the inherent vulnerabilities of global finance to collusion. In economic terms, collusion refers to a secret agreement between competing companies to act together. Their goal is typically to raise profits above the level achievable by competing independently. This cooperation harms consumers by creating a situation known as a deadweight loss, where the overall economic output is lower than it could be. The companies involved need a system to manage their interactions and maintain collusion over time. This system would involve three key elements: first, agreeing on a shared strategy to maximise

[149] Cecconi, above n 148, at 96–98.
[150] US Securities and Exchange Commission "SEC Charges New York-Based High-Frequency Trading Firm with Fraudulent Trading to Manipulate Closing Prices" (Press Release 2014-229, 16 October 2014).
[151] *In the Matter of Athena Capital Research LLC* SEC Administrative Proceeding File No 3-16199 (16 October 2014) at [11]–[14].

profits; second, monitoring each other's behaviour to ensure everyone follows the plan; and third, enforcing the agreement by taking sanctions against any company that tries to undercut the others.[152]

Collusion, in its essence, manifests in two primary forms: explicit and implicit (or tacit). Explicit collusion involves direct agreements between companies to act anti-competitively. The most straightforward example is when firms agree on prices or production levels to maximise their collective profits. In some cases, these companies might even employ algorithms to manage and reinforce their collusive arrangements. Tacit collusion, conversely, achieves anti-competitive outcomes without any direct agreement. In such scenarios, companies recognise their mutual interdependence and adjust their behaviour, accordingly, making independent decisions that ultimately serve to maximise their individual profits. This type of collusion often arises in transparent markets with a limited number of participants, where companies can leverage their combined market influence even without direct communication.[153]

The widespread adoption of trading algorithms and AI-powered pricing agents by firms competing for market share has sparked significant concerns among global regulators and national antitrust authorities.[154] It is particularly pertinent when considering the two primary forms of collusion. Firstly, in explicit collusion, companies might intentionally design or utilise algorithms to directly manage and reinforce pre-existing anti-competitive agreements, such as coordinating real-time price adjustments or production levels. Secondly, and perhaps more concerning in the age of AI, is the heightened risk of tacit collusion. With the increasing use of algorithms and the potential development of more autonomous AI traders, the possibility of algorithmic collusion within financial markets has attracted increasing scholarly and regulatory attention.

Contemporary capital markets involve numerous algorithms interacting through electronic trading systems and responding to market data and changes in electronic order books. This raises the possibility of algorithms, even without explicit human instruction to

---

[152] Donald MacKenzie "How Algorithms Interact: Goffman's 'Interaction Order' in Automated Trading" (2019) 36(2) Theory, Culture & Society 39 at 46–49.
[153] OECD *OECD Handbook on Competition Policy in the Digital Age* (OECD Publishing, Paris, 2022) at 8–11.
[154] OECD, above n 153, at 37–45.

collude, learning to coordinate their behaviour in ways that achieve anti-competitive outcomes. Determining whether such coordinated behaviour results from deliberate and purposeful human design, an emergent consequence of algorithmic interactions, or a more accidental outcome of using sophisticated AI in a complex market environment becomes challenging.[155] Understanding how AI can facilitate explicit and tacit collusion is crucial for regulators seeking to maintain fair and competitive markets.

While the potential for AI-driven collusion in financial markets seems plausible, given advancements in machine learning, the absolute urgency of these risks remains unclear. Ultimately, the ability of algorithms to coordinate behaviour hinges on their capacity to learn and interact effectively within the market environment. Before delving into technical details, it is crucial to examine if specific segments of the increasingly digital capital markets provide fertile ground for the emergence of algorithmic "tacit" collusion. Regulators are undoubtedly becoming more aware of these ecosystems' potential algorithmic disruptions. However, the literature has not established that autonomous algorithmic collusion is an empirically prevalent feature of capital markets.[156]

One factor potentially enabling algorithmic collusion is market transparency. This transparency allows firms to monitor market trends and pricing, facilitating the detection of deviations from artificially high prices. Swift retaliation against such deviations is key to maintaining collusion. The heavily regulated nature of the financial services industry, with existing laws promoting market transparency for investor protection, inadvertently creates a technical environment potentially conducive to algorithmic collusion. Another factor is the frequency of interactions. In algorithmic trading environments, the high frequency of transactions allows for faster adjustments to maintain collusive pricing in terms of speed and cost. Capital markets, especially those involving highly liquid assets and HFT, are prime examples of such fast-paced, interconnected ecosystems.[157]

Finally, product homogeneity can also play a role when investors compare financial instruments based on risk and return, making them appear relatively similar. This inherent similarity can simplify collusion, as there is less need for complex agreements. Furthermore,

---

[155] Azzutti, Ringe and Stiehl, above n 12, at 105.
[156] Azzutti, Ringe and Stiehl, above n 12, at 105–106.
[157] Azzutti, Ringe and Stiehl, above n 12, at 107.

financial engineering tools like synthetic positions, which combine other financial assets to impersonate the performance of a specific instrument, further contribute to this homogeneity. Moreover, the ease of maintaining collusion is often linked to the concentration of a market. With fewer competitors, the economic incentive for rival firms to cooperate increases. While some segments of financial markets, like equity trading, are generally more competitive, there is a broader trend towards market consolidation and interconnectedness, especially among major global players.[158]

Entry barriers and innovation influence market concentration levels, although these factors have opposing effects. High entry barriers, such as licensing requirements and the need for significant reputational and financial capital, are typically seen as making collusion more sustainable. Conversely, innovation-driven markets are less susceptible to collusion. Financial regulations generally aim to promote competition and technological neutrality, but the reality of global capital markets differs. Entry barriers, like licensing and capital requirements, create an exclusive environment. While innovation can potentially disrupt competition, established firms often mitigate this risk by acquiring innovative companies.[159]

These market factors can influence the likelihood of "tacit" collusion in capital markets characterised by repeated pricing interactions. Furthermore, delegating decision-making to algorithms can ease the path to coordination for competing firms, eliminating the need for direct communication. Algorithms offer a significant advantage over traditional human-managed cartels due to their superior speed and analytical capabilities. The increasing sophistication of ML methods and the ability to leverage Big Data adds another layer of complexity. Self-learning AI agents could create new cartel-like behaviour and alter the communication required to achieve illicit agreements or replicate similar market effects. For instance, highly transparent algorithms might be able to predict rivals' strategies and use this information to coordinate their actions.[160]

---

[158] Nicola Cetorelli, Beverly Hirtle, Donald Morgan, Stavros Peristiani and João Santos "Trends in Financial Market Concentration and their Implications for Market Stability" (2007) 13(1) Economic Policy Review 33 at 36–42.
[159] Stefania Vitali, James B Glattfelder and Stefano Battiston "The Network of Global Corporate Control" (2011) 6(10) PLoS ONE e25995 at 4–6.
[160] OECD, above n 153, at 48–52.

Therefore, market segments within the complex network of global capital markets that exhibit these facilitating factors are more susceptible to algorithmic "tacit" collusion risks. However, key questions remain unanswered: Can independent AI traders achieve some form of coordination without human guidance, and if so, how effective would it be? Communication among independent pricing algorithms also seems crucial for establishing and maintaining coordination. To assess the risks of algorithmic tacit collusion by AI traders, a first step would involve evaluating the critical aspects of algorithm properties, functionality, and their resulting effects.[161] While algorithmic collusion represents a potentially significant consequence of AI integration in financial markets, it is not the central focus of this thesis

b) Insider Trading

Insider trading constitutes a form of market misconduct where an "information insider" exploits their access to undisclosed, material information. This information, typically of a company or entity listed on a licensed exchange like the New Zealand Exchange ("NZX"), is not publicly available and could significantly impact the price of financial products if it were. Information insiders can breach regulations in various ways. Based on this inside information, they might directly buy or sell relevant financial products. Alternatively, they could share this information with others, who would act on it by buying or selling. Insider misconduct can also extend to encouraging others to buy, sell, or hold financial products, even without explicitly disclosing the reason behind such advice.[162]

Illustrative examples of insider conduct include a company's chairperson selling their shares before negative company news becomes public; a company employee, privy to confidential whiteboard notes about an upcoming, unannounced offshore expansion, tipping off their father to purchase company shares; or a lawyer sharing confidential information about a client's imminent positive announcement, likely to raise share prices, with a friend involved in share trading.[163] The Financial Markets Authority ("FMA") took action in 2019 against a former chair of Promisia Integrative Limited ("PIL"), a company listed on NZX. The latter admitted to engaging in insider trading and violating his disclosure obligations as a director. These breaches occurred shortly after he resigned from both chair and director positions at PIL.

---

[161] Azzutti, Ringe and Stiehl, above n 12, at 109–112.
[162] Financial Markets Authority *Insider Conduct: Information Sheet* (FMA, Wellington, March 2021) at 2–4.
[163] FMA, above n 162, at 4–6.

He leveraged his prior access to confidential sales information, which was unavailable to the public, to acquire over 2.5 million PIL shares.[164]

Exchanges have adapted by restructuring their architecture to offer co-location services. Firms can purchase physical space to house their servers directly on-site or adjacent to the exchange's servers. In addition to co-location, traders can acquire detailed data feeds directly delivered to their co-located servers. Finally, the dominance of high-speed algorithmic trading and AI integration have fundamentally reshaped the securities market by eliminating the need for human oversight in real-time trade execution. This combination of factors creates a prime position for HFT traders, enabling them to make informed gains ahead of other investors, anticipate market movements, and secure the most favourable trades. While traditional market makers were typically considered uninformed from a finance theory perspective, the rise of HFT suggests a different reality. Studies increasingly reveal the skill of HFT traders in predicting short-term market direction and capitalising on order flow by trading ahead of it.[165]

The architecture of contemporary automated markets represents a significant departure from the past, increasingly granting a substantial informational advantage to a select group of high-speed algorithmic traders. Due to their inherent structural advantages – such as superior connectivity and co-location – these traders possess the distinct ability to be the first to receive and analyse market data. This temporal edge allows them to anticipate future trading behaviour and proactively adjust prices based on this privileged access. Yadav describes this phenomenon as "structural insider trading", reflecting the informational and speed advantages available to certain high-speed traders and the resulting concerns about market fairness. The expression is used as a scholarly concept and is not a category of insider conduct recognised under New Zealand law.

However, a critical distinction lies in its legal character: unlike conventional insider trading, the structural advantages described by Yadav, such as co-location and faster access to market data, do not without more constitute prohibited insider trading. These structural advantages may nevertheless raise broader concerns about market fairness and informational

---

[164] Financial Markets Authority "Former Promisia Chairman Admits Insider Trading Breaches" (Media Release 2019-44, 24 October 2019).
[165] Yesha Yadav "Insider Trading and Market Structure" (2016) 63 *UCLA Law Review* 968 at 971–980.

asymmetry that are distinct from the statutory prohibition on insider conduct. [166] This issue was raised to illustrate a distinct yet related challenge to market fairness arising from technological advantages.

c) Market Manipulation (Thesis Focus)

Market manipulation, a malicious form of market conduct, represents a calculated effort to subvert the fair and transparent trading environment that underpins healthy capital markets. These manipulative tactics, rooted in economic principles, aim to artificially inflate or deflate the prices or supply of financial instruments traded on exchanges like the NZX or ASX. It can be achieved through deceptive means that distort market forces, ultimately tricking investors into misrepresenting trades based on false information. Consequently, it leads to an inefficient allocation of financial resources, thereby representing a form of market failure. The negative externalities, or unintended consequences, generated by market manipulation cannot be rectified by market forces alone. Therefore, such activities are generally prohibited and often prosecuted as a form of market abuse.[167]

Generally, market manipulation falls into two primary categories: information-based and transaction-based. Information-based manipulation occurs when individuals knowingly or unknowingly disseminate demonstrably false or misleading information concerning a specific financial product, such as a particular stock or bond. This deception aims to distort investor behaviour, potentially triggering a surge or decline in trading volume for that financial product. Furthermore, it can artificially inflate or deflate the market price or sway shareholder votes on company matters. For instance, a manipulator might fabricate positive news about a company's stock to artificially inflate its price, allowing them to sell their shares at a profit before the price corrects. Conversely, they might spread negative rumours to depress a stock's price, enabling them to acquire shares cheaply before a potential recovery. Increasingly, some traders are leveraging social media platforms and online forums to execute these information-based manipulative tactics.[168]

---

[166] Yadav, above n 165, at 977–985.
[167] Gina-Gail S Fletcher "Deterring Algorithmic Manipulation" (2021) 74 Vanderbilt Law Review 259 at 291–294.
[168] Financial Markets Authority 'Market Manipulation' [2020] https://www.fma.govt.nz/library/articles/market-manipulation/#:~:text=Penalties%20for%20market%20manipulation&text=In%20the%20most%20serious%20cases,or%20%242.5%20million%20for%20companies.

Another form of market manipulation involves deceptive trading activity, encompassing actions or inactions, such as placing orders, that intentionally create a misleading impression about a financial product. This manipulation can target perceptions of a product's popularity, availability, price, or inherent value. For instance, placing orders to buy shares without the intention of purchasing them creates an illusion of high demand. Similarly, simultaneously buying and selling the same product can distort the true market interest. Furthermore, placing numerous orders at varying prices can also constitute manipulation.[169] There have been few market manipulation cases in New Zealand,[170] with the first case decided in 2014.[171]

AI and machine-learning techniques may alter the traditional landscape of market manipulation by increasing the speed, scale and complexity with which trading strategies can be implemented. Advances in machine learning may also create the possibility that an adaptive system pursuing a specified objective could develop trading strategies that produce manipulative or manipulation-like market effects without those precise strategies having been expressly programmed by a human operator.[172] The opacity of some AI systems may also make potentially manipulative conduct more difficult to detect, reconstruct and attribute.[173] These possibilities should, however, be treated as emerging risks rather than evidence that autonomous AI manipulation is already prevalent in financial markets.

A related concern is that AI-enabled trading systems may be deliberately designed or deployed to generate or exploit misleading market signals to influence prices or investor behaviour. This ability to manipulate human behaviour holds significant weight because, ultimately, the value of financial instruments is intrinsically tied to the underlying human actions that fuel the market. By effectively controlling investor behaviour, AI-enabled trading algorithms could achieve unprecedented accuracy in predicting price movements, allowing them to accumulate significant profits at the expense of others and society and create a one-way market. Three prevalent methods of algorithmic manipulation include pinging, spoofing, and electronic front running.

---

[169] FMA, above n 168.
[170] *FMA v Mark Warminger* [2017] NZHC 327; *FMA v Wei Zhong & Ors* [2022] NZHC 480; *FMA v Wei Zhong* & Ors [2023] NZHC 766.
[171] *Financial Markets Authority v Henry,* above n 44, at [70] – [78].
[172] Azzutti, Ringe and Stiehl, above n 12, at 97–99.
[173] Kiyohiko G Nishimura "Electronic Trading and Financial Markets" (Speech, International Paris-Europlace Financial Forum, Tokyo, 29 November 2010) at 4–6.

Pinging, in the context of market manipulation, refers to the practice of sending a rapid sequence of small-sized limit orders across the order book for a specific security. These orders are intentionally short-lived, typically cancelled within milliseconds before they can be matched and executed. The purpose of this high-frequency, low-volume activity is to observe the immediate responses of other market participants, particularly algorithmic trading systems and market makers. By analysing the changes in the order book – such as shifts in the best bid and offer, the emergence or disappearance of liquidity at different price levels, and the speed of these adjustments – the manipulator attempts to infer the presence and characteristics of larger hidden orders or the price points at which significant buying or selling interest exists. This information advantage gained through minimal market exposure can then be used to inform manipulative trading strategies. [174]

Spoofing is a manipulative trading strategy typically executed in a series of steps designed to create a false impression of market interest and induce other participants to trade at artificial prices. Firstly, a trader places a sizeable, inauthentic order on a specific side of the market (either buy or sell). This large order is not intended for execution but is meant to create artificial buying or selling pressure, causing a temporary price movement in the desired direction. Secondly, once the market price has shifted favourably due to this deceptive order, the trader submits a smaller, legitimate order on the opposite side of the market. This genuine order is then filled at the artificially inflated or deflated price, allowing the spoofer to profit. Finally, the large, inauthentic order that initiated the price movement is cancelled before it can be executed, concluding the manipulative scheme. HFT algorithms can be deployed to automate the spoofing scheme, enabling its repeated execution at a high volume and across multiple asset classes. These algorithms can be programmed to strategically place and subsequently cancel large orders that manipulate market prices on one side while simultaneously submitting orders on the opposing side to capitalise on the artificial price movements generated.[175]

On 6 May 2010, US equity and futures markets experienced a rapid decline and recovery amid stressed market conditions and a severe loss of liquidity. The joint CFTC and SEC report identified a large automated sell programme and interactions among high-

[174] Lin, above n 136, at 1288–1290.
[175] Lin, above n 136, at 1290–1293.

frequency traders as contributing factors. Sarao was subsequently found to have used manual and automated spoofing strategies in the E-mini S&P 500 futures market, including on the day of the Flash Crash. His conduct should not, however, be characterised as the sole cause of the Flash Crash or as an example of autonomous AI manipulation.[176]

Following the Flash Crash, regulatory bodies intensified their scrutiny of manipulative trading practices.[177] In July 2013, the CFTC entered a civil administrative order against Michael Coscia and Panther Energy Trading LLC concerning algorithmic spoofing. The CFTC found that Coscia used algorithmic trading programmes to place and rapidly cancel orders that created a false impression of market interest, while trading on the opposite side of the market. Coscia was subsequently criminally prosecuted and convicted of spoofing and commodities fraud. These proceedings concerned deliberately programmed, rule-based trading strategies and should not be characterised as examples of autonomous AI misconduct.[178] The Flash Crash and subsequent enforcement actions underscore the importance of regulatory controls addressing manipulative practices such as spoofing and the risks associated with automated trading strategies. [179]

In both the *Coscia* and *Sarao* cases, forensic analysis of the trading algorithms yielded crucial evidence that allowed regulators to understand the traders' actions and their specific methods. Moreover, the inherent structure and logic embedded within these rule-based algorithms provided a strong indication of the intended manipulative purpose of the trading strategy. While Coscia and Sarao argued that the algorithms' behaviour did not align with their conscious design as programmers, regulatory access to the algorithms' source code revealed the manipulative intent reflected in the programming logic. In similar cases of deliberate algorithm misuse for manipulation, regulators should be able to meet the legal requirement for proving intent by examining the digital record within the algorithms' source code. If this code demonstrates clear or readily apparent indications of a manipulative purpose, it should provide sufficient grounds to hold the responsible individuals accountable for spoofing.[180]

---

[176] US Commodity Futures Trading Commission and US Securities and Exchange Commission, above n 34, at 1–6.
[177] Fletcher, above n 167, at 292–293.
[178] US Commodity Futures Trading Commission "CFTC Orders Panther Energy Trading LLC and Michael J Coscia to Pay $2.8 Million and Bans Them from Trading for One Year" (Press Release PR6649-13, 22 July 2013).
[179] CFTC and SEC, above n 34, at 7–14.
[180]Fletcher, above n 167, at 294–295.

Another challenge emerges when an algorithm behaves unpredictably, diverging from its intended operation. In such cases, unexpected market distortion does not arise from the algorithm's logical data processing. Instead, this unintended disruption results from flaws in the algorithm's design and pre-deployment testing or insufficient oversight and response to potential warning signals and coding errors. Simultaneously, numerous hypothetical situations exist where trading algorithms could unintentionally manipulate markets, including specific possibilities such as "rational distortion," where algorithms exploit market rules or inefficiencies in a way that, while logical for the algorithm, leads to unintended and destabilising price movements, and "independent misconduct," where complex AI algorithms develop unforeseen and undesirable trading strategies that were not explicitly programmed or anticipated. It is particularly relevant for AI-driven trading algorithms.[181]

Electronic front-running is another form of market manipulation which exploits access to confidential order flow information to gain an unfair advantage. Traditionally, front-running involved human brokers placing personal trades ahead of client orders. However, algorithmic trading has introduced a more sophisticated type of front-running. Brokers may leverage their position within a trading venue to gain a sneak peek at incoming orders, allowing them to execute their trades at a favourable price just before the larger order arrives, manipulating the market in their favour. The speed and automation of algorithmic trading make electronic front-running particularly difficult to regulate.[182]

Another significant incident occurred on 5th February 2018. The Dow Jones Industrial Average plunged a staggering 1,597 points at its lowest point – the most essential single-day drop in its history. While the market did recover somewhat, experts believe lightning-fast trading models and automated algorithms played a significant role in amplifying the sell-off. These rapid-fire programs can amplify fear and panic among investors, leading to a domino effect of selling that exacerbates the initial decline. This incident highlights a growing unease about the increasing reliance on algorithms in the markets. Critics argue that these programs create a system that moves too fast for human comprehension, fuelling volatility and potentially eroding investor confidence.[183]

---

[181] Fletcher, above n 167, at 296.
[182] Morgan Lewis *Algorithmic Trading: Hong Kong Monetary Authority Discusses Supervisory Expectations and Evolving Market Manipulation Risk* (Client Alert, March 2020) at 13.
[183] Gregory Scopino *Algo Bots and the Law: Technology, Automation, and the Regulation of Futures and Derivatives* (Cambridge University Press, Cambridge, 2020) at 411–447.

In an investigation on high-frequency trading tactics in the US, the Financial Industry Regulatory Authority ("FINRA") recently penalised a New York brokerage firm and its employees. The firm, formerly known as Trillium Trading, was accused of using deceptive manoeuvres to trick other traders' automated programs. Over three months, Trillium's traders repeatedly placed and cancelled many fake orders to create a misleading impression of solid buying or selling pressure in specific stocks. This manipulation aimed to lure unsuspecting participants into executing trades at manipulated prices that were either artificially high or low. The regulators highlighted the concerning pattern of these rapid-fire order placements and cancellations, suggesting potentially fraudulent activity.[184]

The US Securities and Exchange Commission ("SEC") has actively pursued those who manipulate the market by spreading false information for their own gain. For instance, in December 2020, the SEC filed a complaint against Barton S. Ross, alleging he inflated his company's stock price over two years by disseminating fake news across various media. Investigators found at least 49 instances of Ross creating and spreading this misinformation, leading to $36,000 in illegal profits.[185] Similarly, in August 2021, the SEC charged Mark Melnick, a live webcast host promoting stock trading, with market disruption for spreading over 100 fabricated rumours. His objective was also to artificially inflate stock prices and illicitly profit, which amounted to $374,000. Ultimately, these offenders were forced to disgorge their ill-gotten profits and pay interest and penalties.[186]

Market manipulation involving trading algorithms is also a known problem in New Zealand and Australia. In recent years, NZX imposed regulatory sanctions and public censures on broking firms, including Craigs Investment Partners Limited and Macquarie Securities (NZ) Limited, concerning orders placed by two clients with Direct Market Access via algorithmic trading, which disrupted the exchange market. Upon investigation, it was found that these two firms failed to monitor a series of anomalous algorithmic-based trades where the share

---

[184] *Department of Market Regulation v Lek Securities Corporation and Samuel Frederik Lek* FINRA Disciplinary Proceeding No 2011029713004; FINRA "FINRA Sanctions Trillium Brokerage Services LLC" (Press Release, 13 September 2010).
[185] US Securities and Exchange Commission "SEC Charges Former Day Trader with Market Manipulation Scheme" (Litigation Release No 24989, 23 December 2020).
[186] US Securities and Exchange Commission "SEC Charges Webcast Host for Role in Market Manipulation Scheme" (Press Release 2021-206, 25 August 2021).

ownership did not change due to the lack of regulatory and technological resources and an appropriate market supervision framework.[187]

In Australia, Gabriel Govinda was sentenced to two and a half years in prison and a fine on charges of manipulation of shares listed on the Australian Securities Exchange (ASX) and dissemination of information relating to manipulation. His manipulation tactic was layered. Firstly, he used fake accounts registered under the names of friends and relatives to buy and trade stocks, giving the impression of high demand and driving the prices up. After that, he promoted these stocks online via the popular share market forum "HotCopper". He essentially "pumped" the stock price even higher by sharing information about his market manipulation activities, which attracted unsuspecting investors to fall into the trap. Govinda sold his holdings at a significant profit.[188] This case was the first person sentenced for an offence under s 1041D of the Corporations Act 2001 (Cth).

The high-speed nature of algo-trading, facilitated by complex algorithms, makes pinging, spoofing and electronic front-running challenging for human traders to detect.[189] Distinguishing between these two scenarios, including deliberate manipulation versus unintended consequences, poses a significant challenge for regulatory bodies and victims.[190] Economic harm and detrimental societal effects caused by market manipulation are well-documented, and a robust global consensus needs to be fostered against them. While the specific legal definitions and the extent of their application may vary across different jurisdictions, the core objective remains unwavering: to protect the market's integrity by curbing manipulative practices.[191]

---

[187] New Zealand Markets Disciplinary Tribunal *Public Censure of Craigs Investment Partners Ltd* (Decision, 2019); New Zealand Markets Disciplinary Tribunal *Public Censure of Macquarie Securities (NZ) Ltd* (Decision, 2019).
[188] Australian Securities and Investments Commission "Gabriel Govinda Sentenced to Two and a Half Years' Imprisonment and Fined for Market Manipulation and Finfluencer Conduct" (Media Release 23-116, 5 May 2023).
[189] Morgan Lewis, above n 182, at 12–13.
[190] Fletcher, above n 167, at 296–298.
[191] Fox, Glosten and Rauterberg, above n 67, at 209–246.

3.2 THE LIABILITY CONUNDRUM: ASSIGNING RESPONSIBILITY FOR AI-DRIVEN HARM

The core challenge posed by AI-powered algorithms lies in their intricate functioning, creating what is often called a “black box”. This term signifies the fundamental inability to fully comprehend the AI's decision-making process and, consequently, the inability to predict its outcomes. The inability to elucidate the algorithm's outputs after the fact restricts our ability to grasp its inner workings or effectively supervise its ongoing use. In the most extreme scenarios, characterised as “strong black boxes,” this lack of transparency becomes absolute. Even detailed retrospective analysis and attempts to reverse engineer the algorithm to decipher its decision-making rationale are rendered entirely impossible. This inherent opacity effectively shields the algorithms' functioning from meaningful human oversight and supervision.

As these AI-enabled algorithms become increasingly integrated into the complex dynamics of financial markets, significant concerns emerge regarding the adequacy of the existing regulatory framework to effectively prevent and deter instances of market manipulation. A key distinction arises when comparing these AI-driven systems to traditional HFT algorithms. While HFT algorithms operate based on clearly defined, rule-based criteria, AI-enabled algorithms possess the capacity to learn from data and make dynamic, intuitive decisions. This inherent adaptability poses a considerable challenge to the current anti-manipulation regime, which is heavily predicated on the legal concept of a demonstrable ‘intent to deceive’ in market conduct. [192]

The “black box” phenomenon introduces a critical transparency issue within financial markets. This lack of transparency arises when human developers themselves struggle to fully understand, predict, and provide clear explanations for the specific logic that an algorithm employs to arrive at a particular decision or solution based on the input data it processes. Given the current legal landscape governing financial markets, regulatory authorities often face significant challenges in deciphering the underlying reasoning behind these complex algorithmic decisions. Such a lack of understanding directly impedes their ability to effectively

[192]Azzutti, Ringe and Stiehl, above n 12, at 118–122.

assign legal liability in cases of potential wrongdoing or market abuse stemming from these automated systems.[193]

The inherent opacity characteristic of "black box" algorithms significantly complicates the practical enforcement of existing legal frameworks. These frameworks were primarily conceived and designed to address human actors' actions and deliberate decision-making processes within the financial system. In essence, the self-learning capabilities exhibited by autonomous AI trading systems raise fundamental and critical legal questions concerning the assignment of liability on account of their inherently opaque nature, particularly with the increasing integration of sophisticated AI and ML techniques. This confluence of factors creates substantial challenges for effectively regulating and preventing potential instances of market abuse within financial markets.[194] Looking ahead, it is conceivable that established liability rules may undergo a process of evolution. The evolution could potentially lead to the imposition of accountability on relevant actors for the ultimate outputs generated by their algorithms, possibly without necessitating a comprehensive understanding of the intricate internal decision-making processes employed by the AI.

A significant conceptual hurdle arises when attempting to apply traditional legal concepts to the realm of algorithmic trading. Foundational notions such as guilt, which are typically intertwined with human intent and the foreseeable consequences of deliberate actions, become considerably blurred when AI systems autonomously "learn" their trading behaviours and strategies from vast and complex datasets. In contrast to human traders, who deliberately design and implement trading strategies based on their interpretation of market dynamics, AI trading systems operate by detecting complex patterns within the data on which they are trained. This fundamental distinction complicates the attribution of malicious intent to any specific individual when an AI system produces a potentially harmful trading outcome.

Furthermore, the traditional legal concept of causation becomes less clear and more challenging to establish in the context of AI-driven trading. While a human programmer initially designs the underlying algorithm and selects the data used for training, the specific trading decisions made by the AI that ultimately led to a problematic market event might be

---

[193] Yesha Yadav "The Failure of Liability in Modern Markets" (2016) 102 Cornell Law Review 1031 at 1039 and 1083.
[194] Azzutti, Ringe and Stiehl, above n 12, at 118–119.

significantly removed from the initial programming. These decisions are often a direct result of the AI's autonomous learning and adaptation within the dynamic and frequently unpredictable market environment.

The perceived absence of direct human “guilt” in the immediate trading decisions of AI systems is further amplified as these systems gain increasing levels of autonomy, operating with progressively less direct and continuous human oversight. Consequently, the process of assigning legal responsibility in cases of AI-driven harm becomes a multifaceted and considerably complex problem for regulatory and enforcement bodies. Theoretically, these bodies might need to implicate numerous individuals who were involved at various stages of the AI's lifecycle. It could include those who initially conceived the AI's purpose, the programmers who wrote the underlying code, the data scientists who curated the training data, the individuals responsible for deploying the system into the market, and those tasked with its ongoing monitoring and maintenance.

However, the sheer number of potentially liable parties involved in the lifecycle of a sophisticated AI trading system makes practical enforcement of legal responsibility incredibly challenging and resource-intensive for authorities. Moreover, the inherent self-learning capabilities that are a core feature of many ML algorithms further obscure the lines of accountability. As the AI's trading strategies and behaviours can evolve organically and autonomously within the market, potentially in unforeseen and unintended ways, it becomes even more challenging to trace a direct and clear line of causation back to a specific human actor's deliberate wrongdoing or negligence.

The inherent opacity of ML algorithms presents a significant challenge for several key reasons that directly impact the effective regulation and oversight of financial markets. Firstly, accurately identifying the root cause of market harm becomes exceedingly difficult when the fundamental decision-making process of an AI is opaque and not readily understandable by human analysts or regulators. When a trading algorithm generates substantial financial losses or contributes to market instability, a comprehensive understanding of why the AI executed those specific trades is crucial for subsequently determining legal liability. Without transparency into the AI's reasoning, pinpointing the responsible party becomes a complex and potentially insurmountable task. As a result, it could be the programmer who introduced a critical flaw, the financial firm that deployed the algorithm without adequate safeguards, or

even the AI itself in more nuanced legal and ethical discussions about the evolving agency of sophisticated systems.

Secondly, this inherent opacity considerably complicates effective regulatory oversight and enforcement within financial markets. Regulatory bodies require a clear understanding of how AI trading algorithms function to ensure their compliance with existing market regulations and to proactively identify and mitigate potential systemic risks to the stability and integrity of the financial system. Suppose the algorithms' decision-making processes remain opaque and inaccessible to human scrutiny. In that case, regulators struggle to assess compliance, detect potentially manipulative behaviours that may not be obvious through traditional surveillance methods, or fully understand the potential for widespread market disruption caused by these robust automated systems. A fundamental lack of transparency can severely undermine the overall effectiveness of regulatory frameworks specifically designed to maintain market integrity and protect the interests of investors.

Finally, the opacity of AI-driven financial systems erodes trust and confidence within the financial system. When the underlying mechanisms driving significant financial decisions are hidden from public view and expert scrutiny, it can breed suspicion and distrust among individual investors and the broader public. When financial losses occur due to the actions of opaque algorithmic trading systems, and the fundamental reasons behind those actions remain unclear and unexplained, it can significantly undermine faith in the financial market's fairness, efficiency, and overall stability. Establishing accountability for such events and effectively rebuilding public trust necessitates a greater degree of transparency in how these increasingly powerful technologies operate within the financial system.

In the specific context of assigning legal liability for harm caused by AI-driven systems, the inherent opacity present in many ML algorithms directly impedes the ability of legal and regulatory bodies to effectively understand the causal chain of events leading to the harm. Lacking transparency also hinders the accurate assessment of fault or negligence, complicates the process of ensuring regulatory compliance, and ultimately undermines the maintenance of public trust in the fairness and stability of the markets – all of which are critical elements in establishing legal responsibility.

Wallach discusses these “black box” problems are not entirely novel concerns within the realm of technology. He points out that specific deterministic algorithms, which operate based on pre-defined rules, have existed for decades and may have already exhibited independent and sometimes unpredictable behaviours within complex systems.[195] Lipton separately examines that the advent of sophisticated AI and ML has introduced significant qualitative changes compared to traditional rule-based algorithms. These changes include a blurring and distortion of the operational parameters that previously distinguished relatively “dumb” algorithms used in HFT from the smarter and more adaptive algorithms employed in other forms of algorithmic trading. [196]

Traditional trading algorithms, including the programs that underpin HFT, operate within a clearly defined and often static framework. They are essentially pre-programmed sets of explicit instructions meticulously designed to execute a very specific and often narrowly defined trading strategy or to operate strictly within predetermined parameters set by human traders and risk managers. These traditional algorithms' logic and decision-making processes are typically transparent and auditable.

In contrast, AI and ML algorithms represent a fundamentally different and more dynamic approach to automated trading. These advanced algorithms are not provided with explicit, step-by-step instructions on how to solve a particular trading problem or achieve a specific financial goal. Instead, human programmers and data scientists equip these AI systems with a high-level goal, such as maximising investment returns or minimising trading risk and provide them with access to vast quantities of relevant historical and real-time market data. The AI then autonomously “learns” through an iterative process of trial and error, continuously analysing past data and the outcomes of its trading decisions to dynamically refine its internal decision-making process and improve its performance over time.

This inherent iterative learning capability allows AI algorithms to execute trades automatically and potentially discover novel and complex trading strategies that human traders might not have explicitly conceived. It significantly reduces the need for direct and continuous human intervention in the day-to-day trading activities of these sophisticated systems. Unlike

---

[195] Wendell Wallach “Implementing Moral Decision-Making Faculties in Computers and Robots” (2007) 22(4) AI & Society 463 at 464–473.
[196] Zachary C Lipton “The Mythos of Model Interpretability” (2018) 16(3) ACM Queue 31 at 35–42.

traditional algorithms, where the precise sequence of steps leading to a particular trading output can be readily traced back through the pre-programmed rules and logic, the internal decision-making process within AI and ML-enabled algorithms often remains somewhat opaque and complex for humans to fully understand. [197]

The inherent opacity is frequently described as “intuitive” in the context of AI. It means that the specific reasoning behind a particular trading output or decision generated by the AI is often not readily apparent or easily explainable in human terms, even to the algorithm's own designers. The complex and interconnected network of relationships and learned patterns within the AI's internal “thinking” makes it exceedingly difficult to pinpoint the exact chain of reasoning that led to a specific trading action. This inherent complexity creates a significant transparency issue that poses challenges for oversight and accountability. As a result of this lack of explainability, these advanced algorithms can be accurately likened to "black boxes" due to our limited understanding of their intricate decision-making processes and the consequent difficulty in reliably predicting their future outputs under various market conditions.

Moreover, the fundamental characteristics of many AI and ML algorithms present a significant challenge because even the human designers and developers of these sophisticated systems often only have a limited understanding of how the AI precisely arrives at its specific trading decisions in any given situation. A lack of complete explainability inherently hinders effective human oversight and supervision of these powerful automated trading tools.[198]

In the opaquest and most challenging “strong black box” scenarios, the level of complexity is such that even sophisticated attempts at reverse engineering the algorithm to understand the underlying reasoning behind a specific AI trading decision are rendered practically impossible due to the intricate and non-linear nature of the learned relationships within the system. Furthermore, the complex and often poorly understood interactions that can occur between different autonomous trading programs operating within the same market introduce another layer of risk and uncertainty. The lack of clear understanding and the inherent

[197] Azzutti, above n 75, at 5.
[198] Azzutti, above n 75, at 5.

secrecy surrounding the specific trading strategies employed by these algorithms make it exceedingly difficult for market participants and regulators to effectively manage systemic risk.

Even in scenarios where advanced AI programs are specifically designed to cooperate and function harmoniously within a portfolio or trading system, these systems' increasing complexity and interconnectedness can lead to unintended and potentially disruptive competitive behaviours. This can occur even when these AI systems' human users and developers act with due care and diligence in their design and deployment. [199]

The use of increasingly autonomous AI systems raises questions about how responsibility for market misconduct should be attributed. AI systems are not presently treated as legal persons capable of bearing legal responsibility for such misconduct. The relevant question is therefore how responsibility should be allocated among the natural persons and corporations involved in designing, developing, deploying, controlling and monitoring AI-driven trading systems.[200] The relevance of concepts such as intent, knowledge, causation and negligence depends on the particular legal rule being applied. These issues are identified here as part of the liability challenges associated with increasingly autonomous trading systems and are examined in greater detail in Chapters 4 to 6.

Identifying potential legally liable parties in cases of harm directly caused by AI-driven trading activity often includes those individuals and entities with decision-making authority over the AI system's deployment and operational parameters. It also extends to those directly involved in the AI's initial conceptualisation, its subsequent development and coding, the crucial process of curating and preparing the training data it learns from, and the ongoing maintenance and oversight of the system once it is deployed in the live market environment. It suggests a potential legal model where accountability for the actions of an AI trading system might be distributed across a range of human actors and corporate entities involved in its lifecycle rather than being solely attributed to a single individual or the AI itself.

An alternative perspective considers AI as simply a highly sophisticated tool that is ultimately wielded and controlled by human users. However, the increasing levels of autonomy

[199] Azzutti, above n 75, at 4.
[200] Azzutti, above n 75, at 4–5.

exhibited by many modern AI systems, particularly those employing advanced Machine Learning techniques, significantly complicates this straightforward analogy of AI as a mere instrument. This increasing autonomy raises the critical and evolving question of precisely where legal liability should be assigned as the capabilities of AI systems advance beyond simple instrumentality and towards more independent decision-making.

The inherent opacity of many advanced AI systems further exacerbates this fundamental challenge of assigning legal liability. Traditional legal methods for proving key elements of liability, such as specific intent or negligence, struggle significantly when the underlying reasoning behind an AI's harmful trading action is inscrutable and not readily explainable in human terms. This lack of transparency makes it exceedingly difficult to pinpoint a specific human party who exhibited the requisite mental state or failed to meet a reasonable standard of care that directly led to the harmful outcome generated by the AI.

The assertion that this fundamental disconnection between the complex decision-making processes of AI and readily identifiable human actors fundamentally undermines the effectiveness of law enforcement in the context of AI-driven market abuse hinges on one's specific understanding of the primary objectives of law enforcement. Suppose the core function of law enforcement is narrowly defined as the identification and subsequent prosecution of human actors based on demonstrable evidence of specific intent or negligence that directly caused harm. In that case, the inherent opacity of AI systems undeniably presents a significant impediment to these traditional investigative and prosecutorial approaches. The difficulty in tracing a harmful market outcome back to a culpable human mind operating within the “black box” of an AI can indeed frustrate established legal processes and evidentiary requirements.

However, if the purpose of law enforcement is viewed through a broader lens – encompassing not only the prosecution of wrongdoing but also the proactive prevention of harm, the maintenance of overall public safety and the integrity of financial markets, and the establishment of clear accountability within increasingly complex technological systems – then the challenges posed by AI opacity might necessitate an evolution and adaptation of legal strategies rather than a complete undermining of law enforcement effectiveness.

Within this broader context, law enforcement and regulatory bodies might adapt their approaches by focusing on regulating the development and deployment of AI technologies in

high-stakes domains like financial markets. This could involve establishing stricter standards of care and due diligence for those who create and utilise highly autonomous systems. Furthermore, it might necessitate the exploration and potential adoption of novel legal doctrines specifically designed to address the unique characteristics of AI agency and their potential for causing harm, even without clear evidence of traditional human intent or negligence. Maintaining human oversight may provide one means of reducing risks associated with increasingly autonomous trading systems. However, the presence of a human in the loop does not itself resolve questions of responsibility or provide an effective safeguard unless that person has sufficient information, authority and opportunity to intervene. [201]

The fundamental legal concept of "causation" in tort law presents significant conceptual and practical hurdles when applied to situations involving harm directly caused by artificial intelligence. Establishing legal causation necessitates demonstrating a direct and clear link between a specific action or omission on the part of a defendant and the resulting harm suffered by a plaintiff. The core challenge of establishing a clear causal link with AI-driven harm lies in the inherent difficulty of tracing the specific sequence of computational steps within the algorithm's complex decision-making process, ultimately leading to harmful outcomes in the real world. This difficulty in precisely pinpointing the chain of algorithmic actions distinguishes AI-related harm from scenarios where the causal pathway between human behaviour and its consequences is often more direct, observable, and attributable.

The inherent "black box" nature of many AI trading tools further complicates the fundamental legal issue of attributing responsibility for market harm caused by these systems. Market abuse perpetrated by an AI could originate from a multitude of unforeseen and often counter-intuitive circumstances. These could include complex reasoning based on patterns in market data that are entirely invisible to human analysts, the extrapolation of obscure and potentially spurious correlations within vast datasets, or the exploitation of novel trading strategies that lie completely beyond the realm of human comprehension or traditional market understanding.

The extraordinary speed and intricate interconnectedness of these sophisticated AI-driven trading systems can also amplify the consequences of unforeseen algorithmic

[201] Azzutti, above n 75, at 4–5.

behaviours across multiple markets with remarkable rapidity and scale. This makes the traditional legal assessments of causation – which often rely on establishing a direct and foreseeable link between a specific human act and a resulting harm – exceptionally difficult, if not practically impossible, to apply effectively in the context of complex AI systems.

Regarding the legal principle of "foreseeability," its direct application to autonomous AI agents is not necessarily straightforward and requires careful consideration. One could plausibly argue that because these advanced systems often operate as "black boxes," a potentially wider range of unpredictable behaviours and market outcomes might be considered foreseeable in a general sense by those who develop and deploy them. The very nature of complex, self-learning systems inherently implies a potential for actions and outcomes that were not specifically programmed or explicitly anticipated by their creators.

Therefore, the argument that the principle of foreseeability becomes entirely inapplicable in the context of autonomous AI might be inaccurate. Instead, the core challenge lies in accurately defining the appropriate scope of what a reasonably prudent developer or deployer of such a complex system should have foreseen as potential consequences of its operation within the dynamic and often volatile financial markets. The established legal framework focuses on what was reasonably foreseeable at the time of the relevant action or omission, not merely on what was foreseen by the defendant. Many events that ultimately proved to be unforeseen consequences of deploying AI systems were, with the benefit of hindsight and expert analysis, potentially foreseeable risks associated with deploying such complex and opaque technologies into critical market infrastructure.[202]

Autonomous AI presents challenges for established market abuse prohibitions and traditional liability rules. The difficulty of assigning responsibility where multiple individuals contribute to an AI system's development and deployment is not entirely novel in law, as doctrines of corporate liability, agency and vicarious liability already address circumstances involving distributed responsibility. However, the opacity and autonomy of AI-driven trading systems may complicate the identification of relevant human conduct, the attribution of

[202] Azzutti, above n 75, at 4.

knowledge or intent, and the establishment of causal links between the design or deployment of a system and resulting market harm.[203]

These difficulties do not establish that existing legal doctrines are necessarily incapable of responding to AI-driven harm. Rather, they raise questions about how those doctrines apply where trading decisions are increasingly automated and the relationship between human conduct and particular market outcomes becomes less direct. The detailed application of the statutory market manipulation framework, relevant case law and private law doctrines is therefore considered in Chapters 4 to 6. Chapter 7 then considers whether the limitations identified in those chapters justify additional regulatory responses, including AI Governance Liability and a bespoke product liability regime.

3.3 CONCLUSION

The increasing integration of AI and ML-driven algorithms into financial markets presents a complex array of risks that significantly threaten market integrity and challenge established legal frameworks. These systems' inherent speed and complexity introduce cyber vulnerabilities and operational failures, amplified by the interconnectedness of global trading platforms. Moreover, malicious actors can exploit AI to deliberately manipulate markets through sophisticated strategies and biased data. Various forms of algorithmic misconduct, including potential collusion, insider trading, and enhanced manipulative tactics, contribute to market inefficiency and erode public trust. The opacity of these “black box” systems further complicates the attribution of intent and causation, straining traditional legal concepts. Addressing these evolving risks requires constant adaptation of regulatory frameworks and technological capabilities worldwide to ensure accountability, maintain market integrity, and safeguard investor confidence in this rapidly changing landscape. The remainder of this thesis focuses on market manipulation because it most directly exposes the challenges that increasingly autonomous and opaque trading systems may present for existing legal rules. Chapter 4 examines the applicable statutory framework, Chapter 5 considers the relevant case law, Chapter 6 examines the application and limitations of private law responses, and Chapter 7 draws these findings together in an integrated regulatory framework, including the proposed AI Governance Liability tort and bespoke product liability regime.

---

[203] Azzutti, above n 75, at 4–5.

Chapter 4: Market Manipulation in New Zealand: The Regulatory Framework of the FMC Act

"A tailored approach that acknowledges the unique characteristics of AI is essential."
(Financial Markets Authority, 2024)

Chapter 4 examines whether New Zealand's existing statutory framework for market manipulation is capable of responding to the risks of AI-driven trading identified in Chapter 3. Building on the risk taxonomy developed there, this chapter undertakes a doctrinal analysis of the Financial Markets Conduct Act 2013 (FMC Act) to assess the extent to which its assumptions, concepts, and enforcement mechanisms remain effective in an environment characterised by high-speed, autonomous, and increasingly opaque trading systems. Section 4.1 establishes the definitional foundations of market manipulation under New Zealand law, providing the conceptual baseline against which AI-enabled manipulation must be assessed.

Section 4.2 then analyses the core statutory prohibitions on market manipulation, with specific attention to their application to algorithmic and AI-driven conduct, including whether concepts such as false or misleading appearance, knowledge, and intent can be coherently applied where manipulation is generated at the level of system design and operation rather than discrete human acts. Section 4.3 then identifies the civil and criminal consequences of contravening those provisions, including the statutory route to compensation. The doctrinal strengths and limitations identified in this chapter directly inform the analysis of judicial responses to algorithmic manipulation in Chapter 5, the assessment of the limits of private law liability in Chapter 6, and the case for the governance-based statutory and tortious reforms proposed in Chapter 7.

4.1. DEFINITION

The Financial Markets Conduct Act 2013 ("FMC Act") is the primary legislation governing the conduct of financial markets in New Zealand. The FMC Act's primary purposes are to ensure that businesses, investors, and consumers have the necessary information to

participate confidently in financial markets and to facilitate the development of fair, efficient, and transparent markets.[204] Moreover, the FMC Act empowers market participants by ensuring they receive timely, accurate, and precise information to make informed decisions about financial products and services. It also aims to establish strong governance practices for these products and services to minimise risk through effective monitoring while balancing the need for regulation with fostering innovation and flexibility within the financial markets.[205] Part 5 of the FMC Act governs dealing in financial products on markets, and subpart 3 prohibits market manipulation on licensed markets.[206]

Part 5, subpart 3 of the FMC Act addresses market manipulation, not by providing a single overarching definition, but by prohibiting specified forms of manipulative conduct.[207] Sections 262 and 265 address false or misleading statements or information and conduct that creates, or is likely to create, a false or misleading appearance of trading, supply, demand, price or value. The broader difficulty of defining market manipulation consistently has long been recognised in the scholarly literature.[208] These provisions are drafted to capture a range of conduct falling within their statutory elements.

However, the absence of a general catch-all prohibition of market manipulation within subpart 3 may present interpretive challenges in technologically advanced trading environments. Conduct that does not satisfy the elements of ss 262 or 265 may still engage other provisions of the FMC Act, including the fair-dealing provisions in Part 2. Nevertheless, economically distortive algorithmic conduct that does not fall within the statutory elements of the market manipulation provisions may expose limits in the application of subpart 3 to emerging forms of trading activity.

The FMA attempts to define market manipulation to benefit new or inexperienced investors. It describes market manipulation as a type of market abuse or misconduct in which someone aims to or attempts to mislead the market, such as the NZX or ASX, through actions or omissions that create a misleading appearance of trading activity, supply, demand, or the overall value of a financial product. According to the FMA, market manipulation also includes

---

[204] Financial Markets Conduct Act 2013, s 3.
[205] Financial Markets Conduct Act 2013, s 4.
[206] Financial Markets Conduct Act 2013, s 230.
[207] Financial Markets Conduct Act 2013, ss 262 and 265.
[208] Fischel and Ross, above n 61, at 510–520.

inflating or deflating market prices, affecting the market by spreading rumours or false information, or engaging in deceptive trade practices.[209]

The FMA's definition, dating back to 2020, serves as a valuable starting point as it can be broadly interpreted to encompass various traditional forms of market manipulation. It addresses misleading the market through actions or omissions that create a false appearance of trading activity, supply, demand, or value. While seemingly broad, its application to the complexities of modern, AI-driven manipulation reveals inherent limitations. For instance, spoofing and layering, executed by AI-driven algorithms in milliseconds, might not always be the result of a single actor deliberately misleading the market. Instead, they generate a pattern of fleeting orders and cancellations that, while ultimately distorting price discovery, can be argued to be simply the automated outcome of a trading program rather than a conscious 'action or omission' intended to mislead a specific individual.

Similarly, AI-driven quote stuffing, which floods the market with non-genuine orders, may be framed as an algorithmic inefficiency or a byproduct of HFT strategies, rather than a deliberate attempt to create a 'false appearance of trading activity' in the traditional sense, where a human actor disseminates misleading information or executes overtly deceptive trades. Because these actions occur at extreme speed and scale, it is often difficult to identify clear intent or link the misleading market signal to a specific human decision - both of which are typically central to establishing a contravention of the market manipulation provisions under the FMC Act.

Furthermore, AI's ability to exploit subtle market signals and engage in sophisticated momentum ignition strategies might not neatly fit the definition's emphasis on false information or rumours, or deceptive trades. These AI strategies often rely on identifying and amplifying genuine market movements through rapid, complex order execution, making it challenging to pinpoint any misleading actions or false information. The manipulation lies in the algorithmic amplification of market dynamics, which may not be easily captured by a definition focused on overt misrepresentation or deceptive trading in the conventional sense. Therefore, while the current FMA definition provides a foundational understanding of market

[209] Financial Markets Authority 'Market manipulation' (2020) <https://www.fma.govt.nz/library/articles/market-manipulation/>.

manipulation, its language and conceptual framework struggle to fully encompass the nuanced and rapidly evolving techniques of sophisticated AI in contemporary financial markets.

The current statutory definition of market manipulation does not explicitly account for the complexities introduced by algorithmic trading, particularly when autonomous or semi-autonomous systems execute strategies without direct human intervention. This raises doctrinal challenges in attributing mens rea to human actors involved in the design, deployment, or oversight of such systems. While machines themselves cannot possess mens rea, the difficulty lies in determining whether and to what extent the requisite intent can be imputed to the programmer, the deploying trader, or another party. The statutory focus on human actions or omissions reflects a narrowly framed conception of the necessary mens rea, which is poorly suited to manipulative conduct facilitated through algorithmic processes. This ambiguity risks creating enforcement gaps and interpretive uncertainty in cases of market abuse arising from the use of trading algorithms.

Moreover, the statutory emphasis on creating a 'false or misleading appearance' of trading activity, supply, demand, or value may not adequately encompass algorithmic strategies that distort market dynamics without generating overtly deceptive signals. Techniques such as latency arbitrage, predictive front-running, and order book layering can subtly undermine price formation and market integrity, even if they do not produce a clearly 'false' representation. In this sense, the difficulty is not a lack of guidance, but that the definition directs attention away from economically distortive conduct, thereby challenging the boundaries of market abuse and calling for a more technologically attuned regulatory response.

Similarly, in Australia, Section 1041A of the Corporations Act 2001 defines market manipulation as conduct that has, or is likely to have, the effect of creating or maintaining an "artificial price" for trading in various financial products on a financial market operating in the jurisdiction.[210] A central theme of this Section is the concept of an "artificial price". The core concept is that the price of a financial product is artificially inflated or deflated, disrupting the regular forces of supply and demand. Since there is no legal definition of 'artificial' or 'artificial price', Courts must assess each case based on the specific actions and circumstances involved.

[210] Corporations Act 2001 (Cth), s 1041A.

The primary inadequacy of this definition lies in the inherent ambiguity of the term "artificial price". Without a clear legal definition, determining what constitutes an "artificial" price becomes highly subjective, relying on complex economic analysis and expert testimony. This subjectivity creates uncertainty for market participants, potentially discouraging legitimate trading as the boundaries remain unclear. For instance, aggressive but legal trading strategies aimed at capitalising on market inefficiencies could be retrospectively deemed manipulative if they are found to have created a price that a court later considers "artificial". Furthermore, establishing the requisite causal link, proving that specific conduct caused the artificial price, can be exceptionally challenging, even in straightforward cases involving human traders and traditional manipulative tactics. Demonstrating this causal connection amidst the wide range of factors influencing market prices is a significant hurdle for regulators.

While the lack of a clear definition for "artificial price" is a fundamental weakness, the advent of AI and algorithmic trading amplifies this inadequacy by introducing new forms of manipulative conduct that are harder to detect, attribute, and link causally to artificial price formation under the existing definition's framework. The speed, complexity, and autonomous nature of AI-driven trading demand a more nuanced regulatory approach that goes beyond a potentially blunt focus on the outcome of an "artificial price" without clearly defining the prohibited conduct in the context of these new technologies.

Amending the current definition could involve broadening the scope of actors to include "automated trading systems," expanding the definition of manipulative conduct to encompass the "deployment, design, or operation" of such systems, and incorporating illustrative examples of algorithmic manipulation. Alternatively, introducing additional definitions could provide greater clarity by specifically addressing manipulation arising from the interaction of multiple algorithms and the "responsibilities for AI governance" within trading firms. Regardless of the chosen approach, key considerations must include ensuring technology neutrality where appropriate, maintaining clarity and precision in the definitions, ensuring enforceability in the context of complex algorithms, and carefully avoiding the inadvertent capture of legitimate trading strategies that contribute to market efficiency.

Adopting this collaborative and adaptive strategy offers numerous benefits for New Zealand's financial markets. Enhanced regulatory effectiveness empowers the FMA to detect and deter AI-driven market manipulation more effectively, thereby bolstering market integrity.

More precise definitions reduce regulatory uncertainty for market participants, fostering responsible innovation within a well-defined framework. In this respect, it is proposed to amend the definition of market manipulation as follows:

"*Market manipulation occurs when any person or automated trading system misleads (or attempts to mislead) the market by actions, omissions, deployment, design, or operation that give a false or distorted appearance of trading activity, supply, demand, or the value of financial products. It includes but is not limited to manipulative strategies such as disseminating false information, deceptive trading practices by individuals, or using algorithms to exploit market microstructure, amplify volatility, or generate artificial momentum. Responsibility for market manipulation may extend to individuals who directly engage in manipulative acts or omissions and those who develop, deploy, or oversee automated systems that result in manipulative outcomes.*"

Building upon the recognition that market manipulation can originate from both human actions and the operation of sophisticated algorithms, the above-revised definition seeks to establish a more encompassing and enduring regulatory framework. By explicitly including "any person or automated trading system" as potential actors and broadening the scope of manipulative conduct to encompass "actions, omissions, deployment, design, or operation," the definition aims to capture the diverse ways in which markets can be misled or an attempt to mislead can be made. Furthermore, by acknowledging "false or distorted" appearances and providing examples relevant to both traditional and algorithmic strategies, the definition seeks to remain relevant amid evolving market practices and technological advancements, ultimately fostering greater clarity and accountability across the spectrum of market participants and automated systems.

For this definition to operate effectively in the context of evolving trading technologies, the FMA must play an active supporting role by engaging in structured consultation with industry experts, particularly those specialising in algorithmic and high-frequency trading. Such collaboration can inform regulatory interpretation and support future legislative reform, ensuring that the definition remains responsive to technological developments and market realities. This engagement should prioritise developing a comprehensive understanding of evolving trading practices and the increasing levels of automation inherent in algorithmic trading. By establishing robust communication channels, such as workshops, consultation

papers, and ongoing industry forums, the FMA can draw on the specialised knowledge of technology providers, algorithmic trading firms, data scientists, market infrastructure providers, and legal professionals.

The primary objective of these collaborations should be to assess the current state of automation, identify emerging capabilities of artificial intelligence in trading strategies, and pinpoint novel algorithmic tactics that may fall outside the scope of traditional conceptions of market manipulation. Furthermore, understanding the potential for unintended market distortions arising from complex algorithms, even in the absence of malicious intent, is essential to shaping a nuanced and workable regulatory response. With a clearer understanding of the technological landscape, the FMA will be better placed to refine the application of the market manipulation definition or, where necessary, develop supplementary guidance or definitions tailored to the specific challenges posed by algorithmic and AI-driven trading.

4.2. PROHIBITION OF MANIPULATION

Part 5, Subpart 3 of the FMC Act prohibits market manipulation on licensed markets. However, the ascendance of sophisticated AI-driven trading necessitates a critical analysis of the ongoing effectiveness of these established provisions in deterring and prosecuting this novel form of market manipulation. AI algorithms can execute complex, high-speed trading strategies that can circumvent traditional detection methods and exploit market vulnerabilities in ways not fully anticipated by current law. Consequently, it is imperative to consider whether specific amendments, such as those addressing algorithmic design, responsibility for AI agents, and the interpretation of manipulative intent in non-human actors, are required to ensure that Part 5, Subpart 3 continues to provide a relevant and robust framework for maintaining market integrity in this rapidly evolving era of AI development within financial markets.[211]

First, section 262 of the FMC Act prohibits a person from making statements or sharing information about financial products where a material aspect of the statement or information is false or misleading, the person making or sharing the information knows or ought reasonably to have known this, and the statement or information is likely to induce trading in the financial product, affect its price, or influence how a person would exercise a voting right attached to

[211] Financial Markets Conduct Act 2013, s 230.

that product.[212] This section essentially shields investors from being misled by false or deceptive information that could manipulate trading decisions or the overall market price of financial products.

Information-based price manipulation involves disseminating false or negative information to artificially inflate (pump and dump') or deflate (poop and scoop') stock prices for personal gain. In pump-and-dump schemes, manipulators spread positive misinformation to drive up a stock's price, allowing them to sell their holdings at a profit before the price falls. Conversely, in poop-and-scoop tactics, negative rumours are circulated to depress prices, enabling the manipulator to buy shares cheaply before the price rebounds. These manipulative tactics can also proliferate on social media platforms like Facebook, Twitter, Instagram, and LinkedIn, where false statements can be posted to influence securities trading. An example would be a user posting false statements about specific companies or a particular market incident to influence the buying and selling of specific securities.[213]

Section 262 of the FMC Act is highly relevant to AI-driven market manipulation. As AI algorithms become increasingly sophisticated, their ability to generate and spread information, including market analysis, predictions, and even synthetic news, at unprecedented speeds and scales, presents a potential avenue for manipulative conduct. If AI-generated content contains materially false or misleading elements and is likely to influence trading decisions or market prices, it can fall under the ambit of this Section.

Furthermore, AI algorithms can be employed rapidly to amplify the reach and impact of misinformation originating from human sources, potentially exacerbating its detrimental effects on market participants and price stability. A developer, deployer or operator is not liable merely because of that role. Section 262 could apply only where the person made or disseminated the relevant statement or information, or is otherwise legally responsible for that conduct, and the statutory knowledge and likely-effect requirements are satisfied. Given that manipulative AI strategies aim to induce specific trading behaviours or artificially influence prices, section 262, which directly addresses information likely to produce such outcomes, remains a valuable legal instrument.

---

[212] Financial Markets Conduct Act 2013, s 262.
[213] Financial Markets Authority 'Market manipulation' (2020) <https://www.fma.govt.nz/news/all-releases/media-releases/formal-warning-market-manipulation/>.

This distinction is critical when considering liability for AI-generated misinformation, as the autonomous nature of algorithmic systems complicates the attribution of either actual knowledge or constructive awareness to human actors. However, applying section 262 to AI-driven market manipulation raises significant attribution challenges, particularly in identifying who should be held responsible for the conduct in question. The statutory requirement under section 262 hinges on establishing that a person had "knowledge" or "ought reasonably to have known" of the relevant conduct. While actual knowledge refers to a mental state, the "ought reasonably to have known" standard operates as a legal test grounded in foreseeability and reasonableness. Where manipulative conduct is executed by autonomous systems, attributing satisfaction of these standards to a natural person, such as a developer, operator, or trader, requires careful analysis of their involvement, oversight responsibilities, and the extent to which the algorithm's behaviour and outcomes were reasonably foreseeable.

Furthermore, determining what constitutes a "material aspect" of AI-generated information demands a sophisticated understanding of how algorithms process and disseminate data, and how market participants, including other algorithms, interpret that output. Emerging forms of algorithmically generated "information," such as latent patterns in high-frequency trading (HFT) data or sentiment scores derived from unconventional sources, may challenge the conventional scope of "statement or information" contemplated by section 262.

Addressing these limitations requires a multi-faceted approach. The FMA may provide interpretative guidance that explicitly outlines how section 262 applies to AI-generated content, clarifies the responsibilities of those overseeing AI systems, and offers examples of "material aspects" in algorithmic contexts. Depending on the evolving landscape, legislative amendments to the FMC Act might ultimately be necessary to explicitly address the unique challenges posed by AI in this context, potentially including specific obligations for those developing and deploying AI in financial markets to ensure the accuracy and non-misleading nature of their systems' outputs. By proactively implementing these measures, regulators can strive to maintain the relevance and effectiveness of section 262 in safeguarding market integrity in an increasingly sophisticated AI-driven trading era.

It is suggested that section 262 should apply not only to statements made by a person but also to information disseminated through automated trading systems. Crucially, it must define the "person" responsible for such information to include those who developed,

deployed, or have oversight of the AI system. To provide a framework for determining whether these responsible parties “ought reasonably to know” about the false or misleading nature of the information, the amendment would specify relevant factors such as:

(i) the due diligence undertaken in the design, testing, and deployment of the automated trading system,

(ii) the robustness of the system's data inputs and processing logic, and

(iii) the existence and effectiveness of monitoring and oversight mechanisms for the system’s outputs, and industry best practices and available technologies for preventing the dissemination of false or misleading information by automated systems.

Moreover, the amendment should empower the FMA to issue specific guidelines tailored to assess materiality and misleading nature in the context of information generated or disseminated by automated trading systems. Considering these suggestions, section 262 can be amended as follows:

“(1) A person must not make a statement or disseminate information, whether directly or through an automated trading system, if—
(a) a material aspect of the statement or information is false, or the statement or information is materially misleading; and
(b) the person (as defined in subsection (2)) knows or ought reasonably to know that a material aspect of the statement or information is false or that the statement or information is materially misleading; and
(c) the statement or information is likely to—
(i) induce a person to trade in quoted financial products; or
(ii) have the effect of increasing, reducing, maintaining, or stabilising the price for trading in those financial products; or
(iii) induce a person to exercise a voting right attached to a quoted financial product in a particular way or to abstain from exercising such a right.

(2) For the purposes of this Section, where information is disseminated through an automated trading system, the "person" responsible for that dissemination includes any person who developed, deployed or has oversight of that system.

(3) In determining whether a person responsible for information disseminated through an automated trading system "ought reasonably to know" that the information is false or materially misleading, factors to be considered include —
(a) the due diligence undertaken in the design, testing, and deployment of the automated trading system;
(b) the robustness and accuracy of the system's data inputs and processing logic;
(c) the existence and effectiveness of monitoring and oversight mechanisms for the system's outputs and industry best practices, and available technologies for preventing the dissemination of false or misleading information by automated systems.

(4) The Financial Markets Authority may issue guidelines specifying factors to be considered when assessing the materiality and misleading nature of information generated or disseminated by automated trading systems from time to time."

A key doctrinal challenge raised by the proposed amendment is whether a person who intentionally designs or deploys an algorithm to generate misleading information can be treated as having made or disseminated the resulting information, even where that person does not know the precise content of each output. Azzutti, Ringe and Stiehl distinguish between the conscious human use of an algorithm and autonomous AI decision-making and explain that black-box systems can complicate the application of conventional legal tests based on human intent and causation.[214]

To resolve this doctrinal gap, the offence itself should be framed so that, for the purposes of section 262, a person who knowingly designs, deploys, or operates an automated trading system with the intention that it produce misleading information is liable for engaging in that conduct. This approach grounds liability directly in intentional system-level behaviour. By focusing on the deliberate creation or deployment of a system intended to mislead, the regulatory framework can capture algorithmic manipulation without requiring proof that the person knew the precise content of the system's individual outputs.

Financial Markets Authority v Hill does not determine the application of ss 262 or 265. The proceeding was resolved under s 22(d) of the FMC Act and concerned anonymous

[214] Alessio Azzutti, Wolf-Georg Ringe and H Siegfried Stiehl "Machine Learning, Market Manipulation, and Collusion on Capital Markets: Why the 'Black Box' Matters" (2021) 43 *University of Pennsylvania Journal of International Law* 79 at 115–121.

promotional posts made by Matthew Geoffrey Hill on the Sharetrader online investor forum without disclosure of his position and financial interests in New Talisman Gold Mines. Although the FMA had originally alleged information-based market manipulation under s 262, that allegation was not pursued following settlement of the proceeding.[215]

Hill is therefore relevant here only by limited analogy. It demonstrates that undisclosed online promotion concerning a financial product may contravene the FMC Act's fair-dealing provisions.[216] It does not establish that AI-generated promotional content constitutes market manipulation or determine how the conduct and knowledge requirements in s 262 apply where content is produced or disseminated through an automated system. Those questions must be determined by reference to the statutory elements of s 262.

Section 263 of the FMC Act establishes a specific exception to the prohibitions in section 262, which concern the making or dissemination of false or misleading statements or information. This exception applies specifically to conduct related to takeover offers for financial products that are regulated under the Takeovers Code or the Takeovers Act 1993. The primary intention of this Section is to prevent regulatory overlapping by acknowledging that the Takeovers Code and Act provide a comprehensive and distinct framework governing the disclosure and conduct requirements during a takeover process. As such, for actions directly related to a takeover offer and regulated by these specific pieces of legislation, the general prohibitions of section 262 do not apply.

The direct relevance of this exception to AI-driven market manipulation is likely limited in most typical scenarios. Takeover offers are generally initiated and conducted by companies or individuals, not AI systems. While AI could be employed as a tool by these entities for various purposes during a takeover, such as analysing data or formulating strategies, the legal responsibility for the statements and information disseminated related to the offer rests with the human actors or the offering company, who are subject to the specific regulations of the Takeovers Code. The Code has its own detailed rules regarding required disclosures to shareholders, which are designed to ensure informed decision-making during a takeover.

---

[215] *Financial Markets Authority v Hill* [2024] NZHC 1353 at [1] – [6].
[216] *Financial Markets Authority v Hill* [2024] NZHC 1353 at [124].

Section 263 essentially prioritises these specific takeover regulations over the general prohibitions of Section 262 for conduct falling within the regulated takeover context.

However, there may be some indirect relevance of section 263 to AI-driven activities. For instance, while the exception might apply to the formal statements made in a takeover offer, AI could be used to generate or spread misleading information about the takeover outside those regulated documents, aiming to influence the target company's share price. In such instances, whether the exception in section 263 would apply could be a matter of interpretation, depending on the directness of the link between the AI-driven misinformation and the conduct regulated by the Takeovers Code. Ultimately, section 263 primarily functions to delineate regulatory responsibilities between the general prohibitions on misleading information and the specific rules governing takeover conduct, with its direct implications for AI-driven manipulation being largely indirect.

Section 265 prohibits a person from doing or omitting anything that has, or is likely to have, the effect of creating a false or misleading appearance concerning active trading in quoted financial products or their supply, demand, trading price or value. The person must know, or ought reasonably to know, that the act or omission will, or is likely to, have that effect. An example of such manipulative behaviour would be an individual placing orders to buy shares to give a false impression of high demand for the shares or selling pressure when they do not want or intend to purchase those shares. Another scenario could involve buying and selling the same financial products simultaneously across accounts controlled by the same person to create a false or misleading appearance of genuine market interest in the product or placing multiple orders at different prices. The global objective of this provision is to prevent market manipulation tactics that rely on creating a false impression of trading activity to influence prices and investor behaviour.

Understanding section 265 is essential in the face of potential market manipulation facilitated by AI. Unlike traditional trading, AI algorithms can generate large volumes of orders and cancellations at high speed, potentially creating the illusion of intense trading interest that does not reflect genuine market sentiment, directly contravening the principles of this section. Similarly, sophisticated algorithmic strategies can manipulate the order book by strategically placing and withdrawing limit orders, thereby distorting perceptions of actual supply and demand and artificially affecting prices. Furthermore, AI could be used to execute coordinated

trading strategies across related accounts to artificially inflate trading volumes and influence price perceptions, thereby undermining the integrity of market signals. The "ought reasonably to know" standard does not automatically extend liability to every developer, deployer or overseer of an AI trading system. Liability under section 265 requires identification of a person whose act or omission had, or was likely to have, the prohibited effect and who knew or ought reasonably to have known that the act or omission would, or was likely to, have that effect.[217]

However, applying section 265 to AI-driven manipulation also raises several complexities. The difficulty arises because attributing the requisite "knowledge" or that a person "ought reasonably to know" in the context of AI-driven trading is not merely a question of whether manipulative outcomes occur, but whether those outcomes were sufficiently foreseeable in light of the system's design, optimisation logic, and oversight structures. Recent scholarship on algorithmic and high-frequency trading underscores that the same technical features which enable legitimate strategies such as market making, liquidity provision, and adaptive order-routing can, under different market conditions, generate patterns that resemble manipulation, thereby complicating the evidential inquiry into intent and knowledge.[218] This creates a tension with the claim, developed in the preceding paragraph, that certain forms of AI-driven manipulation are increasingly predictable at the level of system architecture. If system-level risks are ex ante identifiable, it is necessary to explain why establishing knowledge remains challenging.

The challenge lies not in a conceptual inability to attribute responsibility, but in distinguishing ex post between outputs resulting from lawful algorithmic strategies and those reflecting intentional deception or reckless disregard of market integrity. As Armour and Eidenmüller argue in the context of algorithmic governance in financial markets, liability must be calibrated to capture those who design or deploy systems in ways that make harmful outcomes reasonably foreseeable, without imposing strict liability on all users of complex trading technologies.[219] Similarly, Hildebrandt's analysis of artificial agency and private law liability suggests that responsibility should remain attributable to existing persons or bodies, with the allocation of responsibility informed by control and by the duties associated with

[217] Financial Markets Conduct Act 2013, s 265.
[218] Yadav, above n 193, at 1085–1098.
[219] John Armour and Horst Eidenmüller "Self-Driving Corporations?" (2020) 10 *Harvard Business Law Review* 87 at 110–112.

selecting, operating, monitoring and maintaining the technology.[220] In practice, this implies that regulators and courts must assess whether the system's architecture, data inputs, and monitoring mechanisms rendered the manipulative outcome sufficiently foreseeable to satisfy the statutory threshold, while avoiding over-inclusive liability for legitimate algorithmic trading activity.

The core difficulty is that AI-driven trading systems operate at speeds and levels of computational complexity that fundamentally disrupt the causal and evidential links on which traditional market-manipulation doctrines rely. The problem is not simply that algorithms act quickly, but that their decision-making processes involve layers of optimisation, feedback loops, and adaptive responses that are opaque even to their designers. This makes it difficult to reconstruct, after the fact, how a particular trading pattern emerged, whether it was the product of legitimate high-frequency strategies or an engineered attempt to create a false or misleading appearance of market activity. This is due to the system's internal logic that may generate thousands of micro-decisions per second, each influenced by dynamic market inputs. Regulators face a significant attribution challenge: they must determine whether a specific manipulative effect can be traced to an operator's knowledge or intention, or whether it arose from the autonomous behaviour of the system itself.

Novel algorithmic strategies can generate trading patterns that fall outside the types of conduct contemplated by existing statutory definitions because their optimisation processes allow them to discover behaviours that human traders would not ordinarily devise or recognise as manipulative. The difficulty arises not from technological novelty in the abstract, but from the fact that machine-learning-based systems can identify profitable tactics such as patterns of rapid order placement, cancellation, or sequencing that create the appearance of liquidity or market interest without any human actor formulating that tactic in advance. A hypothetical illustrates the point: an adaptive execution algorithm trained to minimise transaction costs may learn that repeatedly submitting and withdrawing small orders in a thin market induces other traders to reveal their positions. Although no human operator intends to create a false impression of trading interest, the emergent pattern mimics classic spoofing by signalling liquidity that does not exist.

---

[220] Mireille Hildebrandt *Law for Computer Scientists and Other Folk* (Oxford University Press, Oxford, 2020) at 247–250.

Under the current legislative framework, such behaviour may be difficult to classify where the conduct and knowledge required by section 265 cannot readily be attributed to a particular person. This gap explains why the continuous evolution of algorithmic strategies challenges the adequacy of existing definitions: the statutory concept of a "false or misleading appearance" was drafted with human-directed manipulation in mind, whereas AI systems can autonomously generate novel market signals that achieve similar effects without fitting neatly into established doctrinal categories. A more nuanced definition is therefore required, as complex algorithmic systems can generate market signals through mechanisms that differ fundamentally from human-directed trading behaviour. In human-driven scenarios, a false or misleading appearance typically arises from a deliberate act such as placing an order, releasing information, or engaging in a coordinated pattern of trades that a person intends, or at least understands, will distort how the market perceives supply, demand, or liquidity.

To tackle AI-driven market manipulation, interpretative guidance from the FMA on the application of section 265 to algorithmic trading, with clear indicators of what constitutes a "false or misleading appearance" in this context and a clarification of the responsibilities of those deploying, developing or overseeing AI systems, is required. Regulatory efforts should also emphasise the implementation of robust design principles and internal controls for AI trading systems to proactively prevent the generation of deceptive market signals. Investing in advanced AI-powered surveillance tools capable of analysing HFT data to identify patterns indicative of manipulative algorithmic behaviour is crucial.

Regulatory adaptation also requires recognising that what counts as "misleading" in AI-mediated markets is no longer limited to information that would deceive a human trader but must also encompass inputs capable of distorting an AI system's own decision-making processes. Because contemporary trading algorithms rely on vast, rapidly updated data streams, even subtle distortions in those inputs can cause the system to generate trading patterns that create artificial impressions of liquidity, demand, or market sentiment. Strengthening regulatory capacity, therefore, depends on sustained collaboration with data scientists and AI specialists who can clarify how algorithmic models interpret information, how they may be manipulated, and how such distortions manifest in market outputs.

Structured experimentation through regulatory sandboxes, supervised testing environments, and controlled pilot programs would enable regulators to observe how AI

trading systems behave under different informational conditions and to identify the kinds of signals that can mislead the systems themselves. Insights from such initiatives would support the development of legislative amendments that explicitly address manipulative conduct executed through automated systems by focusing on the effects of algorithmic behaviour, rather than relying solely on traditional notions of human intent. An example of how section 265 can be amended is as follows:

"(1) A person must not engage in any conduct (including an act or omission, whether carried out directly or through an automated trading system) if that conduct has, or is likely to have, the effect of creating, or causing the creation of, a false or misleading appearance -
(a) with respect to the extent of active trading in quoted financial products or
(b) with respect to the supply of, demand for, price for trading in, or value of those financial products.

(2) For the purposes of subsection (1), where an automated trading system carries out the conduct, responsibility will be attributed as follows based on what triggered the manipulative outcome:
(a) the person who designed the core programming logic of the automated trading system may be held responsible if the false or misleading appearance was a reasonably foreseeable outcome of that design. Factors to consider include the system's initial specifications, testing procedures, and inherent biases or limitations.
(b) the person who set the specific parameters, provided instructions, or input data into the automated trading system may be held responsible if the false or misleading appearance was a reasonably foreseeable outcome of those actions. Factors to consider include trading limits, the specific algorithms selected, the data sources used, and any manual overrides or interventions.
(c) the person who deployed or uses the automated trading system, or has a responsibility to oversee its operation, may also be held responsible if their failure to adequately monitor and control the system, or to exercise reasonable care and diligence in their oversight, directly contributed to the false or misleading appearance.

(3) (a) Liability in subsection (1) turns on whether the person's conduct made the misleading appearance reasonably foreseeable, rather than on whether they subjectively intended that outcome. For conduct carried out directly by an individual, this means that liability arises only where it can be shown that the person knew, or that a reasonable person in their position ought

to have known, that their actions were likely to produce a false or misleading appearance in the market.
(b) For conduct carried out by an automated trading system and for the attribution of responsibility under subsection (2), it is not necessary to prove that the responsible person had the specific intention to create a false or misleading appearance through each individual action of the system. However, in determining responsibility, the factors outlined in subsection (2)(a), (b) and (c) would be applicable.

(4) Subsection (1) continues to apply to traditional human-driven market manipulation, and the standard for such conduct is governed by subsection (3)(a).”

The proposed amendments to section 265 aim to modernise its application to both traditional and AI-driven market manipulation. Subsection (1) broadens the prohibition to cover any direct or automated conduct that creates a false or misleading appearance in the trading of financial products. Subsection (2) specifically addresses responsibility for actions carried out by automated trading systems, outlining how liability can be attributed to the system's developer based on its design, to the deployer or user based on their input and instructions, or to those who fail to adequately oversee the system's operation, depending on what triggered the manipulative outcome.

Subsection (3) establishes the required mental state for liability, retaining the traditional “knows or ought reasonably to know” standard for direct human actions. AI-driven conduct shifts the focus to the reasonably foreseeable outcomes of the system's design, deployment, and control, without requiring proof of specific intent for each automated action, while still allowing consideration of factors such as design, parameters, and oversight. Finally, subsection (4) confirms the continued applicability of the section to traditional human-driven market manipulation under the requirement that the person knows or ought reasonably to know that the person’s act or omission will, or is likely to have, that effect.

Section 266 of the FMC Act provides that short selling and crossings under the business rules of a licensed market do not contravene section 265 merely by reason of being short selling or crossings. This exception acknowledges the legitimate functions of these standard trading practices within a regulated framework: short selling, defined in section 266(2) as selling a financial product not currently owned but intended to be acquired before settlement through

borrowing or purchase, serves purposes like price discovery and hedging; while crossings, also defined in section 266(2) as a single entity acting as both buyer and seller (often in an agency capacity representing another party to facilitate transactions like large block trades), promote efficient trading. Subject to compliance with the licensed market's business rules, these activities do not contravene section 265 merely by reason of being short selling or crossings.[221]

The increasing use of sophisticated AI algorithms to execute trading strategies, including short selling and crossings, brings a critical dimension to the relevance of section 266 in the context of market manipulation. While legitimate AI-driven short selling and crossing within market rules would fall under the exemption provided by this section, the potential for AI to exploit these exceptions for manipulative purposes is a significant concern. Highly advanced AI could be programmed to conduct a high volume or specific pattern of short sales or crossings, technically within market regulations, but with the primary intention of creating a false impression of market activity or price pressure to benefit the AI's trading positions. In such scenarios, the question arises whether these actions remain protected by section 266, given that the manipulative intent goes beyond the inherent nature of short selling or crossing.

The interpretation of the phrase 'merely by reason of being short selling or crossings' is crucial when considering AI-driven manipulation. If an AI system layers manipulative strategies on top of these activities, the exemption might no longer be applicable. For instance, an AI could coordinate its short-selling activities by disseminating subtly misleading information to amplify the intended manipulative effect. While the short selling itself might be exempt under section 266, the combined effect orchestrated by the AI could constitute a breach of section 262 that transcends the scope of the exception. Regulatory scrutiny will need to focus on distinguishing whether AI-driven short selling and crossing are for legitimate trading within market rules or are being used as tools with the primary aim of market manipulation. The challenges of attributing manipulative intent to the person behind the AI, as discussed in relation to sections 262 and 265, will also be pertinent in determining the applicability of section 266's exceptions in the context of AI-driven trading. To better address the complexities of AI-driven manipulation that could exploit the existing exceptions for short selling and crossings, it is suggested that section 266 be amended as follows:

---

[221] Financial Markets Conduct Act 2013, ss 265 and 266.

"(1) Short selling and crossings under the business rules of a licensed market do not contravene section 265 merely by reason of being short selling or crossings, provided that the primary purpose and effect of such activities are for legitimate trading purposes and not the creation of a false or misleading appearance.

(2) In this Section, -
crossing, in relation to a transaction in financial products, means a transaction where a person acts as —
(a) buyer and seller for that transaction in an agency capacity; or
(b) buyer or seller on one side of that transaction in an agency capacity and as a principal on the other side
short selling means a sale of any financial product where, at the time of the sale, —
(a) the seller does not have a presently exercisable and unconditional right to vest the financial product in the buyer except where the seller has an unconditional agreement to acquire that right before the date required to settle the sale; or
(b) the financial product being sold has been borrowed, and the seller has a presently exercisable and unconditional right to vest the product in the buyer.

(3) For the purposes of subsection (1), when an automated trading system executes short selling or crossings, the assessment of whether the primary purpose and effect are for legitimate trading purposes will consider the design, programming logic, parameters, and oversight of the automated trading system, as well as the reasonably foreseeable outcomes of its operation.

(4) The Financial Markets Authority may issue guidelines specifying factors to be considered when determining whether short selling or crossings executed by automated trading systems fall within the exception in subsection (1), particularly in relation to assessing the primary purpose and effect of such activities".

Section 267 of the FMC Act plays a crucial role in addressing AI-driven market manipulation, particularly tactics such as wash trading and the execution of matched orders, which can create a false or misleading appearance of trading activity prohibited under section 265. Subsection (1) directly targets transactions known as "wash trades," where a person trades financial products of a listed issuer without any actual change in beneficial ownership, such as

buying and selling the same security back and forth between accounts they control. This activity likely creates a misleading appearance of active trading. Consequently, subsection (1) stipulates that a person is automatically considered to have contravened section 265 if they are directly or indirectly involved in such trading. This provision is particularly relevant to AI-driven market manipulation, as algorithms can be readily programmed to execute these round-trip trades, artificially inflating trading volume without genuine market interest. The "deemed contravention" mechanism simplifies enforcement against AI-driven wash trading by focusing on the occurrence of the trading outcome itself, without requiring proof of specific manipulative intent or foreseeability in each instance.

While section 267 provides a strong mechanism against specific manipulative behaviours like wash trading and matched orders that AI could easily facilitate, the exception outlined in subsection (2) for wash trading warrants consideration in the AI context. This exception applies if a person acted on behalf of another and was unaware that no beneficial ownership change would occur. However, this exception is less likely to shield those deploying AI programmed to execute wash trades, as the controlling entity would typically be aware of the system's design and trading activities. Ultimately, section 267 serves as a significant tool in the legal framework to address specific forms of AI-driven market manipulation by establishing clear rules based on the structure and outcomes of trading. However, its effectiveness in comprehensively addressing all potential avenues of AI-driven manipulation likely depends on its interpretation and application alongside the broader provisions of section 265 and any amendments designed to address the unique challenges posed by sophisticated automated trading systems, particularly regarding the attribution of responsibility and the demonstration of manipulative intent. [222]

Section 267(3) of the FMC Act highlights another tactic of transaction-based manipulation that contravenes section 265. It applies when a person (A) places an order to buy or sell a financial product of a listed issuer. Either A or someone closely connected to them (referred to as an "associated person," defined in section 12 of the FMC Act as someone with a close familial or corporate connection, including those related by blood or marriage, connected through corporate structures, acting in coordination, following each other's lead, sharing mutual associations, or having significant influence over each other's actions) places

[222] Financial Markets Conduct Act 2013, s 267.

an opposite order, such as to buy if A wants to sell, and vice versa, with substantially matching details in terms of the number and price of financial products. This matched order will likely create an artificial impression of solid buying and selling pressure.

Subsection (3) directly addresses the issue of matched orders, where a person makes an offer to trade, and they or an associated person (potentially also controlled by AI or the same entity) make a substantially matching opposite offer in terms of quantity and price. This tactic can create a deceptive impression of supply and demand at particular price points. Given AI's capability to generate and execute precise, coordinated orders at high speed and scale, this subsection directly applies to AI-driven manipulation through matched orders. If an AI system is found to be engaging in this behaviour, the controlling entity will be deemed to have contravened Section 265. The provision's reference to "an associated person" is also pertinent, as a single actor could deploy multiple AI agents or control-related accounts to carry out such coordinated trading strategies.[223]

Subsection (4), Section 267, identifies specific actions deemed to create a false or misleading appearance of trading under Section 265 (like wash trades and matched orders), which is subject to the defence outlined in Section 268. According to Section 268, a person (A) accused of such an infringement can defend themselves, even in the context of AI-driven trading, by proving that the algorithmic trading activity or offer (i) was conducted in compliance with established and generally accepted practices in financial markets and (ii) had a proper purpose.

For instance, legitimate algorithmic execution strategies that incidentally result in no change of beneficial ownership during specific automated rebalancing processes might be defensible if they adhere to accepted market practices and serve a proper, non-manipulative purpose. However, determining "accepted market practices" and "proper purpose" for complex AI algorithms will require careful consideration of the algorithm's design, its intended function, and the prevailing norms of algorithmic trading, emphasising that the mere use of AI does not automatically negate manipulative intent if the primary aim was to distort the market.[224] It is

[223] Financial Markets Conduct Act 2013, s 12; *Financial Markets Authority v Wei Zhong* [2023] NZHC 766 at [56] – [58].
[224] Financial Markets Conduct Act 2013, s 268.

suggested that Sections 267 and 268 be amended to better account for the nuances of AI-driven trading and manipulation, particularly concerning wash trades and matched orders, as follows:

“Section 267

(1) A person (A) must be treated as contravening Section 265 if A is directly or indirectly a party to trading in the quoted financial products of a listed issuer from which no change in beneficial ownership results. This includes trading executed by an automated trading system controlled by A.

(2) Subsection (1) does not apply if—
(a) in trading the financial products, A was acting on behalf of another person, and
(b) A did not know and ought not reasonably to have known (including through reasonable oversight of any automated trading system deployed) that no change in beneficial ownership would result when trading the financial products.

(3) A person (A) must be treated as contravening Section 265 if—
(a) A has made an offer to trade the financial products of a listed issuer; and
(b) either A or, to A’s knowledge (including the known or reasonably foreseeable behaviour of an associated automated trading system), an associated person of A has made or proposes to make an opposite offer (the opposite offer) to trade financial products of the listed issuer; and
(c) The opposite offer substantially matches A’s offer regarding the number and price of the financial products.

(4) For the purposes of this Section, in the context of trading by automated trading systems, the “person” responsible includes the developer, deployer, and any person with oversight responsibilities for the system. The design and reasonably foreseeable outcomes of the automated trading system's operation may be considered when determining whether a contravention occurred.

(5) This Section is subject to Section 268.

Section 268

(1) In any proceeding against a person (A) for contravention of Section 265, it is a defence if the trading in the financial products occurred or the offer to trade was made, —

(a) in conformity with —

(i) established and generally accepted practices in financial markets and the trading activity, and

(ii) in the case of trading or offers executed by automated trading systems, the design, parameters, and operational logic of the system were consistent with prudent industry standards for algorithmic trading and did not have the purpose or reasonably foreseeable effect of creating a false or misleading appearance of trading; and

(b) for a proper purpose, in the context of automated trading systems, the algorithm was designed and deployed for legitimate trading objectives (such as genuine price discovery, hedging, or efficient order execution) and not primarily to manipulate the market or mislead other participants.

(2) When determining whether the defence in subsection (1) applies to trading or offers executed by automated trading systems, the court may consider factors such as —

(a) the documented design objectives and intended functionality of the algorithm;

(b) the due diligence undertaken in the development, testing, and deployment of the algorithm to prevent manipulative behaviour;

(c) the existence and effectiveness of monitoring and control mechanisms to detect and prevent the algorithm from creating a false or misleading appearance of trading;

(d) whether the person responsible for the automated trading system took reasonable steps to modify or cease its operation upon becoming aware of trading patterns that could be construed as creating a false or misleading appearance; and

(e) relevant industry best practices and regulatory guidance concerning algorithmic trading."

Section 545 of the FMC Act permits regulations, made by Order in Council on the recommendation of the Minister, to exempt conduct from being market manipulation that would otherwise fall within section 262 or 265. Any such exemption may be subject to terms and conditions.[225] This regulation-making power may therefore provide flexibility to address particular forms of algorithmic trading where an exemption is considered appropriate.

[225] Financial Markets Conduct Act 2013, s 545.

However, section 545 does not in itself prescribe accepted algorithmic trading practices or a proper-purpose test; any relevant criteria would need to be specified in the regulations.

4.3. CONSEQUENCES OF BREACH

Sections 262 and 265 are Part 5 market provisions under s 385(3)(b) of the Financial Markets Conduct Act 2013 (FMC Act) and are therefore civil liability provisions for the purposes of subpart 3 of Part 8.[226] On the application of the FMA or any other person, the High Court may make a declaration of contravention under s 486 if it is satisfied that a person has contravened one of those provisions. Under s 487, the declaration may be relied on as conclusive evidence of the matters stated in it in a subsequent application for a civil liability order.[227]

The FMA may apply for a pecuniary penalty order under s 489. Under s 490, the maximum pecuniary penalty is the greatest of the consideration for the relevant transaction, three times the gain made or loss avoided where that amount can be readily ascertained, or NZ$1 million for an individual and NZ$5 million in any other case.[228] Separately, under s 494, the FMA or any other person may apply for a compensatory order where an aggrieved person has suffered, or is likely to suffer, loss or damage because of the contravention. Section 495 permits the Court to make any order it considers just to compensate that person or to prevent or reduce the loss or damage.[229]

Criminal liability is addressed separately. A contravention of s 262 may constitute an offence under s 264 where the person knows that a material aspect of the statement or information is false or that the statement or information is materially misleading. A contravention of s 265 may constitute an offence under s 269 where the person knows that the act or omission will, or is likely to, create the prohibited false or misleading appearance. In either case, the maximum penalty for an individual is five years' imprisonment, a fine not exceeding NZ$500,000, or both. In any other case, the maximum fine is NZ$2.5 million.[230]

---

[226] Financial Markets Conduct Act 2013, ss 385(2)–(3)(b) and 485(1).
[227] Financial Markets Conduct Act 2013, ss 486(1) and 487.
[228] Financial Markets Conduct Act 2013, ss 489(1) and 490(1).
[229] Financial Markets Conduct Act 2013, ss 494 and 495.
[230] Financial Markets Conduct Act 2013, ss 264 and 269.

The FMC Act therefore provides a statutory route to compensation as well as public enforcement. In an AI-driven case, the practical difficulty may lie in establishing the contravention and proving that the relevant loss or damage resulted from it, rather than in the absence of a statutory compensatory remedy. The private law and governance reforms considered in Chapters 6 and 7 should accordingly be understood as potential supplements to, rather than substitutes for, the remedies already available under the FMC Act.

4.4. CONCLUSION

This chapter has demonstrated that New Zealand's market manipulation regime under the FMC Act remains largely effective in addressing traditional, human-led forms of misconduct, but is increasingly strained when applied to contemporary AI-driven trading practices. The statutory framework is grounded in assumptions of identifiable human actors, discrete acts or omissions, and actual or constructive knowledge. While these assumptions align well with conventional manipulation techniques such as false statements, wash trading, or matched orders executed by individuals, they sit uneasily with the realities of algorithmic and high-frequency trading, where manipulation may arise through speed, scale, and system-level behaviour rather than deliberate human deception.

The analysis of the definitional foundations of market manipulation reveals a significant doctrinal gap. Neither the FMC Act nor the FMA's guidance provides a comprehensive conceptual framework that fully captures how autonomous or semi-autonomous trading systems can distort market integrity. Existing definitions focus on creating a "false or misleading appearance" or disseminating false information. Yet, many AI-driven strategies operate by amplifying genuine market signals, exploiting microstructural features, or generating emergent trading patterns that undermine price formation without producing overtly false representations. As a result, economically harmful conduct may evade regulatory scrutiny, not because it is benign, but because it does not fit neatly within categories designed for human-directed manipulation.

The chapter further shows that these difficulties are not merely technological, but deeply doctrinal. The attribution of actual or constructive knowledge under sections 262 and 265 becomes increasingly complex where outcomes are generated by adaptive systems whose internal logic is opaque even to their designers. While the FMC Act's reliance on standards

such as “knows or ought reasonably to know” offers some flexibility, its application to AI-driven conduct exposes tensions between foreseeability, responsibility, and fairness. Holding individuals liable for every harmful output of a complex algorithm risks over-inclusive regulation yet failing to attribute responsibility where system-level risks were reasonably foreseeable creates enforcement gaps that threaten market integrity.

The comparative discussion of Australia’s “artificial price” concept underscores that outcome-based definitions alone are insufficient in the AI context. Ambiguity about what constitutes an “artificial” price, coupled with the evidentiary difficulty of establishing causation in fast-moving, multi-factor markets, highlights the limitations of frameworks that focus primarily on effects rather than on the design, deployment, and governance of trading systems. AI exacerbates these weaknesses by accelerating market interactions and obscuring causal chains, making post hoc attribution increasingly difficult.

In response to these challenges, the chapter argues for recalibrating New Zealand’s market manipulation regime to better reflect the realities of modern trading. This includes explicitly recognising automated trading systems as potential sources of manipulation; expanding the scope of relevant conduct to include system design, parameter setting, deployment, and oversight; and grounding liability more firmly in intention and reasonable foreseeability at the system level rather than in precise control over individual outputs. The proposed amendments to sections 262, 265, 266, and 267 illustrate how the FMC Act could evolve to capture AI-driven manipulation while preserving space for legitimate algorithmic trading that contributes to market efficiency.

Crucially, legislative reform alone is insufficient. The chapter emphasises the central role of the FMA in supporting any revised framework through interpretative guidance, technical engagement, and sustained collaboration with industry experts, data scientists, and market infrastructure providers. Regulatory sandboxes, supervisory experimentation, and advanced surveillance capabilities are essential to understanding how algorithmic strategies function in practice and how they may mislead both human and machine market participants. Such engagement enables regulators to distinguish harmful manipulation from lawful innovation and to respond dynamically to emerging risks.

Ultimately, this chapter concludes that the challenge posed by AI-driven market manipulation is not one of regulatory obsolescence, but of regulatory adaptation. By re-orienting market manipulation law toward system-level responsibility, foreseeable risk, and technological realism, New Zealand can preserve the core objectives of the FMC Act - fairness, efficiency, and transparency while ensuring that its markets remain resilient in the face of increasingly autonomous trading technologies. More precise definitions, principled attribution of responsibility, and active regulatory stewardship together offer a path toward a market manipulation regime that is both robust and future-proof.

Chapter 5: Case Law Trends in Addressing Modern Market Manipulation

5.1. INTRODUCTION

The previous chapter examined the statutory framework governing market manipulation in New Zealand, focusing on the Financial Markets Conduct Act 2013 ("FMC Act"). This chapter builds directly on that analysis by examining how courts and regulators have applied market manipulation law to conduct involving electronic and algorithmic trading. The cases are analysed through four recurring themes: the objective effect of trading conduct on the market; the role of purpose and knowledge; the attribution of automated conduct to human or corporate actors; and the evidential difficulty of distinguishing legitimate algorithmic trading from market manipulation.

The chapter first considers the developing New Zealand authorities under s 11B of the Securities Markets Act 1988 and s 265 of the FMC Act. It then considers selected international judgments and enforcement outcomes to identify how comparable legal principles have been applied to increasingly automated forms of trading. In doing so, the analysis distinguishes between contested judgments, consent orders, administrative settlements, regulatory final notices and unresolved proceedings, recognising that these sources do not carry the same precedential weight. Although the cases considered largely concern human-directed or conventional algorithmic trading rather than autonomous AI, they provide a basis for examining how existing principles concerning market effect, purpose, knowledge and attribution may apply where trading is increasingly automated.

The analysis does not assume that these authorities establish liability for autonomous AI-driven market manipulation. Rather, it identifies the doctrinal and evidential questions that arise where manipulative market effects may be observable, but the relevant human purpose, knowledge or responsibility is more difficult to establish. These questions provide the link between the statutory analysis in Chapter 4 and the remaining chapters of the thesis. Chapter 6 considers whether existing private law doctrines can respond to harm arising from AI-driven

trading, while Chapter 7 draws together the statutory, case law and private law findings to assess whether additional governance-based regulatory responses are justified.

5.2. NEW ZEALAND

While New Zealand jurisprudence currently lacks specific case law directly addressing market manipulation perpetrated by sophisticated AI trading algorithms, analysing existing precedents concerning traditional forms of market manipulation remains valuable for understanding how established legal principles might be applied to AI's unique characteristics. *FMA v Brian Henry* is the first market manipulation case (transaction-based) brought by the FMA under Section 11B of the now-repealed Securities Markets Act 1988. The FMA's claims stemmed from a referral by the NZX regarding Henry's trading activities in shares of the NZX-listed company, Diligent Board Member Services ('Diligent'). Henry had direct access to the NZX trading system through online platforms provided by ASB and Direct Broking Limited. Henry traded in the Diligent shares with himself ("wash trades"), influencing the share price without changing the beneficial ownership.[231]

In addition, Henry created a false and misleading appearance of trading in Diligent shares on four occasions by placing multiple buy and sell orders without completing the trades (referred to as "layering"). It gave an artificial impression of trading interest, forcing other buyers to bid higher and setting the market closing price of Diligent shares. He continued these activities despite warnings from ASB. It was established that Henry should have reasonably known that his actions would likely create a false or misleading appearance regarding the extent of active trading, supply, demand, price, or value of Diligent's shares.[232] Henry's misconduct constituted market manipulation, leading to a false supply-and-demand imbalance in shares and undermining the integrity and trust in the fair and orderly operation of the equities market.[233]

The *Henry* case, decided under the SMA, demonstrates that market manipulation is not confined to completed transactions. The placement of an order may itself create a false or misleading appearance concerning trading, supply, demand, price or value.[234] The case also

[231] *Financial Markets Authority v Brian Peter Henry* [2014] NZHC 1853 at [13]– [21].
[232] *Henry,* above n 231, at [33]–[39].
[233] *Henry*, above n 231, at [33]–[39] and [70]–[78].
[234] *Henry*, above n 231, at [33]–[39] and [70]–[78].

demonstrates the importance of the statutory knowledge requirement, which may be satisfied where the defendant knew or ought reasonably to have known that the conduct would or was likely to have the prohibited effect.

Applied to algorithmic trading, *Henry* supports the narrower proposition that an order generated through an algorithmic trading system may fall within the statutory prohibition even if it is not executed, provided the relevant requirements concerning market effect and knowledge are established. The case does not determine how conduct generated by an autonomous AI system should be attributed to a developer, deployer or operator. That attribution question is considered further in the later chapters of this thesis.

FMA v Warminger is another case where Mr Warminger, while employed by Milford Asset Management Limited ("Milford"), was found to have contravened s 11B of the Securities Markets Act 1988 by creating deceptive appearances of demand and price for Fisher & Paykel Healthcare Limited ("FPH") and a2 Milk Company Limited ("ATM") shares. His trading orders increased the offer quote and price for FPH and ATM shares and maintained them at higher levels than they otherwise would have been by entering a crossing to sell these shares. This conduct created a misleading appearance regarding the supply, demand, price, or value of FPH and ATM shares. Mr Warminger tried to explain that when he engaged in the trades, he thought he was trading against an algorithm programmed to place sell orders in the market. He later admitted that it could not have been an algorithm, as the time differences were not constant.[235]

It was also concluded that the transaction could not have been genuine, because only one party to the trade was a genuine buyer or seller. That said, if the trader's purpose is to manipulate the market, the trade can constitute market manipulation even if the other party is a genuine buyer or seller. It was accepted that, as a very experienced portfolio manager and trader with knowledge of the buying interest for large volumes, Mr Warminger reasonably ought to have known the effect his actions would have on the market. The Court held that Mr Warminger's actions lacked any genuine commercial purpose and manipulated the market in relation to the FPH and ATM shares, on the balance of probabilities.[236]

---

[235] *Financial Markets Authority v Warminger (No 3)* [2017] NZHC 327 at [44] – [51].
[236] *Warminger (No 3)*, above n 235, at [73] – [80].

The *Warminger* case has significant implications for AI-driven trading and manipulation, particularly in highlighting the types of conduct likely to be viewed as market manipulation, regardless of whether executed by a human or an algorithm. The case establishes that creating a false or misleading appearance of trading activity, supply, demand, or price is prohibited. It directly applies to AI strategies that could mimic Warminger's tactics, such as algorithms designed to execute small "feeder" trades to induce price movements before larger trades, or those programmed to manipulate closing prices to secure advantageous positions.

Furthermore, *Warminger* confirms that the statutory prohibition is principally effects-based. The FMA was not required to prove that Mr Warminger had a specific intent or purpose to create a false or misleading appearance. However, the Court recognised that the purpose of the trader may nevertheless be relevant when determining whether the conduct created a false or misleading appearance. Venning J observed that the trader's purpose may be the key factor distinguishing culpable manipulation from a genuine trade.[237] Accordingly, specific manipulative intent or purpose is not an independent element of liability, although the trader's purpose remains relevant evidence in determining whether the conduct created a false or misleading appearance.

The relevance of *Warminger* to AI-driven trading is therefore that liability does not depend upon attributing a specific manipulative intention to the technology itself. However, market effects alone are not sufficient, because the statutory knowledge requirement remains applicable to the relevant human or corporate actor. Where trading is increasingly automated, the more difficult question is how that knowledge and responsibility can be attributed to the persons responsible for designing, deploying or controlling the trading system.

FMA v Wei Zhong & Ors ("Wei Zhong") is a significant case in New Zealand history in which proceedings were brought under the FMC Act against four former executives of Oceania Natural Limited for market manipulation and failure to meet disclosure standards. The defendants were alleged to have executed trades which resulted in no genuine change in ownership and a misleading appearance of trading ("wash trades"), placed buy and sell orders at the same price and volume, resulting in a false impression of market demand ("matched

[237] *Warminger (No 3)*, above n 235, at [60] and [62]–[63].

orders") and traded close to the end of the trading day to manipulate the closing price of specific securities ("marking the close"). These actions misled other investors about the securities' true value and liquidity. The NZX referred the trading to the FMA in December 2016.[238]

The former key personnel of Oceania Natural Limited ("ONL") involved in this case were Mr Zhong (executive chairman & CEO), Ms Ding (marketing director), Mr Meng (previous director) and Mr Qian (chartered accountant). In June 2016, Ms Ding attempted to sell ONL shares using a securities account under her mother's name. Simultaneously, Mr Qian placed buy and sell orders through the Qian/Lu Securities account. Due to technical issues at ASB Securities, one of Ding's sell orders and Qian's buy orders did not match. Had they matched, it would have led to a trade that increased the price of ONL shares. Other aggressive and insignificant buy orders were placed, ultimately inflating the ONL share prices. Similarly, in October 2016, the parties used different securities accounts to layer four sell orders to create a false impression of trading interest, aggressive buy orders that did not aim for the best price, and insignificant order volumes. Whilst admitting to market manipulation, Meng and Qian acknowledged that they should have reasonably known these orders were likely to create a misleading appearance regarding the level of active trading, supply, demand, or the price of ONL shares.[239]

In April 2023, the High Court dealt with the defended market manipulation proceedings against Ms Ding, Mr Zhong and Mr Qian. The Court's approach was to first address the onus and the standard of proof. It emphasised that the usual Court rules and procedures for civil cases apply, including the standard of proof, meaning the FMA had to prove the allegations of market manipulation on a balance of probabilities. In *FMA v Warminger,* the Court noted that compared to criminal standards, the civil standard required evidence of severe misconduct but required less certainty. Under s265 of the FMC Act, the FMA had the burden of establishing on a balance of probabilities: (i) the defendants placed orders for ONL shares into the trading system (ii) those orders must have had, or were likely to have had, the effect of creating a false or misleading appearance regarding the extent of active trading, supply, demand, price, or value of ONL shares and (iii) the defendants knew or ought reasonably to have known that their orders were likely to have this effect.[240]

---

[238] *Financial Markets Authority v Wei Zhong* [2022] NZHC 480 at [11] – [20].
[239] *Wei Zhong (2022),* above n 238, at [21] – [32].
[240] *Financial Markets Authority v Wei Zhong & Ors* [2023] NZHC 766 at [30] – [39].

The Court highlighted that placing an order to buy or sell shares can constitute market manipulation, even if no actual trade occurs. Such orders can affect share prices and mislead others about the actual supply and demand dynamics for a particular stock. The FMA's claims were based on the likely and exact effects of market manipulation. The Court adopted the approach in Warminger that, in the context of s 265, "likely" means that the prohibited effect "might well happen", rather than that it was more likely than not to occur.[241] Regarding the term 'effect', it was observed that a false or misleading appearance can be created directly or indirectly. The mere fact that one party is a genuine buyer or seller does not imply the transaction is authentic. If the intention in offering to sell shares is to manipulate the market, the trade can still be considered manipulative, even if the purchaser genuinely offered to sell shares. Expert evidence was tendered to the effect that a share's liquidity, or the frequency at which it is being traded, influences investment decisions. The rule is that more actively traded shares are considered better investments, which affects their demand, supply, and prices.[242]

The Court clarified the test for liability regarding the knowledge element of market manipulation by stating that the FMA did not need to prove that the defendants acted with malicious intent or a specific manipulative purpose.[243] Instead, the requisite knowledge that the defendants knew or ought reasonably to have known their actions would create a false or misleading appearance of trading could be inferred from the evidence. This inference could arise directly from the defendants' trading conduct or through the application of constructive knowledge, meaning their actions would be judged against what a reasonable person in their position should have understood. While proof of malicious intent was not a prerequisite for liability, the Court noted that understanding the purpose behind the defendants' trades remained relevant in determining whether the conduct was indeed manipulative, as a trader's intention in placing orders could be a key factor in distinguishing between legitimate trading and actions intended to create a false impression of market activity.

Furthermore, a trader's purpose can be relevant in ascertaining their actual or perceived knowledge.[244] The expert evidence also addressed the relevance of whether the trading had a legitimate economic purpose. Mr Zhong and Ms Ding were found to have contravened ss 265

---

[241] *Wei Zhong* (2023), above n 240, at [44]–[52]; *Warminger (No 3)*, above n 235, at [39]–[41].
[242] *Wei Zhong (2023),* above n 240, at [53] – [55].
[243] *Wei Zhong (2023),* above n 240, at [56] – [58].
[244] *Warminger (No 3),* above n 235, at [60].

and 267 of the FMC Act in relation to market manipulation and the relevant disclosure provisions.[245]

In 2025, the FMA brought an action against Kok Ding Cheng based on allegations of market manipulation under s 265 of the FMC Act. The FMA argued that Mr Cheng's placement of five buy orders for shares in Rua constituted a breach of this Section. The legal framework for determining liability under s 265 requires the FMA to demonstrate that Mr Cheng's actions either had or were likely to have created a false or misleading appearance regarding the extent of active trading or the supply, demand, price, or value of the shares and that Mr Cheng knew or ought reasonably to have known that his actions would produce such an effect.[246]

In assessing the likelihood of a false or misleading appearance, the Court clarified that 'likely" signifies an effect that "might well happen" rather than "more likely than not."[247] The evaluation of this likelihood is objective, based on the circumstances prevailing at the time of the orders, without the benefit of hindsight. Notably, the FMA was not obligated to prove that any market participant was deceived. While the intention behind the actions is not a formal element of liability under s 265, the Court recognised its significance as a factor in whether market manipulation occurred. The purpose of a trade can be instrumental in distinguishing legitimate trading activity from manipulative conduct.

The Court also highlighted that the purpose and the requisite knowledge can be inferred from the totality of circumstances, including behaviour patterns, any irregularities observed, and the trading data itself. The standard applied to determine constructive knowledge is inherently objective, focusing on what a person "ought reasonably to know". Ultimately, the Court inferred that Mr Cheng's strategy of placing small, late-day buy orders, often at prices that crossed the spread and were uneconomical due to brokerage costs, coupled with his unconvincing explanations, indicated a purpose to manipulate the closing price and create a misleading impression of demand. The Court also concluded that Mr Cheng, given his experience as an investor, knew or should have known that his actions were likely to have this manipulative effect.[248]

---

[245] *Wei Zhong (2023),* above n 240, at [70] – [78].
[246] *Financial Markets Authority v Cheng* [2025] NZHC 1017 at [1]–[10].
[247] *Kok Ding Cheng*, above n 246, at [37] – [41].
[248] *Kok Ding Cheng,* above n 246, at [60] – [66].

Given Mr Cheng's failure to file a statement of defence, the FMA sought judgment through formal proof. This required the FMA to present affidavit evidence that satisfied the Judge, to the standard of the balance of probabilities, that each element of the alleged contravention of s 265 had been established. The Court was satisfied on the balance of probabilities that the elements of s 265 were established in relation to all five orders, including that the orders had, or were likely to have had, the prohibited effect and that Mr Cheng possessed the requisite knowledge.[249]

These judgments reinforce that the FMA need not prove a specific manipulative intent under s 265. However, liability does not depend solely on the objective market effect. The defendant must also have known, or ought reasonably to have known, that the conduct would or was likely to have the prohibited effect. Purpose may remain relevant evidence when determining whether the conduct was manipulative and assessing the defendant's knowledge.[250] In the context of AI-driven trading, the difficulty is therefore not whether the AI itself possesses intent, but how the statutory knowledge requirement can be attributed to the relevant human or corporate actor.

Taken together, the New Zealand authorities reveal several recurring themes. The prohibited market effect is assessed objectively; actual or constructive knowledge remains required; specific manipulative intent is not an independent element of liability; and the purpose of the trading may nevertheless provide relevant evidence when distinguishing genuine trading from manipulation and assessing knowledge. These principles provide a foundation for technologically mediated trading, but the existing cases do not resolve how knowledge and responsibility should be attributed where increasingly autonomous AI systems generate trading conduct that was not specifically directed or understood by an identifiable human actor.

5.3. INTERNATIONAL CONTEXT

From the market manipulation cases tried in New Zealand, it is evident that the legal precedents established thus far primarily address traditional trading methods involving direct

[249] *Kok Ding Cheng*, above n 247, at [67]–[72].
[250] *Wei Zhong (2023),* above n 240, at [50]–[51]; *Warminger (No 3*), above n 235, at [60].

individual involvement. Notably, New Zealand's legal history has yet to see a specific case in which sophisticated trading algorithms and/or AI technologies played a significant role in market disorder. Given the accelerating pace of technological advancement and the pervasive integration of AI across sectors, including financial markets, closely monitoring international developments in this evolving landscape is essential. Examining the international context is crucial for several reasons. Firstly, jurisdictions with more active trading environments, such as the United States and Europe, might have already begun grappling with the potential for AI-driven market manipulation. Understanding their experiences, including the manipulative tactics algorithms employ, can provide valuable insights for New Zealand. Secondly, as financial markets are increasingly interconnected globally, manipulative activities originating in other jurisdictions from AI systems could impact New Zealand's markets. Therefore, awareness of international precedents is vital for anticipating and mitigating cross-border risks.

The 2015 case of *U.S. Commodity Futures Trading Commission v Nav Sarao Futures Limited PLC and Navinder Singh Sarao* stands as a significant marker in the history of market manipulation, particularly in understanding the application of existing legal frameworks to increasingly sophisticated trading techniques.[251] The action brought by the CFTC against British national Navinder Singh Sarao for his manipulation of the Chicago Mercantile Exchange's E-mini S&P 500 futures contract (E-mini S&P) provided a detailed examination of a modern "spoofing" scheme, one that leveraged algorithmic capabilities to create a false impression of market activity for personal profit. The case underscores the enduring relevance of intent and effect in determining market manipulation, even in the realm of automated trading. The proceeding was resolved by consent order rather than a contested judgment, and Sarao admitted the relevant allegations and findings for the purposes of that order.

At the heart of Sarao's manipulative strategy was the design and utilisation of an automated "layering" algorithm.[252]This program was engineered simultaneously to generate and display exceptionally large sell orders in the E-mini S&P market. These orders, often exceeding a notional value of $140 million, were systematically modified thousands of times and ultimately cancelled in over 99% of cases, never resulting in actual trades. The sheer

[251] *U.S. Commodity Futures Trading Commission v Nav Sarao Futures Ltd & Sarao* No 1:15-cv-03398 (ND Ill, 17 November 2016) (Consent Order for Permanent Injunction, Civil Monetary Penalty and Other Equitable Relief) at [3] – [18]; see also CFTC and SEC, above n 36, at 1–6.
[252] *CFTC v Nav Sarao Futures Ltd PLC,* above n 251, at [10] – [15].

volume of these visible but fleeting orders was intended to create a false perception of significant selling pressure, thereby influencing other market participants' trading decisions. During his period of activity, Sarao allegedly accrued approximately $40 million in profits by exploiting the artificial price movements he induced through his "layering" technique.

The CFTC successfully argued that these actions violated the Commodity Exchange Act ("CEA") and Commission regulations prohibiting manipulation, the use of manipulative devices, and spoofing. The Consent Order, a settlement agreement Sarao entered without a trial, detailed key findings pertinent to determining liability. Sarao admitted to all the allegations presented in the CFTC's Complaint and all the Findings of Fact and Conclusions of Law outlined in the Consent Order. This admission established the factual basis of the manipulative scheme for the specific legal outcome of the settlement.

The findings explicitly stated that Sarao's purpose and intent in utilising his various spoofing techniques were to place orders he never intended to execute, thereby creating a deceptive sense of supply and/or demand. This manipulation was designed to influence other market participants' actions and generate artificial prices, ultimately allowing Sarao to profit from his genuine trades. This intent was evidenced by his communications with software developers, requesting specific programming codes to facilitate order cancellation based on market proximity, rapid order placement at varying prices, and the ability to adjust order aggressiveness. The findings also detailed how Sarao's actions frequently caused artificial price movements in the E-mini S&P futures market.[253]

Furthermore, the CFTC demonstrated a clear temporal proximity and a direct link to profitability between Sarao's spoof order activity and his subsequent genuinely executed trades, suggesting that the former was intentionally designed to benefit the latter. The sophistication of Sarao's techniques, which combined automated programs specifically engineered for layering and queue manipulation with manual methods, indicated a deliberate, intricate effort to manipulate the market.

---

[253] *CFTC v Nav Sarao Futures Ltd PLC*, above n 251, at [22]–[28].

While the Consent Order represents a settlement rather than a litigated judgment with a court-articulated "test", the findings admitted by Sarao strongly suggest that liability for spoofing under the relevant provisions of the CEA and CFTC Regulations hinges on several elements. These include the act of placing orders with the specific intent not to execute them, the creation of a false or misleading appearance of supply or demand through these orders (regardless of execution), and the manipulative intent to induce others to trade based on this false impression, ultimately benefiting the spoofer's genuine trades.

The *Sarao* case demonstrated that the defendant's state of mind (intent) was a core legal requirement of the claims. Because Sarao consented to the findings of fact, the case determined that he acted with the intent to manipulate and defraud the market. This intent was critical for proving the specific violations: first, he was found to have the intent to cancel orders before execution, which is the statutory definition of spoofing. Second, he had the broader intent to fabricate a false market impression of supply and demand to induce other traders to act, which is the core of market manipulation. Therefore, although not contested in court, his admitted intent served as the foundational basis for the CFTC's successful enforcement action.

As the *Sarao* case was resolved through a consent order, the Court did not formally determine the standard of proof following a contested trial. However, in civil enforcement actions initiated by the CFTC, the standard of proof is typically the preponderance of the evidence, meaning the defendant must be found more likely than not to have engaged in the alleged conduct. In this instance, the CFTC would have needed to present evidence showing that it was more likely than not that Sarao engaged in the suspected manipulative conduct with the requisite intent and that his actions had or were likely to have a prohibited effect on the market.

In summary, the *Sarao* proceeding demonstrates the application of existing spoofing and market manipulation rules to the deliberate human use of automated trading technology. The proceeding was resolved by consent order rather than a contested judgment, and Sarao admitted the relevant allegations and findings for the purposes of that order.[254] Those admissions concerned the placement of orders with the intent to cancel them before execution

[254] *U.S. Commodity Futures Trading Commission v Nav Sarao Futures Ltd & Sarao* No 1:15-cv-03398 (ND Ill, 17 November 2016) (Consent Order for Permanent Injunction, Civil Monetary Penalty and Other Equitable Relief) at [3]–[18].

and the use of automated trading techniques to create a false or misleading appearance of market activity. *Sarao* does not concern an AI system independently developing a manipulative strategy and should not be treated as authority for liability arising from autonomous AI conduct.

The 2018 civil complaint filed by the CFTC against Jitesh Thakkar, a software engineer, and his company, Edge Financial Technologies, Inc. ("Edge"), provides a critical perspective on the scope of liability in market manipulation cases. Unlike the action against the primary manipulator, Navinder Singh Sarao, this case focused on the responsibility of those who allegedly facilitated the manipulative scheme. The CFTC sought to compel Thakkar and Edge's compliance with the CEA and pursue civil monetary penalties and other equitable relief, arguing that they aided and abetted Sarao's "spoofing" activities in the E-mini S&P 500 futures market. This case is crucial as it delves into the liability tests for those who contribute to market manipulation without directly executing the manipulative trades themselves. It implicitly addresses the standard of proof required in such civil enforcement actions.[255]

The CFTC's complaint centred on the custom trading software application developed by Thakkar and Edge for Sarao. A key feature, added specifically at Sarao's request, was the "Back-of-Book Function." This functionality automatically adjusted Sarao's orders to keep them behind other market participants' orders at the same price level, significantly reducing the likelihood that these large orders would be executed. Additionally, the software included an automatic cancellation function that would cancel Sarao's orders as soon as any portion was executed. This combination of features allowed Sarao to maintain the appearance of substantial trading interest without the genuine intention of completing the trades, a hallmark of spoofing. The complaint further highlighted Thakkar's opportunity to study Sarao's trading patterns in the E-Mini S&P and discuss improving the Back-of-Book function, suggesting a deeper understanding and involvement in the manipulative strategy.

The CFTC's argument for aiding and abetting liability rested on demonstrating that Thakkar and Edge knowingly provided the means for Sarao to commit market manipulation. The complaint noted that despite never being registered with the CFTC, Thakkar and Edge had

---

[255] *U.S. Commodity Futures Trading Commission v Thakkar* No 1:18-cv-00619 (ND Ill, 26 January 2018) (Complaint) at [27] – [34].

extensive experience developing algorithmic trading systems. Crucially, the CFTC highlighted that during the design of the Back-of-Book function, Thakkar and Edge recognised its potential for spoofing and for injecting false market supply and demand information into the market. This knowledge, coupled with the specific design features tailored to Sarao's requests and their continued troubleshooting and improvement of the software in communication with Sarao, formed the basis of the aiding and abetting claim.

The liability test for aiding and abetting, as outlined in the CFTC's complaint, hinges on satisfying three key elements. First, there was a primary violation by Trader A in the form of spoofing, which he admitted to committing. Second, knowledge of the primary violation by Thakkar and Edge. The CFTC argued that their understanding of Trader A's intentions was evident in the specific design features of the Back-of-Book function, their communications with Trader A, and Thakkar's expertise in HFT strategies and spoofing through his involvement with the CFTC's Technology Advisory Committee. The CFTC contended that the Back-of-Book function directly enabled Trader A's spoofing by minimising the risk of order execution and automating cancellations, thus facilitating the injection of false market signals. Even after Trader A encountered issues, the continued support and development of the software by Thakkar and Edge further underscored their ongoing assistance. The CFTC alleged that Thakkar and Edge knowingly assisted Sarao by developing functionality used in his spoofing strategy. These were allegations rather than findings made against Thakkar following a contested trial. The civil claims against Thakkar were subsequently dismissed with prejudice, while Edge separately entered into a consent order.[256]

Regarding the standard of proof, a civil enforcement action was required, and the CFTC was required to demonstrate its claims by a preponderance of the evidence. This means they needed to show that it was more likely than not that Thakkar and Edge knowingly provided substantial assistance to Trader A's spoofing activities. The evidence presented included the design and functionality of the Back-of-Book function, communications between Thakkar and Trader A discussing its intended use, and the continued support provided by Thakkar and Edge for the software. The CFTC sought various forms of relief, including permanent injunctions against future violations and against the sale of the manipulative software, trading restrictions, civil monetary penalties, disgorgement, and payment of costs and fees. The numerous instances

[256] *CFTC v Thakkar*, above n 255, at [11], [23]–[36] and [42]–[44].

of Trader A's spoofing using the Back-of-Book function between 30 January 2013 and 30 October 2013, were alleged to constitute separate and distinct violations for which Thakkar and Edge were responsible as principals under the CEA.

As AI algorithms become increasingly sophisticated and capable of executing complex trading strategies at high speeds, this precedent for holding enablers liable for manipulative conduct becomes crucial. The case establishes that even those who do not directly engage in manipulative trading but provide the technological means with the knowledge and intent to facilitate it can be held responsible. This principle could be directly applied to developers and providers of AI trading platforms, or to specific AI algorithms designed or foreseeably used for manipulative purposes, such as spoofing or creating artificial market signals.

Furthermore, the liability test, which emphasises knowledge and substantial assistance, is particularly relevant in the context of AI. While attributing "intent" to an AI remains a complex legal challenge, focusing on the designer's intent and the foreseeable manipulative capabilities embedded within the AI's architecture could form the basis for liability. The "Back-of-Book" function in the Thakkar case serves as an analogy for sophisticated AI algorithms that could be designed to subtly manipulate order books or exploit market microstructure in ways that constitute spoofing or other forms of manipulation. The standard of proof, the preponderance of the evidence, will likely remain the benchmark in civil cases, requiring the CFTC to demonstrate that it is more likely than not that the developers of AI trading tools had the requisite knowledge of their potential for manipulation and provided substantial assistance in facilitating such activities. The proceedings therefore illustrate the evidential difficulty of establishing responsibility against a technology developer who did not personally execute the manipulative trades, rather than establishing a general rule of developer liability.

The Financial Conduct Authority ("FCA") initiated a claim against investment firms Da Vinci Invest Limited ("DVI"), Mineworld Limited ("Mineworld"), and three individual traders for market abuse under Sections 381 and 129 of the Financial Services and Markets Act 2000 ("FSMA"). This claim stemmed from alleged market manipulation in 2010 and 2011 involving high-volume trading of shares on the London Stock Exchange ("LSE"). The core allegation was that the traders used algorithmic trading to repeatedly and automatically place

immediate orders on the LSE, then cancel the initial orders to capitalise on the price movement.[257]

The High Court focused on whether market abuse occurred and whether the individual traders' actions could be attributed to the corporate defendants. The Court emphasised that under s118(5) of FSMA, the key consideration is "the likely perception of other parties" rather than the traders' subjective intentions. Thus, the traders' state of mind was deemed immaterial in determining if market abuse occurred. Regarding attribution, the court concluded that FSMA allows a company to be held accountable for its agents' market abuse, although the company can present evidence to mitigate or avoid liability.[258]

To clarify the liability test, the court assessed whether the defendants' actions gave or were likely to give a false or misleading impression of the supply, demand, or price of shares. This is an objective test, focusing on the behaviour's impact on the market rather than its intent. For corporate defendants, this involved an additional layer of attributing traders' actions, requiring careful analysis of the traders' authority and the company's oversight. Liability is attributed under agency principles, holding them responsible for the actions of those authorised to trade on their behalf, regardless of the precise nature of their relationship (e.g., employer and contractor). The crucial factor is whether the behaviour occurred on behalf of the company. [259]

The High Court confirmed that the civil standard of proof, on the balance of probabilities, applies in market abuse cases, meaning that the FCA had to prove that it was more likely than not that market abuse occurred. The Court acknowledged the need for strong evidence due to the inherent improbability of the allegations. The FCA's expert witness, who was involved in designing the 'SMARTS' market surveillance system, provided the crucial evidence. The expert's analysis of the trading patterns revealed that the traders' orders created a false or misleading impression of increased supply or demand, as their true purpose was not to trade but to manipulate prices for profit. The Court concluded that the transactions contravened s118(5) of the FSMA, enabling the defendants to trade from an advantageous position. [260]

---

[257] *Financial Conduct Authority v Da Vinci Invest Ltd* [2015] EWHC 2401 (Ch) at [40] – [52].
[258] *Da Vinci Invest,* above n 257, at [85] – [104].
[259] *Da Vinci Invest*, above n 257, at [105] – [114].
[260] *Da Vinci Invest,* above n 257 at [115] – [125].

Financial Conduct Authority v Da Vinci Invest Ltd is particularly relevant because the High Court focused on whether the trading behaviour gave, or was likely to give, a false or misleading impression to other market participants. The statutory inquiry was therefore directed to the objective character and market effect of the conduct rather than requiring proof of a subjective intention to manipulate.[261] The Court also considered the attribution of the traders' conduct to the corporate defendants.

The decision provides a useful comparison with the New Zealand approach because it demonstrates how an effects-based prohibition can operate where trading is executed through automated systems. However, *Da Vinci Invest* does not establish strict liability for algorithmic trading or determine responsibility for autonomous AI behaviour. Its relevance is that objective market effects can be assessed independently of the technology used to execute the trades, while responsibility must still be attributed to the relevant legal person.

In comparison, *Panther Energy Trading LLC and Michael J. Coscia* feature an explicit example of an administrative proceeding under Sections 6(c) and (d) of the CEA targeting AI-driven market misconduct that requires subjective intent, thereby reinforcing the two-track nature of anti-manipulation enforcement. Since spoofing is statutorily defined by the intent to cancel, the CFTC had to prove Michael Coscia's intent, which it achieved by demonstrating that the algorithm itself was deliberately designed to spoof. The key takeaway is that the relevant intent was therefore fixed at the time the algorithm was designed and deployed, rather than at the moment each order was placed or cancelled.[262] This stands in contrast to *Da Vinci Invest*, where liability is determined by the conduct's objective market effect rather than by proof of a specific intent to manipulate.

In a related matter, the United Kingdom's Financial Conduct Authority (FCA) issued a Final Notice regarding its enforcement action against Mr Coscia relating to his market abuse activities on the ICE Futures Europe (ICE), a Recognised Investment Exchange. The FCA concluded that Mr Coscia's behaviours amounted to market abuse (market manipulation) for breach of Section 118(5) of the Financial Services and Markets Act 2000. It was noted that Mr Coscia was not a member of ICE or an approved person and traded from the United States

[261] *Da Vinci Invest,* above n 257 at [107], [156] and [194] – [201].
[262] *In re Panther Energy Trading LLC and Michael J Coscia* CFTC Docket No 13-26 (22 July 2013) at [4] – [9].

through a Direct Market Access (DMA) provider. This Notice was also limited to Mr Coscia's trading activity on his proprietary account, where an automated algorithmic program was deployed to place orders for Brent Crude, Gas Oil, and Western Texas Intermediate futures from the US on the ICE in the UK at much faster rates than manual order entry, like HFT. The large orders were then automatically cancelled following a partial execution. As a starting point, the FCA analysed the effect of trading activity on the market and whether trading activity that utilised cancellations created a false or misleading impression.[263]

To investigate Mr Coscia's trading activity, the ICE used a monitoring tool to identify the layer trading pattern. The ICE found that Mr Coscia's trading pattern breached the ICE rule on Disorderly Trading. His trading activities created a false impression of liquidity rather than genuine market supply and demand. Furthermore, the buy and sell orders were found to lack a legitimate justification and to be contrary to accepted market practices. Consequently, he profited substantially from his trading to the detriment of other market participants. Compared to the approach adopted in US civil actions, the FCA did not directly emphasise the trader's wrongful trading intent to create a false or misleading effect. Instead, they stressed that the trader used specific cancellation techniques.

The CFTC initiated an enforcement action against Panther and Coscia for engaging in 'spoofing', a prohibited practice under Section 4c(a)(5)(C) of the CEA. The CFTC's findings, to which Panther and Coscia consented without admitting or denying, covered activities from 8 August 2011 to 18 October 2011. The violation's core involved algorithmic trading, a technological precursor to modern AI-driven trading systems. Panther and Coscia employed a computer program designed to place bids rapidly and offers on the CME Group's electronic trading platform (Globex) with the specific intent to cancel those orders before they could be executed.[264]

They would place a smaller order on one side of the market, positioned near the best available price, while simultaneously entering larger orders on the opposite side at progressively more favourable prices. This tactic aimed to create a false impression of significant buying or selling interest, thereby artificially influencing the market to increase the

---

[263] Financial Conduct Authority *Final Notice: Michael Coscia* (3 July 2013) <https://www.fca.org.uk/publication/final-notices/coscia.pdf>.
[264] *Panther Energy and Coscia*, above n 262, at [10] – [18].

likelihood that their smaller order on the other side would be filled. Once this initial small order was successfully executed, the larger, misleading orders were promptly cancelled, and the algorithm would reverse its operation. This pattern of activity was repeated across a range of commodity futures contracts, generating approximately $1.4 million in net profits for the respondents.

The CFTC's finding of a violation of Section 4c(a)(5)(C) of the Act directly addresses liability for spoofing. The Act defines spoofing as "bidding or offering with the intent to cancel the bid or offer before execution." Liability was established by demonstrating that the respondents placed bids or offers on a registered entity's trading platform (Globex) and that these orders were placed with the intent, facilitated by their algorithmic program, to be cancelled before execution. The conduct was explicitly identified as the disruptive practice of spoofing. While this case predates the common understanding of sophisticated AI in trading, the principle of liability is highly relevant.

In the context of AI-driven trading, the "intent to cancel before execution" is effectively embedded in the AI's design and parameters. This is true even if a human does not step in at the time of the trade. This application of liability is not a legal anomaly but rather an extension of established tort principles. For example, in *Katko v Briney 183 N.W.2d 657* (Iowa 1971), the defendants were held liable for setting up a "spring gun" that shot a trespasser. Even though they did not pull the trigger themselves, they were liable because they set the machine to do so.

The AI is just like that spring gun, that is, the tool used to carry out the wrongdoing. While the AI is infinitely more complex than a mechanical gun, the focus remains on the demonstrable functionality and the intended outcome encoded into the system. If an algorithm is designed to generate orders primarily to influence prices with the rapid-fire capability to cancel them, it represents a digital "spring gun" calibrated for spoofing. The focus shifts to what the code was made to do and the result it was meant to achieve.

The Court's decision included an Order Instituting Proceedings based on a settlement offer. In such settlements, the respondents typically consent to the findings without admitting or denying them. Consequently, the Order does not provide an exhaustive account of the evidentiary burden the CFTC would have had to meet in a contested proceeding. However, it

is essential to note that in CFTC enforcement actions, the standard of proof for establishing violations of the Commodity Exchange Act is generally the preponderance of the evidence. This requires the CFTC to demonstrate that it is more likely than not that the respondent engaged in the alleged conduct (i.e., spoofing).

In the AI trading context, meeting this standard would likely involve presenting evidence such as the AI's code, trading logs showing the pattern of order placement and cancellation, expert testimony on market impact, and potentially any data relating to the AI's design and intended function. Although the *Panther* and *Coscia* case do not explicitly detail the evidence used beyond the description of the algorithm's operation, it sets a legal precedent that utilising algorithmic programs to create a false impression of market interest through rapid order cancellation constitutes spoofing under the Act. This precedent is crucial for future enforcement actions against manipulative AI trading strategies.

The Australian Securities and Investments Commission (ASIC) brought a claim in July 2024 against COFCO International Australia Pty Ltd ("COFCO Australia") and COFCO Resources SA ("COFCO Resources") in the Federal Court concerning their trading activities in Eastern Australia Wheat Futures Contracts for the delivery of wheat in January 2023 ("WMF3 contracts") on the ASX 24 market. ASIC alleged that on 34 occasions between January and March 2022, shortly before the close of the trading day, the defendants, through their employee Nicholas McGaw, made offers to sell WMF3 contracts with the undisclosed intention of influencing the Daily Settlement Price ("DSP") calculated by the Australian Securities Exchange ("ASX"). ASIC's media release noted that this conduct, placing orders shortly before the close of the day session of the ASX24 to create artificial prices for WMF3 contracts, is commonly referred to as 'marking the close' and that the resulting price did not reflect the forces of genuine supply and demand. The DSP is a crucial daily quotation used for margining and settling these derivative contracts, and its integrity is vital for public confidence in the commodity derivatives market.[265]

The ASIC contended that these offers to sell did not constitute genuine participation in the supply-and-demand process. Instead, Mr McGaw's purpose, as indicated in an email, was

[265] Australian Securities and Investments Commission "ASIC Sues COFCO International Australia Pty Ltd and COFCO Resources SA for Alleged Market Manipulation" (Media Release 24-163MR, 11 July 2024).

to protect COFCO International's short position in WMF3 contracts from other parties potentially bidding up the DSP. When these offers were made, it was likely that trades would occur, and indeed, most of the offers (22 out of 34) resulted in trades and had the intended effect of lowering the DSP.

The ASIC's primary legal grounds for relief were the civil penalty provisions under Sections 1041A and 1041B(1)(b) of the Corporations Act 2001. Section 1041A prohibits transactions that have or are likely to have the effect of creating or maintaining an artificial price for trading in financial products on a financial market. Section 1041B(1)(b) prohibits acts or omissions that have or are likely to have the effect of creating a false or misleading appearance concerning the market for, or the price for trading in, financial products on a financial market. The ASIC emphasised that the WMF3 contracts were financial products traded on the ASX24 market and thus fell within the purview of these Sections.[266]

The ASIC asserted that Mr McGaw's conduct, attributed to both COFCO entities due to his authority, constituted market manipulation as it undermined the integrity of the ASX 24 market and potentially caused other participants to suffer financial detriment due to increased margin requirements on their WMF3 contract trades. ASIC sought declarations of contravention and pecuniary penalty orders against the defendants. To establish liability against COFCO, ASIC's primary task is to prove, on the balance of probabilities, that the defendants' trading activities were likely to result in either an artificial price for WMF3 contracts or a false and misleading impression of the market's genuine supply and demand. Importantly, the ASIC does not need to demonstrate that COFCO acted with a specific malicious intent to manipulate the market to secure civil penalties under the relevant Sections of the Corporations Act.

The COFCO case, though focused on human-driven manipulation, provides a vital framework for AI-driven trading in Australian markets by shifting the legal focus from subjective malicious intent to the objective effect of trading actions on market integrity. In the context of AI, debating whether an algorithm can possess “intent” is a legal distraction. Because AI lacks legal personhood, cannot be a defendant, and possesses no assets, attributing intent to it is not just tricky, it is legally impossible. However, this impossibility is irrelevant to the core legal question. Just as a shooter cannot avoid liability by blaming the firearm, a

[266] ASIC, above n 265.

market participant cannot shift responsibility to an algorithm. AI is a sophisticated tool, and the legal standards remain anchored to the human actors, the designers who build it and the users who deploy it, who are responsible for their own behaviour and the tools they choose to employ.

The argument for focusing on the objective consequences of AI behaviour acknowledges that the primary challenge is evidentiary, not theoretical. AI often operates as a "black box," making it difficult to establish the evidence needed to prove what a human controller "intended" for a machine to do at a specific millisecond during high-frequency trading. By emphasising the market result, the COFCO approach allows the law to apply familiar standards to the human controller based on the outcomes they set in motion. This treats the AI as part of the chain of causation, where the human's decision to deploy or design the system is the primary act of responsibility.

However, the complexity of AI-executed strategies raises a distinct issue of fault. A defendant may attempt to invoke the doctrine of *novus actus interveniens* by arguing that an AI's autonomous learning or "emergent behaviour" constituted an intervening cause that broke the chain of responsibility between the original design and the resulting market distortion. That argument misconceives the nature of fault in this context. Where a person deploys a tool, whose behaviour is highly unpredictable, that very unpredictability makes harmful outcomes more readily foreseeable, not less so. The greater the autonomy and opacity of the system, the more foreseeable it is that its operation may generate misleading market effects, and the greater the responsibility on the human operator to anticipate and guard against that risk.

Novus actus interveniens is concerned with whether responsibility for the harm should be displaced because it is truly the fault of another person. It does not apply where the alleged intervening cause is merely the operation of a tool chosen and deployed by the defendant. An AI system is not a person capable of independent fault, agency, or legally meaningful decision-making. To invoke *novus actus* in this context is therefore to assert, in substance, "I did not do; the AI did it," an argument that fails because the AI cannot bear responsibility. The relevant fault remains with the human actor who designed, deployed or continued to operate a system whose foreseeable operation distorted the market.

5.4. LESSONS

The core challenge AI poses for market regulation is not a change in legal theory but a significant evidentiary hurdle. The core challenge AI poses for market regulation is partly evidential. As AI-driven trading systems become increasingly complex and opaque, establishing the knowledge, purpose or intent of the relevant human or corporate actor may become more difficult. The authorities considered in this chapter demonstrate that objective market effects remain important, but they do not establish that those effects alone are sufficient for liability. The applicable statutory requirements concerning knowledge, purpose or intent must still be established against the relevant legal person.

A comparative examination of the international authorities also demonstrates that the relevant mental element differs between statutory regimes. In the United States, the specific prohibition on spoofing requires proof of an intent to cancel an order before execution. By contrast, the United Kingdom authorities considered in this chapter place greater emphasis on the actual or likely market effect of the conduct. New Zealand adopts an effects-based approach under s 265 of the FMC Act, while retaining the requirement that the defendant knew or ought reasonably to have known that the conduct would or was likely to have the prohibited effect.

New Zealand's constructive knowledge requirement may therefore be relevant to AI-driven trading because liability under s 265 does not depend upon proof of a specific manipulative intent. However, the existing authorities generally concern conduct attributable to identifiable human or corporate actors and provide limited guidance where an increasingly autonomous AI system generates trading behaviour that was not specifically directed or understood by a human actor. This creates potential evidential and attribution difficulties where the market

5.5. CONCLUSION

In conclusion, while there are currently no legal precedents that directly address the intersection of artificial intelligence and market manipulation, the existing statutory framework provides a robust foundation for guidance. In the New Zealand context, the FMC Act remains the primary regulatory vehicle. The authorities considered in this chapter generally attribute market manipulation to identifiable human or corporate actors. They do not, however,

determine how responsibility should be attributed where increasingly autonomous AI systems generate trading conduct that was not specifically directed or understood by a human actor.

The application of New Zealand case law, including *FMA v Brian Henry*, *FMA v Warminger* and *FMA v Wei Zhong & Ors*, demonstrates how the courts apply the statutory requirements concerning market effect and actual or constructive knowledge. Specific manipulative intent is not an independent element of liability under s 265 of the FMC Act, although purpose may remain relevant evidence. The unresolved issue for increasingly autonomous AI-driven trading is how the required knowledge and responsibility should be attributed where the relationship between the relevant human actor and the trading conduct becomes less direct.

Internationally, the authorities similarly demonstrate the application of existing market manipulation rules to technologically mediated trading, although their procedural status differs. *Da Vinci Invest* is a judicial decision, while the *Sarao* and *Panther Energy* matters involved regulatory enforcement outcomes and the COFCO proceeding remains unresolved. These authorities therefore provide comparative guidance on automated and algorithmic trading but do not establish a uniform approach to autonomous AI-driven market manipulation.

Overall, the authorities demonstrate that existing market manipulation principles can apply to technologically mediated trading where the relevant statutory requirements can be established against an identifiable human or corporate actor. However, the authorities considered in this chapter do not resolve how those requirements should apply where increasingly autonomous and opaque AI systems make the relevant human knowledge, purpose or responsibility more difficult to establish. This unresolved issue informs the analysis in Chapter 6 and the assessment in Chapter 7 of whether additional governance-based regulatory responses are justified.

Chapter 6: Negligence and Product Liability

6.1. INTRODUCTION

Chapters 4 and 5 examined the public-law response to AI-driven market manipulation, focusing on the statutory prohibitions in the Financial Markets Conduct Act 2013 and their judicial application. While public enforcement plays a critical role in maintaining market integrity and is principally directed at deterrence and sanction, the FMC Act is not exclusively punitive. Part 8 permits declarations of contravention under ss 486 - 487 and compensatory orders under ss 494 - 495 for loss or damage resulting from a contravention of a civil liability provision. These statutory routes may nevertheless be difficult to use in AI-related cases where opacity, multiple actors and market movements make contravention, causation and loss difficult to prove. This chapter therefore turns to private law to consider whether negligence and product liability can deliver meaningful accountability and redress to investors and other market participants harmed by AI-driven trading systems. The central question is not whether victims can access remedies in the abstract, but whether the existing fault-based and strict-liability frameworks are doctrinally and practically suited to govern harms that are predominantly financial, often diffuse, and arise from technologically opaque and operationally distributed systems.

Section 6.2 examines negligence, focusing on the conditions under which developers, operators, or firms deploying AI may be held liable for failing to exercise reasonable care in the design, oversight, or use of AI trading systems. Section 6.3 turns to product liability under the Consumer Guarantees Act 1993 and the Fair Trading Act 1986, assessing the extent to which these statutes impose strict liability for defective or misleading AI products. Section 6.4 draws the two frameworks together and identifies the contexts in which each is most appropriate, while Section 6.5 evaluates their combined adequacy. The analysis exposes significant doctrinal limits, particularly the rights-based constraints on recovery for pure economic loss, the evidentiary difficulties created by AI opacity, and the gaps left by the consumer-focused architecture of the CGA and FTA. These limits are not analytical detours; they directly justify the bespoke regulatory, tortious, and product-liability reforms developed in Chapter 7, including the proposed AI Governance Liability tort and a targeted

product-liability framework for AI financial products. A detailed examination of those reforms follows in the next chapter.

The shift to private law requires an important qualification. Part 8 of the Financial Markets Conduct Act 2013 already provides a statutory route to compensation where loss results from a contravention of a civil liability provision. A declaration of contravention may provide conclusive proof of the contravention for a later compensatory claim, and a compensatory order may also be sought without a prior declaration where the Court is satisfied that the statutory requirements are met. The question addressed in this chapter is therefore not whether the FMC Act contains any compensatory mechanism. Rather, it is whether that mechanism adequately responds where harm arises from defective system design, negligent development or supervision, or an autonomous system failure that cannot be proved to constitute market manipulation or another civil liability contravention. Negligence and product liability are examined as possible responses to that residual accountability gap.[267]

6.2. NEGLIGENCE

The application of negligence law to AI-driven market manipulation is complex and faces fundamental conceptual barriers under current principles. The primary challenge is that an AI system lacks legal personality and cannot be sued directly; therefore, liability must be directed towards the natural or legal persons responsible for its development, deployment or operation.[268] A broad range of parties could be considered, including the AI developer, the programmer, the operator, and the firm that employs the AI for trading. The tort of negligence provides a civil remedy for harm caused by a failure to exercise a reasonable standard of care. To establish a successful negligence claim, a plaintiff must prove four core elements: a duty of care, a breach of that duty, causation, and resulting damage.

This section does not argue that negligence supplies a general, market-wide remedy for losses attributable to algorithmic trading or price distortion. Its purpose is instead to (i) explain

[267] *Financial Markets Conduct Act 2013*, ss 486–488 and 494–495.
[268] Pınar Çağlayan Aksoy "AI as Agents" in Larry A DiMatteo, Cristina Poncibò and Michael Cannarsa (eds) *The Cambridge Handbook of Artificial Intelligence: Global Perspectives on Law and Ethics* (Cambridge University Press, Cambridge, 2022) 153 at 159–160.

why negligence is constrained in this setting, primarily because the loss is typically pure economic loss and because the potential claimant class may be indeterminate and (ii) identify the limited (and truly exceptional) circumstances in which negligence remains applicable on orthodox principles. Properly framed, negligence remains relevant because AI is typically deployed as a tool within human decision-making systems: liability focuses on whether human actors exercised reasonable care in the design, testing, deployment, monitoring, and maintenance of those systems, and on whether there were identifiable points of human control at which reasonable precautions could and should have been taken.[269]

In the specific context of AI-driven market manipulation, the negligence inquiry cannot be approached in isolation from the nature of the loss suffered. The harm alleged will ordinarily be financial and will typically constitute pure economic loss, rather than loss consequential upon physical injury or damage to property. Although New Zealand law does not impose an absolute bar on recovery for economic loss, its presence remains a significant factor weighing against recognition of a duty of care. Foreseeability alone is insufficient. The existence and scope of any duty must be justified consistently with the controls New Zealand courts apply to pure economic loss claims, particularly concerns relating to indeterminacy, proximity, reliance, and coherence with other areas of private law.[270]

6.2.1. A Rights-Based Account of Negligence

Before turning to the elements of negligence, it is necessary to identify the theoretical lens through which they will be examined. This thesis adopts a rights-based account of negligence of the kind developed most fully by Beever and supported by Stevens.[271] On this view, the law of negligence does not protect persons against loss simpliciter. Rather, it protects pre-existing primary rights and recognises a duty of care only where the defendant's conduct is capable of infringing such a right.[272] Reasoning in recent Canadian and Australian economic loss cases is consistent with aspects of this relational approach, although those courts have not adopted Beever's rights-based theory as such. In *Deloitte & Touche v Livent Inc (Receiver of)*,

---

[269] Bryan H Choi 'Negligence Liability for AI Developers' (26 September 2024) *Lawfare* <https://www.lawfaremedia.org/article/negligence-liability-for-ai-developers>.

[270] Allan Beever "A Rights-Based Approach to the Recovery of Economic Loss in Negligence" (2004) 4(1) Oxford University Commonwealth Law Journal 25 at 26–27.

[271] Allan Beever *Rediscovering the Law of Negligence* (Hart Publishing, Oxford, 2007) at 31–66; Robert Stevens *Torts and Rights* (Oxford University Press, Oxford, 2007) at 20–35 and 174–192.

[272] Beever, above n 271, at 211–250.

1688782 Ontario Inc v Maple Leaf Foods Inc in Canada,[273]and *Mallonland Pty Ltd v Advanta Seeds Pty Ltd*, the courts emphasised the need to confine liability by reference to matters such as the purpose and scope of the defendant's responsibility, reliance, vulnerability and the nature of the interest protected.[274] It provides a principled explanation for why pure economic loss in an AI-driven trading context generally falls outside the scope of negligence in the absence of an assumption of responsibility.

The rights-based reading is not uncontested. Stapleton, Witting, and others have argued for a more openly policy-driven approach to the existence and scope of a duty of care, particularly where novel technological harms are concerned.[275] Bagshaw and Goudkamp have likewise mapped a range of more functionalist accounts of duty in modern negligence.[276]These rival accounts are engaged with at the relevant points below. However, the analysis ultimately aligns with Beever's position that the rights-based account offers the most doctrinally coherent explanation for why negligence is structurally ill-suited to govern the diffuse economic harms arising from AI-driven trading, and why a dedicated cause of action of the kind proposed in Chapter 7 is therefore required.

The distinction between these approaches is important in the present context. A rights-based analysis explains the difficulty of identifying a pre-existing private right held by every market participant to protection against careless interference with market prices. A more policy-oriented analysis reaches a similar practical difficulty by focusing on the consequences of recognising such a duty, particularly the prospect of liability to an indeterminate and constantly changing class of investors. The two approaches therefore differ in their theoretical explanation, but both expose the difficulty of extending negligence to diffuse economic losses arising from AI-driven trading. The rights-based account is useful in explaining that difficulty, but it does not replace the proximity and policy analysis required by New Zealand law.

---

[273] *Deloitte & Touche v Livent Inc (Receiver of)* 2017 SCC 63, [2017] 2 SCR 855 at [21]–[40] per McLachlin CJ and Karakatsanis J; *1688782 Ontario Inc v Maple Leaf Foods Inc* 2020 SCC 35, [2020] 3 SCR 504 at [21]–[33] per Brown and Martin JJ.

[274] *Mallonland Pty Ltd v Advanta Seeds Pty Ltd* [2024] HCA 25 at [22]–[26] per Gageler CJ, Gordon, Steward and Gleeson JJ, and at [115]–[121] per Edelman J.

[275] Jane Stapleton "Duty of Care Factors: A Selection from the Judicial Menus" in P Cane and J Stapleton (eds) *The Law of Obligations: Essays in Celebration of John Fleming* (Oxford University Press, Oxford, 1998) 59 at 70–82; Christian Witting "Duty of Care: An Analytical Approach" (2005) 25 Oxford Journal of Legal Studies 33 at 42–55.

[276] Roderick Bagshaw "The Edges of Tort Law's Rights" in Donal Nolan and Andrew Robertson (eds) *Rights and Private Law* (Hart Publishing, Oxford, 2012) 405 at 410–422; James Goudkamp *Tort Law Defences* (Hart Publishing, Oxford, 2013) at 1–18.

6.2.2. Duty of Care

The first and most fundamental element is a duty of care. This is a legal obligation to avoid acts or omissions that could reasonably be foreseen to cause harm to a 'neighbours.' This "neighbourhood principle," as established in the landmark case of *Donoghue v Stevenson,*[277] is based on the idea of proximity: a duty of care is owed to persons who are so closely and directly affected by a defendant's act that they should have been in the reasonable contemplation of the defendant. This principle was later reframed in *Anns v Merton London Borough Council*, which established a two-stage test: first, is there a sufficient relationship of proximity between the parties such that the defendant's carelessness might cause damage? Secondly, are there any policy considerations which might negate or limit the scope of that duty?[278]

While the UK later adopted a different three-stage test in *Caparo Industries plc v Dickman*,[279] New Zealand's position, as confirmed in *South Pacific Manufacturing Co Ltd v New Zealand Security Consultants & Investigations Ltd* and *Attorney-General v Carter*, is that the Anns two-stage test remains relevant and is not at odds with the *Caparo* test.[280] Although the terminology used in the authorities has varied, New Zealand courts do not treat *Anns* and *Caparo* as competing rigid formulae. The modern inquiry considers reasonable foreseeability, proximity and relevant policy considerations. In *North Shore City Council v Attorney-General (The Grange),* the Supreme Court explained that proximity principally concerns factors bearing on the relationship between the plaintiff and defendant, while the policy stage considers wider consequences of recognising the proposed duty. The categories may overlap, and the inquiry remains sensitive to the particular relationship and type of loss alleged.[281]

In the AI trading context, it is necessary to move beyond a general assertion that harm is foreseeable. It is reasonably foreseeable that a poorly designed, developed, or managed AI trading system could cause significant financial harm. However, where the loss alleged is pure economic loss, foreseeability cannot by itself justify the imposition of a duty of care. The duty analysis must confront the risk that recognising a duty owed to "all other market participants"

---

[277] *Donoghue v Stevenson* [1932] AC 562 (HL) at 580 per Lord Atkin.
[278] *Anns v Merton London Borough Council* [1978] AC 728 (HL) at 751–752 per Lord Wilberforce.
[279] *Attorney-General v Carter* [2003] 2 NZLR 160 (CA) at [18]–[22]; *Caparo Industries plc v Dickman* [1990] 2 AC 605 (HL) at 617–618 per Lord Bridge.
[280] *South Pacific Manufacturing Co Ltd v New Zealand Security Consultants & Investigations Ltd* [1992] 2 NZLR 282 (CA) at 295–301 per Cooke P.
[281] *North Shore City Council v Attorney-General* [2012] NZSC 49, [2012] 3 NZLR 341 at [157] – [160].

collapses negligence into an unbounded mechanism for reallocating trading losses, exposing defendants to indeterminate liability. These are precisely the concerns that economic loss controls are designed to manage.

A principled way to reconcile negligence with these constraints is to refine what "foreseeability" is doing. Negligence does not require that a defendant anticipate the precise manner in which harm will occur or predict each specific AI output in advance. Rather, the relevant inquiry is whether the defendant's conduct created a foreseeable class of risk and whether reasonable care could have reduced that risk at identifiable points of human control. In the AI context, those points of control are generally located upstream and downstream of model inference: in system design and training choices, deployment context, governance and monitoring, safeguards and constraints, and response protocols. Where such control points exist and are exercised carelessly, the fact that particular outputs are probabilistic or opaque does not itself defeat negligence; it shifts the analysis from output-level "why did the model do that?" to system-level "was it reasonable to deploy and operate this system in this way?[282]

Contemporary commentary has noted that innovation-driven, intangible and networked economies place pressure on traditional pure economic loss distinctions, particularly where parties are embedded in identifiable "platform" or network relationships.[283] The relevance of this literature here is limited but supportive: it reinforces that negligence analysis remains workable where responsibility is anchored in identifiable control points and confined relational settings, while leaving intact the orthodox concern to avoid indeterminate liability.

In New Zealand, economic loss is not automatically excluded from recovery in negligence. Nevertheless, its presence intensifies the court's concern with proximity and policy. Economic loss raises particular difficulties because it is often suffered by a potentially large and indeterminate class of plaintiffs, many of whom may be able to protect themselves through contract, pricing, insurance, or risk management strategies.[284] These concerns are especially acute in financial markets, where losses are an ordinary incident of trading and where

[282] Choi, above n 269.
[283] Moshood Abdussalam "Redefining Legal Responsibility for Pure Economic Loss in the Innovation Economy" (2020) 26 Auckland University Law Review 33 at 34–36 and 38–41.
[284] *South Pacific Manufacturing,* above n 280, at 296; Beever, above n 270, at 27–29.

price movements typically reflect the aggregated behaviour of many participants rather than the conduct of a single actor.

The indeterminacy problem is particularly acute in an open market because the defendant may not know who will trade during the period of distortion, whether those persons will buy or sell, the volume of their transactions, or whether the same price movement will benefit some participants while harming others. A duty framed simply as an obligation not to cause foreseeable financial loss could therefore extend to a fluctuating and potentially very large class. This distinguishes open market loss from the more confined economic loss cases in which the defendant knows the claimant, the relevant class, or the purpose for which its information or services will be used.

The English relational economic loss decision in *Spartan Steel & Alloys Ltd v Martin & Co (Contractors) Ltd* illustrates the underlying concern. Recovery was permitted only for loss consequential upon damage to the claimant's property, and denied for broader market-connected economic loss unanchored to any primary right.[285] As Beever explains, the structural difficulty is not a dislike of "economic loss" as such, but the absence of a coherent primary right infringed by the defendant's carelessness: duties in negligence imply correlative rights, and it is incoherent to use negligence to manufacture in rem rights that the law of property or contract does not recognise.[286]

If no prior right, or a responsibility analogous to a right, is engaged, the defendant's conduct, however careless, does not infringe any legally protected interest of the plaintiff. In such circumstances, liability cannot be justified. The consequence is that no duty of care arises, and a negligence claim cannot succeed, unless the case falls within a recognised exception in which the law treats the plaintiff's economic interest as deserving of protection.

Reasoning consistent with a rights-based account appears in recent Canadian authority. In *Deloitte & Touche v Livent Inc (Receiver of)*[287] and *1688782 Ontario Inc v Maple Leaf Foods Inc*,[288] the Supreme Court of Canada emphasised that negligence does not generally

[285] *Spartan Steel & Alloys Ltd v Martin & Co (Contractors) Ltd* [1973] QB 27 (CA) at 38–39 per Lord Denning MR; Beever, above n 270, at 28–33.
[286] Beever, above n 270, at 25–29.
[287] *Deloitte & Touche v Livent Inc (Receiver of)*, above n 273, at [21]–[24] and [30]–[31].
[288] *1688782 Ontario Inc v Maple Leaf Foods Inc*, above n 273, at [23]–[24] and [47]–[58].

protect a plaintiff's expectation of economic advantage absent a recognised relationship or protected interest The High Court of Australia has taken a similar position, with Edelman J observing in *Mallonland Pty Ltd v Advanta Seeds Pty Ltd* that the ordinary rule is that a person does not owe a duty of care to avoid causing reasonably foreseeable economic loss to another.[289] New Zealand courts have not adopted the rights based framework as the governing explanation of pure economic loss. The comparative authorities are nevertheless useful because their emphasis on the scope of responsibility, the interest protected and the need to avoid indeterminate liability is consistent with concerns already recognised within New Zealand's proximity and policy analysis.

Accordingly, any duty of care argument in this setting must be framed narrowly and with precision. One route is to confine the class of potential plaintiffs to those with a sufficiently close transactional nexus to the defendant's conduct (for example, direct counterparties or parties who can demonstrate a specific and identifiable interaction with the relevant system within a defined segment and timeframe). Another route, often the most doctrinally secure where economic loss is alleged, is to rely on the established controls for economic loss, particularly assumption of responsibility (frequently arising in the context of negligent misstatement) rather than attempting to ground liability solely in foreseeability.[290]

Assumption of responsibility plays a central role in New Zealand negligence law where the loss alleged is pure economic loss. It is most associated with negligent misstatement, where liability arises because a defendant with special skill or knowledge undertakes to provide information or advice to another person, knowing that the recipient will reasonably rely upon it in making economic decisions.[291] This approach has long been reflected in New Zealand authority. In *Scott Group Ltd v McFarlane* and *Boyd Knight v Purdue*, the Court of Appeal emphasised that liability for economic loss depends upon a relationship of sufficient proximity grounded in reliance and assumption of responsibility, rather than mere foreseeability.[292]

---

[289] *Mallonland,* above n 274, at [29]–[41] and [42]–[57] per Gageler CJ, Gordon, Steward, Gleeson, Jagot and Beech-Jones JJ, and at [58]–[71] and [104]–[111] per Edelman J.
[290] Beever, above n 270, at 28–33.
[291] *Hedley Byrne & Co Ltd v Heller & Partners Ltd* [1964] AC 465 (HL) at 502–503 per Lord Morris.
[292] *Scott Group Ltd v McFarlane* [1978] 1 NZLR 553 (CA) at 566–567 per Woodhouse J; *Boyd Knight v Purdue* [1999] 2 NZLR 278 (CA) at 287–289.

Although AI-driven market manipulation primarily involves conduct rather than express statements, assumption of responsibility remains relevant in delineating the boundaries of negligence liability. Ordinary market signals generated by trading activity cannot sensibly be treated as actionable representations without opening the door to indeterminate liability. However, assumption of responsibility may arise in structured relationships such as between an AI system operator and its clients, between a broker/execution provider and a client, or where a platform or intermediary supplies services and assurances in circumstances inviting reliance. In these confined contexts, the economic loss hurdle is less severe because the claimant class is limited and the responsibility nexus is clearer.

Assumption of responsibility must be assessed objectively and in context. *Attorney-General v Carter* confirms that the concept does not necessarily require a voluntary acceptance of legal liability, because tortious obligations may be imposed rather than consciously assumed. Nevertheless, the purpose for which information or a service is supplied remains critical. On the facts of *Carter*, no duty was owed to protect the plaintiffs' economic interests because the ship survey certificates were issued for safety and seaworthiness purposes, not to support the purchasers' financial reliance.[293] Courts in other jurisdictions now require a clear express or implied undertaking to find an assumed responsibility, and the High Court of Australia in *Mallonland v Advanta Seeds* confirmed that any assumption of responsibility can be negated by explicit disclaimer.[294] Therefore, in the AI trading context, it is unrealistic to suggest that a person who merely deploys an algorithmic trading system thereby tacitly assumes legal responsibility to all other market participants. Absent some explicit undertaking (for example, a contractual obligation or a regulatory commitment to others regarding the AI's behaviour), the assumption of responsibility pathway will rarely, if ever, be triggered in open-market scenarios.

Assumption of responsibility must also be confined to the purpose and scope of the undertaking. A defendant who undertakes to provide an execution service, trading recommendation or portfolio management service does not necessarily assume responsibility for every economic consequence associated with that service. The relevant question is the particular economic interest for which responsibility was undertaken. This limitation is

[293] *Attorney-General v Carter* [2003] 2 NZLR 160 (CA) at [22]–[26].
[294] *Mallonland*, above n 274, at [22] – [26].

important in AI cases because the same system may perform several functions, while the defendant's undertaking to the claimant may concern only one of them.

Nonetheless, it is conceivable that external regulatory measures could play a role. For instance, if industry regulations or exchange rules required algorithmic traders to explicitly accept certain duties or to certify their compliance with specified safeguards as a condition of market access, that might create a clearer avenue for negligence liability by effectively importing an element of assumed responsibility. Such mechanisms, however, lie in the realm of regulatory and statutory design rather than the common law of negligence, and they highlight that truly plugging the liability gap may require solutions beyond ordinary tort doctrine.

The following hypothetical examples (drawn from practical scenarios) illustrate confined settings in which negligence might plausibly operate notwithstanding the economic loss hurdle, because proximity, reliance, and assumed responsibility can be coherently identified.[295] First, a trading firm providing discretionary portfolio management services may deploy an AI system to generate trade signals or rebalance portfolios. If the firm negligently fails to validate the model, fails to monitor drift, or blindly executes AI outputs contrary to reasonable professional practice, a client suffering loss may plausibly claim negligence grounded in the firm's assumed responsibility for the client's economic interests, rather than a general duty to the market.

Secondly, an execution broker may represent that it uses AI to achieve "best execution" (for example, to minimise market impact or optimise timing). If the AI is negligently configured or monitored so that it systematically worsens execution quality, the duty analysis is more plausibly framed as a confined undertaking to the client - an economic loss claim grounded in reliance and assumed responsibility rather than as a claim based on a duty to all market participants.

Thirdly, a firm may "top up" or materially modify an AI trading system post-deployment (for example, adding new data sources, adjusting constraints, or introducing reinforcement-learning components to enhance profitability). If these changes materially degrade risk controls or create foreseeable manipulation-like patterns and the firm fails to test

[295] *Scott Group*, above n 292, at 566–567; *Boyd Knight*, above n 292, at 287–289.

or supervise the changes appropriately, liability (if any) is most coherently anchored in the antecedent human governance failures associated with the modification. Any duty in such scenario would arise from the relationships the firm has with its clients or counterparties, for example, a duty to clients to exercise reasonable care when altering the tools used to manage their investments.

Fourthly, where a firm supplies AI-assisted research, execution recommendations, or compliance assurances to clients (for example, AI-generated market outlooks, stress scenarios, or surveillance tools described as detecting manipulative conduct), negligence may be framed through the orthodox reliance route. The relevant duty is confined to the party to whom the information or service is supplied in circumstances inviting reliance and is assessed by reference to whether responsibility for that economic interest was assumed and whether reliance was reasonable.

6.2.3. Breach of Duty and Standard of Care

The second element, breach of duty, occurs when a defendant's conduct falls below the standard of care expected of a" reasonable person" in their position. Where the defendant operates in a specialised and high-risk field, such as algorithmic trading or AI system development, the applicable standard is that of a reasonably competent professional in that field. A breach of duty may arise in numerous ways, including errors in system design that ought to have been detected through reasonable testing, the absence of adequate safeguards or constraints, failures in monitoring and supervision, inadequate governance arrangements, or a failure to update systems considering known vulnerabilities. In this context, it is important to distinguish between two separate points. The first is the general negligence principle that compliance with regulatory or internal approval standards does not, by itself, establish reasonable care. The second is the narrower issue of whether reasonable care required warnings or disclosures to those who would rely on the system. These are analytically distinct.

In *Kociemba v GD Searle & Co*, the court held that regulatory approval did not exhaust the manufacturer's duty of care.[296] The relevance of this principle to AI trading lies in the broader proposition that permission to operate does not excuse careless conduct. Just as a

[296] *Kociemba v GD Searle & Co* 683 F Supp 1579 (D Minn 1988) at 1582–1585.

driver's licence permits driving but does not excuse negligent driving, regulatory approval to deploy an AI trading system does not excuse negligent design, testing, or supervision. The standard of care depends on what a reasonably competent developer or operator would do considering known risks, not merely on formal compliance.[297]

The literature has begun to engage with what "reasonable care" looks like in the algorithmic-trading context. Selbst's "Negligence and AI's Human Users" provides a detailed account of how opacity, automation bias, system interaction and organisational responsibility complicate the application of ordinary negligence principles to AI-assisted decisions. His analysis supports focusing on system design, validation, monitoring and the allocation of human responsibility, but it is not framed as an adaptation of the *Bolam* or *Bolitho* tests.[298] Hacker's "European Approach to Liability for AI" similarly argues that the standard of care for AI operators should be informed by recognised technical standards such as ISO/IEC 23053 and ISO/IEC 42001 on AI management systems.[299]

Regulatory frameworks such as the European Union's MiFID II/MiFIR regime, which require investment firms to maintain effective systems and risk controls to prevent algorithmic trading from creating disorderly markets, are relevant in articulating the content of the standard of care, even though they are not directly applicable in New Zealand.[300] They help identify what reasonable care looks like for sophisticated market participants deploying high-speed automated systems.

These sources should not, however, be treated as establishing the New Zealand common law standard of care. Technical standards, overseas regulation and accepted industry practice are better understood as evidence of the precautions that may reasonably have been available to a developer or operator. Their significance will depend on the defendant's role and the particular risk. For a developer, reasonable care may concern training data, model architecture, validation and foreseeable failure modes. For a trading firm, it may instead concern deployment parameters, pre trade limits, monitoring, model drift, updates, kill

[297] Beever, above n 270, at 25–26.
[298] Andrew D Selbst "Negligence and AI's Human Users" (2020) 100 Boston University Law Review 1315 at 1341–1356.
[299] Philipp Hacker "The European AI Liability Directives — Critique of a Half-Hearted Approach and Lessons for the Future" (2023) 51 Computer Law & Security Review 105871 at 105874–105878.
[300] MiFID II, above n 49, art 17.

switches and incident response. The standard should therefore be assessed by reference to the function over which the particular defendant exercised control rather than by imposing an undifferentiated duty on every participant in the AI supply chain.

A further point relevant to AI but not unique to it, is that AI tools can disrupt ordinary negligence analysis because errors may be difficult to detect in real time and human oversight can be impaired by interface design, automation bias, or cognitive limitations. In contexts where negligence continues to govern (including finance and financial advice), the law's focus remains on whether those deploying the technology acted reasonably in managing known categories of risk associated with its use, including validation, oversight, and response protocols.[301]

There is also a limit to treating human oversight as the solution to AI related negligence. Selbst's analysis is important in this respect because automation may alter rather than simply remove human responsibility. In high-speed algorithmic trading, a nominal human supervisor may have insufficient time or information to understand and interrupt an emerging strategy before it affects the market. Requiring a person to remain "in the loop" therefore does not necessarily amount to effective risk control. Where meaningful real time intervention is impracticable, reasonable care may instead require stronger ex ante constraints, automated monitoring and effective shutdown mechanisms. The relevant negligence question is consequently not simply whether a human was supervising the AI, but whether the overall system of supervision and control was reasonable given the speed, complexity and foreseeable risks of its operation.

6.2.4. Causation

The third element, causation, presents particular challenges in claims arising from AI-driven market manipulation. The plaintiff must establish that the defendant's breach was both the factual and legal cause of the loss suffered. The orthodox "but for" test remains the primary analytical tool, asking whether the loss would have occurred but for the defendant's negligent act or omission. In a trading environment, this enquiry is complicated by the nature of financial

[301] Selbst, above n 298, at 1319–1321, 1331–1332, 1346–1349 and 1350–1353.

markets, where prices emerge from the interaction of numerous trades, strategies, algorithms, and external factors.

The difficulty is not merely evidential. Market prices are an emergent product of complex systems, meaning that a plaintiff must identify a counterfactual price path absent the defendant's conduct, demonstrate that their own trading would have occurred on materially different terms in that counterfactual, and distinguish manipulation-induced distortion from lawful trading and background volatility. Where multiple algorithms interact and react to one another, the causal chain between a single defendant's conduct and a plaintiff's loss may be mediated by the independent decisions of other market participants and systems.

Where market prices emerge from the interaction of many algorithmic actors, the "but for" inquiry may need to be supplemented by causal doctrines familiar from other complex-system contexts. In *Fairchild v Glenhaven Funeral Services Ltd* the House of Lords accepted that, where scientific uncertainty makes "but for" causation impossible to establish, a defendant who has materially contributed to the risk of harm may be liable.[302] "The *Fairchild* exception was developed in the specific context of indivisible mesothelioma where orthodox 'but for' proof was impossible. Couch v Attorney-General concerned whether a novel duty of care could properly be struck out and did not adopt a material-increase-in-risk test for causation.[303] No current New Zealand authority establishes that the Fairchild exception applies to diffuse financial loss. Applying such an approach to AI-driven market loss would therefore require a significant doctrinal development rather than a routine application of existing law.

Sharkey and Bambauer have argued that an analogous "contribution-to-risk" framework is appropriate for AI-induced harms,[304] and Ben-Shahar has argued more generally that complex-causation problems should be resolved through evidential burden-shifting rather than denial of liability.[305] Whether New Zealand courts would extend Fairchild-style reasoning to AI-driven market manipulation is open; the more conservative

---

[302] *Fairchild v Glenhaven Funeral Services Ltd* [2002] UKHL 22, [2003] 1 AC 32 at [21] – [35].
[303] *Couch v Attorney-General* [2008] NZSC 45, [2008] 3 NZLR 725 at [32] – [37] and [118] – [130].
[304] Catherine M Sharkey "Product Liability in the Digital Age: Liability for Algorithmic Harms" in Kerry McCausland (ed) *Research Handbook on Law and Artificial Intelligence* (Edward Elgar, Cheltenham, 2024) 211 at 218–227; Derek E Bambauer "Negligence and AI's Human Users" — A Reply" (2021) 101 BU L Rev Online 99 at 102–106.
[305] Omri Ben-Shahar "Causation and Foreseeability" in Jennifer Arlen (ed) *Research Handbook on the Economics of Torts* (Edward Elgar, Cheltenham, 2013) 83 at 99–105.

pathway, on which this thesis ultimately relies, is to anchor causation upstream in identifiable human governance failures.

The contribution to risk and burden shifting approaches proposed in this literature are therefore better understood as possible responses to the evidential difficulties created by AI than as statements of existing New Zealand causation law. Their attraction lies in preventing informational opacity from operating systematically in favour of the party controlling the system and its records. Their difficulty is that relaxing ordinary proof of causation risks imposing liability where the defendant's conduct cannot be shown to have caused the particular claimant's loss. Any such departure from orthodox causation therefore requires separate justification rather than being assumed to follow from the technological complexity of AI.

The autonomy of AI can further complicate causation, particularly where an AI system adopts strategies not explicitly programmed by its designers. Nevertheless, causation need not be impossible to establish. The more defensible approach is to focus on antecedent human negligence, such as failures in system design, governance, monitoring, or intervention controls, and to link those failures to identifiable manipulation mechanisms with demonstrable market effects. In this way, causation is anchored in human conduct rather than the AI's internal reasoning.[306]

That does not, however, dispense with proof of causation. Identifying an upstream governance failure establishes a possible source of negligence, not that the failure caused the particular investor's loss. The claimant must still establish that, absent the relevant failure, the manipulative orders or strategy would probably not have occurred or continued; that those orders materially affected the relevant price or execution conditions; that the claimant traded during that distortion; and that the claimed loss resulted from the distortion rather than ordinary market movements or other independent causes.

In this respect, the appropriate causation inquiry is analogous to other negligence contexts involving complex decision-making systems. For instance, in *Joiner v Mitchell County Hospital Authority*, the hospital was held liable not because it directly caused the patient's harm, but because it failed to exercise reasonable care in selecting a competent

[306] Turner, above n 101, at 105–110.

physician; the negligence lay in the hospital's earlier act of hiring (a human decision), which in turn allowed the harm to occur. The case illustrates that causation in negligence does not require the defendant to be the immediate or sole actor causing harm; it is sufficient that the defendant's breach materially contributed to the conditions under which the harm occurred.[307]

Similarly, in AI trading cases, if a loss can be traced to identifiable failures in system governance such as inadequate safeguards, deficient monitoring, or a failure to respond to known warning signals causation may be established by linking those failures to the emergence of manipulative behaviour and consequent market distortion. The inquiry should not treat the AI system as an independent causal agent. Instead, the focus is on whether the defendant's negligence allowed the AI system to operate in a way that caused the loss. By anchoring causation in human action (or inaction), the legal test remains the familiar one of asking whether the plaintiff's loss would have occurred "but for" the defendant's breach, and whether that breach was a proximate cause of the damage.

6.2.5. Damage

The final element of negligence is damage. In the context of AI-driven trading and market manipulation, the damage alleged will ordinarily be financial loss. Where a duty of care has been established on orthodox grounds, most plausibly through a confined relationship of proximity or assumption of responsibility, the fact that the loss is economic does not bar recovery. The admissibility of pure economic loss has already been addressed at the duty stage; the damage inquiry instead concerns whether the plaintiff has suffered actual, measurable loss that is sufficiently certain and not too remote.

In trading contexts, losses will typically be quantified by reference to execution prices, foregone opportunities, increased transaction costs, or avoidable losses attributable to negligent system design, deployment, or supervision. While financial markets are inherently volatile, this does not preclude recovery where the plaintiff can demonstrate that, but for the defendant's breach, the relevant transactions would have occurred on materially different terms. Difficulties of proof or valuation do not negate the existence of damage; they bear on evidential sufficiency and quantification.

---

[307] *Joiner v Mitchell County Hospital Authority* 186 SE 2d 307 (Ga Ct App 1971) at 309–310.

Questions may also arise where several developers, operators, algorithms or market participants contributed to the circumstances in which the loss occurred. The existence of multiple actors does not prevent recovery itself, but neither can the issue be treated simply as one of apportionment or contribution. The plaintiff must first establish the necessary causal connection between the breach of each defendant pursued and the loss claimed. Questions of contribution or allocation between defendants arise only once the relevant basis of liability has been established.

The quantification of trading loss may itself be difficult. A claimant may seek the difference between the price actually paid or received and the price that would probably have prevailed absent the distortion, additional transaction costs, or another transaction specific loss. Claims for lost trading opportunities or hypothetical profits require particular caution because financial markets are inherently uncertain. Difficulty of quantification does not necessarily defeat a claim, but the claimant must provide a sufficiently reliable basis for distinguishing loss caused by the defendant's breach from ordinary market risk.

6.2.6. Comparative Case Law

There are few, if any, cases directly addressing negligence liability in the context of AI trading. Overseas cases concerning professional misuse of AI are nevertheless instructive for a limited purpose. In *Mata v Avianca Inc* (United States), and in *Ayinde v London Borough of Haringey* and *Al-Haroun v Qatar National Bank* (England), the English High Court similarly addressed the professional consequences of relying on fictitious or incorrect AI generated authorities without proper verification. These cases do not establish new principles of negligence; rather, they confirm that the use of AI does not dilute existing professional duties of care. The relevance of these decisions lies in reinforcing the principle that responsibility for the deployment and supervision of AI tools remains with human actors.[308]

By contrast, cases such as *Moffatt v Air Canada* (Canada) are better understood as turning on misrepresentation and reliance rather than on general negligence principles.

[308] *Mata v Avianca Inc* No 1:22-cv-01461 (SDNY, 2023); *Ayinde v London Borough of Haringey* [2025] EWHC 1383 (Admin); *Al-Haroun v Qatar National Bank* [2025] EWHC 1383 (Admin).

Liability arose because information was communicated in circumstances inviting reliance, and the claimant reasonably relied on that information to their detriment. This illustrates that AI-generated outputs can attract liability where assumption of responsibility and reliance are present, but it does not support a broad negligence duty for market-wide trading losses. Its precedential significance should not be overstated, but the decision is useful as an illustration of the much closer relational setting in which AI related economic loss may fit orthodox negligence principles.

Negligence therefore offers a credible but confined private law response. It is strongest where a developer, operator, broker or adviser exercises identifiable control, the claimant falls within a relationship of proximity or assumed responsibility, and breach and transaction-specific causation can be proved. It is weakest for diffuse open-market loss, where indeterminacy, pure economic loss and the opacity of interacting systems make duty and causation difficult to establish. The central difficulty is therefore not that negligence is incapable of responding to AI, but that its relational structure is poorly suited to losses arising through diffuse market interaction rather than a defined relationship between defendant and claimant. These limits explain why negligence cannot substitute for the FMC Act and provide part of the justification for the targeted duties and evidential reforms considered in Chapter 7.

6.3 PRODUCT LIABILITY

6.3.1. The Consumer Guarantees Act

New Zealand's approach to product liability has traditionally been statutory, rather than through a generalised tort of strict product liability. As Todd notes, New Zealand did not follow other jurisdictions in developing a standalone tort for defective products, preferring to address product-related harms via targeted legislation and existing tort doctrines.[309] The two key statutes are the Consumer Guarantees Act 1993 ("CGA") and Fair Trading Act 1986 ("FTA"). The CGA and FTA impose strict obligations on suppliers and manufacturers regarding product quality and truthful marketing. Although neither statute was drafted specifically for AI-driven trading systems, both use language capable of applying to digital products and services. The CGA expressly includes computer software within its definition of "goods" and defines

[309] Stephen Todd (ed) *The Law of Torts in New Zealand* (7th ed, Thomson Reuters, Wellington, 2016) at [24.1.01]– [24.1.03].

"services" broadly. Its application to any particular AI trading system nevertheless depends on whether there has been a supply of goods or services to a "consumer" within the statutory definition, and whether the transaction has validly been contracted out under s 43. The FTA raises different questions because its principal prohibitions apply to conduct in trade rather than only to supplies to a consumer.

Before examining the CGA and the FTA, it is worth situating the analysis within the broader scholarly debate over how product-liability concepts should be applied to AI systems. Three lines of scholarship are particularly relevant. First, there is what 'counts as a product' debate. In New Zealand, the threshold classification problem is narrower than in jurisdictions whose product-liability legislation was historically confined to tangible movables, because the CGA expressly includes computer software within "goods".[310] The more difficult question is not whether software can ever be a product, but whether the particular transaction supplies software as goods, supplies access or ongoing functionality as a service, or combines both elements.[311] By contrast, Twigg-Flesner and Wagner have argued that traditional product-liability regimes cannot accommodate AI without statutory reform, because the dynamism of machine-learning models destabilises the "moment of supply" baseline against which defect is conventionally assessed.[312]

Secondly, there is the 'defect' debate. Hacker criticises the EU's revised Product Liability Directive for failing to articulate what counts as a defect in a self-learning system, particularly where the system has updated itself post-supply.[313] Wendehorst, by contrast, has defended the EU's solution of treating substantial modifications and updates as a renewed moment of supply, which she argues is workable provided national courts adopt a robust burden-shifting framework.

Thirdly, opacity and technical complexity create difficulties in proving both defect and causation. *Boston Scientific Medizintechnik GmbH v AOK Sachsen-Anhalt* is relevant only by analogy. The Court held that, given the function of the medical devices, the vulnerability of

---

[310] Consumer Guarantees Act 1993, s 2(1), definition of "goods".
[311] Choi, above n 24, at 572–586.
[312] Christian Twigg-Flesner "Guiding Principles for Updating the Product Liability Directive for the Digital Age" (Briefing for the European Law Institute, 2021) at 11–17; Gerhard Wagner "Liability for Artificial Intelligence: A Proposal of the European Parliament" (2021) 28 Maastricht J Eur & Comp L 793 at 797–805.
[313] Hacker, above n 299, at 105878–105881.

patients and the abnormal potential for harm, a device belonging to a product group or production series with a potential defect could itself be classified as defective without proof that the individual device contained that defect. The decision therefore relaxed proof of individual defect, not causation generally. The claimant still had to establish compensable damage and the necessary connection between the defect and the remedial intervention.[314] The reasoning should not be treated as creating a general class-risk presumption for AI systems, still less one under the CGA, which contains no equivalent statutory presumption. Cabral's suggested relevance of class-level evidence is nevertheless useful in identifying why common model architecture, training data or updates may be evidentially important in proving non-compliance with a statutory guarantee.[315]

These debates provide the conceptual framework for assessing the CGA's application to AI trading systems. The analysis therefore turns first to the threshold question of whether the CGA applies to transactions involving such systems at all. The CGA's protections apply only to goods or services of a kind ordinarily acquired for personal, domestic, or household use or consumption (in general, "consumer" goods or services).[316] Critically, the CGA defines "consumer" not only by reference to the type of goods or services in question, but also by excluding persons who acquire goods or services for certain purposes. Section 2(1) provides that a person is not a consumer if they acquire goods or services "for the purpose of resupplying them in trade; or consuming them in the course of a process of production or manufacture or, in trade, repairing or treating other goods or fixtures on land".[317] These are specific statutory exclusions rather than a general exclusion of acquisitions motivated by financial or commercial gain. A purchaser's intention to obtain a financial benefit does not, without more, establish that the acquisition falls outside the definition of "consumer". This distinction is important when assessing AI trading systems because it cannot simply be assumed that using such a system to generate investment returns amounts to consuming it in a process of production or manufacture.

In particular, if goods are normally supplied for business use (for example, a specialised algorithmic trading system marketed exclusively to professional investors), they clearly fall

---

[314] *Boston Scientific Medizintechnik GmbH v AOK Sachsen-Anhalt* (C-503/13) EU:C: 2015:148, [2015] 3 CMLR 6 at [37] – [41].
[315] Tiago Sérgio Cabral "Liability and Artificial Intelligence in the EU: Assessing the Adequacy of the Current Product Liability Directive" (2020) 27 MJ 615 at 622–633.
[316] Consumer Guarantees Act 1993, s 2(1).
[317] Consumer Guarantees Act 1993, s 2(1) (definition of "consumer"), para (b).

outside the default scope of the Act.[318] Moreover, the CGA permits contracting out in limited circumstances. Section 43 permits contracting out only where the agreement is in writing, the goods or services are supplied and acquired in trade, all parties are in trade and agree to contract out, and it is fair and reasonable that they be bound by the exclusion. Accordingly, a supplier cannot exclude the CGA merely by describing an individual investor's acquisition as commercial. Where an AI trading system is supplied to an institutional investor or other business user acting in trade, however, s 43 may permit the parties to exclude the relevant CGA guarantees if each of the statutory requirements is satisfied.[319]

The question of whether the CGA applies to AI trading systems acquired by retail investors turns on the proper interpretation of the statutory consumer definition and the purpose for which the goods or services are acquired. The leading authority on the interpretation of "ordinarily acquired" is the Court of Appeal's decision in *Nesbit v Porter*,[320] in which the Court held that "ordinarily" is used in the Act in the sense of "as a matter of regular practice or occurrence" or in the "ordinary or usual course of events or state of things."[321] *Nesbit v Porter* establishes that the first limb of the consumer definition focuses on whether goods or services of the relevant kind are, as a matter of regular practice or occurrence, acquired for personal, domestic or household use or consumption. The test does not require personal use to be the predominant use in the market. Nor does an acquisition cease to be consumer in character merely because the purchaser also expects a commercial or financial benefit. The purchaser's actual purpose remains relevant, but through the specific exclusions in s 2(1)(b), rather than through an open-ended distinction between personal and commercial motives.[322]

The statutory definition therefore involves two distinct questions. First, the goods or services must be of a kind ordinarily acquired for personal, domestic or household use or consumption. Secondly, the particular acquisition must not fall within one of the specific exclusions in s 2(1)(b). The second inquiry should not be converted into a broader question of whether the purchaser had a commercial or profit-making motive, because s 2(1)(b) identifies particular excluded purposes rather than excluding commercially motivated acquisitions

---

[318] Consumer Guarantees Act 1993, s 2(1).
[319] Consumer Guarantees Act 1993, s 43.
[320] *Nesbit v Porter* [2000] 2 NZLR 465 (CA) at [29].
[321] *Nesbit*, above n 320, at [29]; *Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec)* [2012] NZHC 3204 at [34].
[322] *Nesbit*, above n 320, at [29] – [36]; Consumer Guarantees Act 1993, s 2(1).

generally.[323] For instance, a person who purchases building cladding for use on their own home as part of a self-directed renovation project would ordinarily be a consumer, as they are the end user acquiring the product for a domestic purpose. Conversely, if the same person purchases identical cladding for the purpose of resupplying it in trade, the acquisition falls within the express exclusion in s 2(1)(b). The purchaser's purpose is therefore relevant where it engages one of the specific statutory exclusions, rather than because a commercial motive is independently determinative of consumer status.[324]

Whilst the retail investment market is increasingly populated by algorithmic "robo-advisers", trading apps, and automated platforms offered directly to everyday consumers (for instance, AI-driven investment apps or services marketed to individuals), the AI trading systems under consideration in this thesis are predominantly designed and marketed for the purpose of facilitating trading in financial products. The core function of these systems is to assist users in making investment decisions, executing trades, and managing portfolios with the objective of generating financial returns. That function does not, however, determine whether the CGA applies. The relevant threshold remains whether systems of the particular kind are ordinarily acquired for personal, domestic or household use or consumption. A specialised algorithmic trading system marketed predominantly to professional traders may therefore fall outside the consumer definition, whereas a mass-market trading app or robo-advice service commonly supplied to individuals may present a different case.

A retail investor who acquires an AI trading system to develop trading strategies, analyse market data and engage in the ongoing buying and selling of financial products is undoubtedly using the system for the purpose of generating financial returns. It does not necessarily follow, however, that the investor is consuming the system "in the course of a process of production or manufacture" within s 2(1)(b).[325] Neither the statutory language nor the New Zealand authorities considered here establish that investing personal funds, managing a portfolio or accumulating wealth constitutes a process of production or manufacture. The stronger basis for excluding specialised trading systems from the CGA is therefore that systems of that kind may not be ordinarily acquired for personal, domestic or household use. The

---

[323] Consumer Guarantees Act 1993, s 2(1).
[324] *Sleight v Beckia Holdings Ltd* [2020] NZHC 2851 at [113] – [115]; *Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec)* [2012] NZHC 3204 at [34] – [35].
[325] Consumer Guarantees Act 1993, s 2(1)(b); *Body Corporate Number DPS 91535 v 3A Composites GmbH* [2023] NZCA 647 at [89].

position of mass-market systems supplied to retail investors is less certain and should depend on the statutory test rather than on characterising financial gain itself as a product being manufactured or produced.

This approach is consistent with the fact-specific inquiry undertaken in subsequent decisions. In *Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec)*, Duffy J considered whether shrink-wrapping services supplied in relation to a superyacht were services ordinarily acquired for personal, domestic or household use. The decision illustrates the importance of identifying the relevant kind of service and examining how services of that kind are ordinarily acquired. It does not establish a broader principle that services fall outside the CGA merely because their acquisition has a commercial or financial purpose.[326]

Similarly, in *Body Corporate Number DPS 91535 v 3A Composites GmbH*, the Court of Appeal held that building materials acquired by construction professionals for incorporation into buildings, even residential buildings owned by individuals, were not goods ordinarily acquired for personal, domestic or household use because they were consumed in a commercial or productive process and transformed into something else. The Court observed that "the owner of the building does not acquire the panels as such; rather, the owner of the building acquires a building or part of a building as a result of the materials consumed and work performed by the building professional".[327] The reasoning should, however, be confined to that context. The case concerned building materials acquired by building professionals for incorporation into buildings and does not establish a general principle that any goods or services used in an activity intended to generate an economic benefit are consumed in a process of production or manufacture. In particular, it does not determine whether personal investment or portfolio management falls within that exclusion.

By contrast, in *Sleight v Beckia Holdings Ltd*, the High Court found that project management services engaged by homeowners to oversee repairs to their own home fell within the CGA because the homeowners were genuine end users acquiring a service for a domestic purpose, and the service was not consumed within a productive or commercial process.[328] The contrast with Body Corporate reinforces that the inquiry turns on the statutory character of the

---

[326] *Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec)* [2012] NZHC 3204 at [34] – [35].
[327] *Body Corporate Number DPS 91535 v 3A Composites GmbH* [2023] NZCA 647 at [89] – [92].
[328] *Sleight v Beckia Holdings Ltd* [2020] NZHC 2851 at [113] – [120].

particular goods or services and the circumstances of their acquisition, rather than simply on whether the activity has financial consequences.

These authorities demonstrate that consumer status depends on the statutory character of the goods or services and the application of the specific exclusions in s 2(1)(b), rather than on a general distinction between personal and commercially motivated activity. As Davison J observed in *Body Corporate No DPS 91535 v 3A Composites GmbH*, the consumer is "a person who is the 'end user' of the goods, and the provisions of the Act imposing the statutory guarantees are directed at the kind of goods and services 'ordinarily acquired for personal, domestic, or household use or consumption'."[329] The analogy with building materials should therefore not be extended directly to retail investment. A specialised AI trading system may fall outside the CGA because systems of that kind are ordinarily acquired for professional rather than personal use, but the same conclusion cannot necessarily be drawn for a consumer-facing trading or investment tool merely because its user seeks a financial return.

It is conceivable that an AI system might be used in a manner that falls within the CGA's scope, such as where a person acquires a financial planning tool for the one-off purpose of assessing whether their existing savings are adequate for retirement, without any intention to engage in active trading or ongoing investment management. In such a scenario, a stronger argument may arise that the tool is of a kind ordinarily acquired for personal use, particularly where comparable financial planning services are commonly supplied to individuals for personal financial management. However, such cases would be highly fact-specific, and the distinction has not been tested in New Zealand courts in the context of AI financial products.[330]

A specialised AI trading system supplied to an institutional investor or used internally by a financial firm is unlikely to engage the CGA where systems of that kind are ordinarily acquired for professional rather than personal, domestic or household use. This highlights how the CGA's reach depends principally on the statutory character of the supply rather than on the use of AI technology itself. The uncertainty about whether the CGA applies to AI trading systems used by retail investors for wealth accumulation becomes policy relevant as similar AI

---

[329] *Body Corporate Number DPS 91535*, above n 327, at [52] – [56].
[330] Consumer Guarantees Act 1993, s 2(1); *Sleight v Beckia Holdings Ltd* [2020] NZHC 2851 at [113]– [115]; *Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec)* [2012] NZHC 3204 at [34]– [35].

tools migrate between retail and wholesale contexts, raising the question of whether consumer protection principles should inform commercial law in this area.

The question of whether a retail investor using an AI trading system for active trading can qualify as a consumer under the CGA appears not to have been directly tested by New Zealand courts. The answer is likely to depend on the particular market for the system. Bespoke or institutional trading systems are unlikely to satisfy the requirement that goods or services of that kind be ordinarily acquired for personal, domestic or household use. Consumer-facing trading apps, robo-advice services and similar products may be more difficult to exclude categorically where they are commonly supplied to individuals. The CGA's protection of AI-related investment loss is therefore uncertain rather than necessarily non-existent. This uncertainty itself exposes a limitation in relying on general consumer law to address AI-driven trading risks and supports consideration of the more targeted regulatory framework discussed below.

The CGA applies to both goods and services, but the content of the guarantees differs. Determining which category an AI trading system falls into can be complex. The CGA defines goods as "personal property of every kind (whether tangible or intangible)," explicitly including computer software.[331]Accordingly, an AI-driven trading system supplied as computer software may constitute goods for CGA purposes. By contrast, where the user receives continuing access to cloud-based functionality, algorithmic execution or portfolio management rather than software supplied as such, the transaction may instead involve services. If, on the other hand, the AI system is provided as a service (for example, a cloud-based algorithmic trading service, or an AI-driven portfolio management service), then the CGA's service guarantees (such as the guarantee of reasonable care and skill) may apply if the transaction falls within the CGA's scope.[332]

In practice, many AI trading products are hybrid: they comprise software tools, continuous digital services (like data feeds or updates), and possibly human support. Where the CGA applies, a consumer's rights under the CGA can arise in respect of both aspects. For instance, a retail investor might purchase a subscription-based AI trading service that uses a

[331] Consumer Guarantees Act 1993, s 2(1) (definition of "goods").
[332] Consumer Guarantees Act 1993, ss 28 and 29.

software algorithm; the reliability of the software itself can be judged under the guarantees for goods, while the overall service (including any human oversight or auxiliary services) is subject to the guarantees for services. In any event, the CGA's remedies do not depend on a classical sale of physical goods, a recognition by the legislature that even in 1993 "goods" could be as intangible as software code. Whether those guarantees are available to a retail investor will nevertheless depend on the threshold consumer definition discussed above, including whether the particular goods or services are of a kind ordinarily acquired for personal, domestic or household use.

In the limited circumstances where the CGA does apply (for example, where a court determines that the system is ordinarily acquired for personal, domestic or household use and the investor is an "end user" within the meaning of the Act), the cornerstone of the CGA's protection is the guarantee of "acceptable quality" in section 6. This guarantee means that goods sold to a consumer must be of a standard that a reasonable consumer would find acceptable, considering factors such as the nature of the goods, price, and any representations made about them.[333]

Section 7 of the Act elaborates that "acceptable quality" in relation to goods (including digital goods like software) means the goods are: (a) fit for all the purposes for which such goods are commonly supplied, (b) of acceptable appearance and finish, (c) free from minor defects, (d) safe, and (e) durable, as judged by the reasonable consumer with full knowledge of the goods' condition and any statements about them.[334]

For AI-driven trading systems, were the CGA to apply, these criteria would require translation to the digital, algorithmic context. An app that automatically executes trades should reliably perform the functions for which goods of that kind are commonly supplied; hidden software flaws may be relevant to freedom from defects; and the system should remain functional for the period a reasonable consumer would regard as acceptable having regard to its nature and surrounding circumstances. Whether financial trading loss falls within the separate concept of "safety" requires greater caution and is considered below.

---

[333] Consumer Guarantees Act 1993, s 6.
[334] Consumer Guarantees Act 1993, s 7(1).

These applications of the CGA standards may be novel, but they spring from familiar principles. The fit for purpose guarantee, for instance, has long been applied to products whose performance can only be assessed over time or in variable conditions. The fact that an AI system learns or evolves does not exempt it from the requirement of fitness for its represented purpose; if anything, continuous learning might be considered part of the "nature of the goods" that shapes a reasonable consumer's expectations.

A more difficult question is whether "safe" in s 7(1)(d) extends to protection against pure financial loss caused by the operation of an AI trading system.[335] The CGA does not define "safe", and the existing New Zealand authorities do not clearly establish a distinct concept of economic safety. It is therefore preferable not to characterise every financially harmful output as a safety defect. A trading system that executes unauthorised or erratic transactions because of a software defect may instead fail the acceptable quality guarantee because it is unfit for the purposes for which goods of that kind are commonly supplied or because it contains defects. Financial loss may then become relevant to the remedies available for that failure. This distinction avoids extending the statutory concept of safety beyond the present authorities while still allowing the CGA's existing guarantees to respond to defective digital performance.

The CGA's durability requirement, in the context of AI, translates to the notion that the system should not rapidly or continuously degrade to an unfit state.[336] One challenge with AI trading algorithms is that their performance can change over time as market conditions shift or as the model "learns" (or drifts). However, these changes are not wholly mysterious or unforeseeable. Developers are aware that models require updates or retraining to maintain performance. If an AI trading system ceases to perform the function for which it was bought (for example, reliably executing a certain strategy) because the model degrades or becomes obsolete after a short period, a consumer could argue the product was not as durable as a reasonable consumer would expect in the unlikely event that the CGA were held to apply.

It might appear that the notion of "durability" is inappropriate for software. Yet the CGA's flexible standard is precisely capable of adaptation: just as durability of a second-hand car is assessed considering its age and mileage,[337] the durability of an AI service can be

[335] Consumer Guarantees Act 1993, s 7(1)(d).
[336] Consumer Guarantees Act 1993, s 7 (1)(e).
[337] *Drummond v Van Ingen* (1887) 12 App Cas 284 (HL) at 297 per Lord Macnaghten.

assessed considering how long a reasonable user would expect the system to remain functional and supported. Moreover, the CGA's guarantee of ongoing compliance (and available repairs or refunds if a product fails prematurely) arguably encourages proactive maintenance: if continuous updates are necessary to keep the AI system functioning safely, suppliers should provide them to avoid breaching the durability guarantee where the CGA applies.

Apart from the general guarantee of acceptable quality, the CGA contains separate fitness guarantees for goods and services. Section 8 applies to goods acquired for a particular purpose in the circumstances specified by that provision, while s 29 applies where a consumer makes known the particular purpose for which services are required or the result, they wish the services to achieve.[338] In the context of AI trading, in the unlikely circumstances where the CGA might apply, it is foreseeable that consumers will rely on suppliers' expertise and representations to select appropriate products. If a consumer tells an AI software vendor that they need a tool to perform a specific trading strategy or risk management function, and the vendor recommends an AI product, the law guarantees that the product will be fit for that particular purpose (unless the consumer unreasonably refused to rely on the vendor's advice). This creates a strong incentive for vendors to accurately represent what their AI systems can and cannot do and indeed overlaps with the FTA's requirements not to misrepresent a product's characteristics.

One of the concerns sometimes raised about AI liability is the so-called "black box" problem: AI systems (particularly those using machine learning) may be so complex or opaque that it is extremely difficult to explain why they made a particular decision. In the unlikely cases where the CGA were to apply, the statutory guarantees create a form of no-fault liability for suppliers (and, in some cases, manufacturers).[339] Unlike negligence, the statutory guarantees do not require the consumer to establish that the supplier or manufacturer failed to exercise reasonable care or to identify a particular negligent coding decision. This reduces the significance of the black box problem, but the claimant must still establish that the goods or services failed to comply with the relevant statutory guarantee.

[338] Consumer Guarantees Act 1993, ss 8 and 29.
[339] Consumer Guarantees Act 1993, ss 6–17.

An adverse AI output may provide evidence that the relevant statutory guarantee has not been met, but it does not necessarily establish breach by itself. A trading loss or unsuitable recommendation may result from a defect in the system, but it may also reflect ordinary market risk, user inputs, third-party data or disclosed limitations of the technology. The focus under the CGA is on whether the goods or services complied with the applicable statutory standard rather than on whether the producer acted negligently. In this respect, the claimant does not face the same need to establish producer fault as in negligence. However, where consequential loss is claimed, the claimant must still establish the necessary connection between the failure to comply with the guarantee and the loss suffered.

Another challenge posed by AI trading systems is the distributed nature of their development and operation. An AI trading product might involve a software platform from one provider, a predictive model from another, data streams from several sources, and user-specific customisations. When something goes wrong, this raises the question: who is responsible? New Zealand's consumer law addresses this through pragmatic rules on supplier and manufacturer liability. Under the CGA, where the statutory guarantees applying to goods supplied to a consumer are not complied with, the consumer may pursue the remedies provided against the supplier without having to prove that the supplier personally caused the underlying technical defect. This is significant for AI systems because the immediate supplier may not have developed the model, supplied the data or designed the relevant component.[340] The supplier, in turn, may have rights of indemnity or contribution against those upstream parties, but the consumer is not required to sort out the chain of contracts or technology providers.

The CGA also provides direct remedies against manufacturers in relation to specified guarantees concerning goods. These include acceptable quality, compliance with certain descriptions, the availability of repair facilities and parts, and express guarantees.[341] Whether a software developer, model trainer or other upstream participant is directly liable as a manufacturer will depend on whether that participant falls within the CGA's statutory definition of "manufacturer". Participation in the development of an AI system does not itself establish manufacturer status. It is true that AI supply chains add layers of complexity: liability may be factually fragmented and contested among multiple companies (for instance, the data

[340] Consumer Guarantees Act 1993, s 18(1).
[341] Consumer Guarantees Act 1993, ss 12–14 and 25–27.

provider might blame the algorithm developer, who in turn blames the system integrator). Although distributed responsibility is not unique to AI, AI systems may intensify the problem because software, models, data, updates and deployment decisions can originate from different actors and change over time. Complex products with components from multiple sources have long raised issues of apportioned responsibility, and legal systems have tools such as indemnity claims and joint-defendant procedures to handle them.[342]

6.3.2 The Fair Trading Act

The Fair Trading Act 1986 offers a parallel avenue of redress that can be especially relevant for AI trading tools. Section 9 broadly prohibits misleading or deceptive conduct in trade. Liability does not depend on an intention to mislead; the question is whether the conduct, viewed objectively and in context, was misleading or deceptive or likely to mislead or deceive. Sections 10 and 11 address conduct liable to mislead the public about specified characteristics or suitability of goods and services respectively, while s 13 prohibits specified false or misleading representations about matters including standard, quality, performance characteristics and benefits.[343]

If an AI trading system is marketed as having a "proven algorithm" or "guaranteed strategy" which it in fact does not, or if promotional materials misstate its capabilities or risks, the maker or vendor may be liable under the FTA. Unlike a negligence claim, liability for misleading conduct does not depend on establishing a failure to exercise reasonable care, and an FTA claim does not necessarily depend on the existence of a contract. For example, a developer claiming "our AI has a 90% success rate" or "this trading bot outperforms the market" when it does not can face enforcement action by regulators or a civil claim by purchasers on the basis that these statements were misleading. In the context of AI, misrepresentation scenarios might include claiming that a trading algorithm has been "extensively tested" or "uses proven strategies" if those claims are unsubstantiated.

---

[342] Contract and Commercial Law Act 2017, s 48.
[343] Fair Trading Act 1986, ss 9–13; *Red Eagle Corporation Ltd v Ellis* [2010] NZSC 20, [2010] 2 NZLR 492 at [26] – [29].

The applicable provision does not depend on characterising every AI system solely as a good. Section 10 concerns conduct liable to mislead the public about the nature, manufacturing process, characteristics, suitability for a purpose or quantity of goods, while s 11 performs a corresponding function for services. The distinction may depend on whether the system is supplied as downloadable software, a hosted trading service, or a combined product and service.[344]

However, the FTA's remedies differ from the CGA. While the CGA gives consumers direct rights to demand repair, replacement, or refunds (and to recover consequential losses) from suppliers, the FTA is enforced either by regulatory action or through courts in civil proceedings (often by a class of affected consumers). The court has discretion to grant a range of orders, including damages, contract voiding or variation, and injunctions.[345] Thus, for an individual claimant seeking compensation, the FTA can be powerful when misrepresentations are evident, but it does not guarantee the kind of on-the-spot remedies that the CGA provides. A defect or harmful output does not, by itself, establish misleading or deceptive conduct. The FTA becomes relevant where the supplier's or developer's words, demonstrations, disclosures or other conduct convey a misleading impression about the system, or where a known limitation or change in performance renders an earlier representation misleading in context. The FTA is therefore not a general product-performance regime, but neither is it confined to statements made before sale.

For a private claimant, proving a contravention is not sufficient by itself. Section 43 requires loss or damage suffered "by" the prohibited conduct. In *Red Eagle Corporation Ltd v Ellis*, the Supreme Court separated the objective breach inquiry under s 9 from the remedial inquiry under s 43. A claimant seeking compensation must show that the claimant was actually misled and that the contravening conduct was an effective or operating cause of the loss, although it need not be the sole cause.[346] The claimant's own carelessness may be a sole or contributory operating cause, and the court retains a discretion to reduce or decline relief where appropriate. Available orders include avoiding or varying a contract, refunding money, supplying specified goods or services, and paying the amount of the loss or damage. Section

[344] Fair Trading Act 1986, ss 2, 10 and 11.
[345] Fair Trading Act 1986, ss 40–43.
[346] *Red Eagle Corporation Ltd v Ellis* [2010] NZSC 20, [2010] 2 NZLR 492 at [26]–[31].

43A requires an application to be brought within three years after the loss or damage, or its likelihood, was discovered or ought reasonably to have been discovered.[347]

In practice, the FTA may apply where developers or vendors make misleading representations about matters such as back-testing methodology, historical performance, expected return, accuracy, risk controls, data quality, model limitations, the degree of autonomy, the presence of human oversight, or the system's ability to detect or prevent manipulative trading. Section 13 may apply where those matters amount to representations about the system's standard, quality, performance characteristics or benefits. Section 12A is also relevant where a trader makes a representation without reasonable grounds at the time it is made, although only the Commerce Commission may commence proceedings in relation to a contravention of that section.[348] Claims about future performance must therefore be framed carefully: the issue is not merely whether the forecast proved wrong, but whether the trader had reasonable grounds for it and whether the overall representation was misleading when made.[349]

An omission will not invariably be misleading, because the FTA does not impose a general duty to disclose every product limitation. However, silence, a half-truth or an inconspicuous qualification may form part of misleading conduct where, viewed as a whole, it creates or preserves a false impression. This is particularly relevant where headline claims about returns, accuracy, autonomy or risk controls are materially qualified only in technical documentation that an ordinary user is unlikely to connect with the claim.[350]

The FTA occupies a complementary position alongside the CGA, contract, negligence and the Financial Markets Conduct Act 2013 ("FMC Act"). The CGA focuses on statutory guarantees attaching to qualifying supplies of goods or services. Contract addresses obligations assumed between the parties, while negligence addresses unreasonable conduct causing legally recoverable harm. The FTA instead addresses misleading or otherwise prohibited conduct in trade and may apply even where no contractual warranty has been breached. For conduct

---

[347] Fair Trading Act 1986, ss 43 and 43A; *Red Eagle Corporation Ltd v Ellis*, above n 346, at [34]–[39].
[348] Fair Trading Act 1986, ss 12A–12C and 13(a), (b) and (e).
[349] Fair Trading Act 1986, s 12B.
[350] Fair Trading Act 1986, s 2(2)(a); *Godfrey Hirst NZ Ltd v Cavalier Bremworth Ltd* [2014] NZCA 418, [2014] 3 NZLR 611 at [59]–[60] and [71]–[84].

concerning financial products, financial services or market manipulation, the FMC Act is generally the more directly aligned statutory regime. That does not, however, exclude the FTA's potential application to representations made by a developer or vendor about the AI trading system itself. Sections 48P–48S coordinate the two regimes: s 48P requires the Commerce Commission to obtain the FMA's consent before commencing proceedings under ss 9–13 concerning a financial product or financial service, while s 48S prevents conduct contravening specified FMC Act provisions from also contravening ss 9 - 13.[351]

The most challenging category of harm from AI trading systems is pure economic loss on the part of users or third parties – for example, lost investment value or broader market impact not accompanied by any physical damage. Neither the CGA nor the FTA is expressly limited to physical harm; the CGA in fact allows recovery of reasonably foreseeable consequential economic losses from a failure of a guarantee.[352] Yet, certain types of financial harm may not fit neatly into the current framework. The CGA's guarantees were conceived with the paradigm of a product being "defective," a notion that maps readily onto, say, a toaster that bursts into flames or a car that fails to start, but less easily onto an AI that triggers a market anomaly. In extreme cases, an AI trading algorithm might cause losses to a broad class of investors or market participants by contributing to a flash crash or by engaging in manipulative behaviour. Such harms may not be traceable to a "defect" in the traditional sense: the AI might have operated as designed but in a way that harms market integrity or investor interests.

Even if a retail investor's own AI trading tool malfunctioned (for example, if it executed trades far outside the specified parameters), whether the CGA applies would depend on whether the investor satisfies the statutory definition of "consumer" in the circumstances of the particular acquisition. However, if the complaint is that an AI trading firm's algorithm caused a general market distortion that affected others, the CGA provides no remedy to third parties (who did not purchase the product), and an FTA claim would require misleading or otherwise prohibited conduct and, for a private claim for loss, the necessary causal connection between that conduct and the claimant's loss. Simply put, these consumer law regimes are not well-equipped to handle systemic economic harms or harms that do not arise from a discernible product failure or misstatement.

---

[351] Fair Trading Act 1986, ss 48P–48S; Financial Markets Conduct Act 2013, pt 2 and pt 5 subpt 2.
[352] Consumer Guarantees Act 1993, s 18(4).

These gaps underscore that whilst New Zealand consumer law may provide some protection in limited circumstances where the CGA applies, it does not fully address all AI-driven harms, especially those that resemble securities law or competition law issues more than traditional product defects. In those situations, public law (financial markets regulation) and common law (fraud or negligence) may become relevant, each bringing its own limitations (for example, negligence for pure economic loss faces restrictive duty rules). Given these significant gaps and limitations, there is a compelling case for developing a bespoke regulatory framework for AI-driven financial products under the FMC Act to address more directly the specific characteristics and risks of AI-driven activity in financial markets.[353]

New Zealand's consumer protection model shares a common heritage with other jurisdictions (notably the UK and Australia). Developments in the European Union's product liability regime are therefore instructive as to how the law can evolve to address AI risks. In 2023–24, the EU moved to modernise its decades-old product liability directive specifically to accommodate digital technology and AI. The Product Liability Directive 2024 introduces several important changes: it expands the definition of "product" to encompass software, including AI system software, ensuring that intangible digital products are clearly subject to strict liability for defects.[354]

The Directive does not, however, extend the harmonised product liability regime to stand-alone pure economic loss. Recoverable damage includes death or personal injury, specified damage to property, and destruction or corruption of data not used for professional purposes. Accordingly, trading or investment loss, without one of the recognised forms of damage, does not fall within the Directive's harmonised liability regime.[355] Perhaps most notably, the new EU rules empower courts to alleviate the burden of proof for claimants in complex technology cases. Where technical or scientific complexity makes it excessively difficult for a claimant to prove defectiveness or causation, the Directive permits rebuttable presumptions where the specified statutory conditions are satisfied.[356] This may reduce some of the evidential difficulties associated with the "black box" problem by recognising the

[353] Jessica C Lai "The Consumer Guarantees Act and Digital Products" (2019) 50 Victoria University of Wellington Law Review 385 at 391–398.
[354] Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products [2024] OJ L, art 4(1).
[355] Directive (EU) 2024/2853, above n 354, art 6.
[356] Directive (EU) 2024/2853, above n 354, arts 9-10.

informational disadvantage that a claimant may face in establishing how a complex digital product failed. These reforms reflect a broader concern in the literature that conventional product liability rules may place significant evidential burdens on claimants where the operation of complex AI systems is difficult to reconstruct or explain.[357]

Across the Atlantic, the United States has seen an emerging trend of courts finding ways to apply product liability principles to software-driven services, particularly where harm results from design features rather than content. Two prominent examples illustrate this. In *Garcia v. Character Technologies Inc*, a federal court allowed a strict product liability claim to proceed against the developer of an AI chatbot, treating the chatbot's algorithm and interface as a product for purposes of a defective design claim (notwithstanding arguments that the AI's output was mere "speech").[358] Similarly, in *Lemmon v. Snap Inc* (9th Cir. 2021), the US Ninth Circuit Court of Appeals held that a claim of a design defect in a social media app's "Speed Filter" (which allegedly encouraged reckless driving) was not barred by the immunity usually granted to online content publishers, because the claim targeted the design of a feature of the app (a product component) rather than any third-party content.[359]

These cases suggest that, even absent new legislation, courts can adapt traditional product liability concepts to hold AI system providers accountable; for example, by distinguishing between the outputs of an AI (such as content or decisions, which might be considered beyond the scope of product liability or protected by free-speech principles) and the design of the AI system itself (which can be subjected to scrutiny under product safety standards).While the US and New Zealand legal systems differ in significant ways (for instance, New Zealand does not have an equivalent to the US regime of tort-based strict liability for products), these developments are indicative of a broader willingness to reconsider legal categories (like the definition of "product" or "defect") in light of new technology. New Zealand courts, when faced with cases involving AI, may look to such overseas precedents for inspiration in interpreting domestic statutes such as the CGA and FTA.[360]

---

[357] Miriam C Buiten "Product Liability for Defective AI" (2024) 57 *European Journal of Law and Economics* 239.
[358] *Garcia v Character Technologies Inc* No 6:24-cv-01903 (MD Fla, 21 May 2025) at [12] – [20].
[359] *Lemmon v Snap Inc* 995 F 3d 1085 (9th Cir 2021) at 1090–1093.
[360] See generally Lai, above n 353, at 391–398 (on the adaptability of New Zealand consumer law to digital products).

6.4. CHOOSING BETWEEN NEGLIGENCE AND PRODUCT LIABILITY

The preceding sections have examined negligence and product liability as distinct doctrinal frameworks. This section draws them together to assess how each framework aligns with the realities of AI-driven trading systems and to identify the contexts in which one may be more appropriate than the other. The aim is not to advocate for a singular approach, but to clarify how each framework functions, what it demands of claimants and how it distributes responsibility across the actors involved in AI development and deployment.

In the AI trading context, the choice between negligence and product liability is not merely a matter of legal form. It reflects deeper assumptions about the nature of AI systems and the roles of human actors. Negligence presupposes that harm results from a failure to take reasonable care, and it is best suited to contexts where human oversight, professional judgment, or governance failures are central to the harm. Product liability, by contrast, treats the AI system as a product and focuses on whether it meets objective standards of quality and safety, regardless of the supplier's fault.[361] However, as analysed in sections 6.3.1 and 6.3.2, the applicability of product liability frameworks to AI trading systems used for wealth accumulation is highly constrained. The CGA is unlikely to apply where the system is consumed in a commercial or productive process, and the FTA, whilst applicable to misleading representations, is structurally unsuited to addressing the complex design and performance issues that arise with AI financial products.

The following table summarises the key distinctions between the two frameworks as they apply to AI trading systems:

Feature	Negligence	Product Liability (CGA / FTA)
Liability Basis	Fault-based: requires proof of breach of duty of care (e.g., poor coding or lack of oversight) by a human actor.	CGA: Statutory liability for failure to comply with applicable consumer guarantees, where the AI trading system and its acquisition fall within the Act. Coverage depends on the statutory definition of "consumer"

[361] Selbst, above n 298, at 1367–1369.

		and the circumstances of the acquisition. FTA: Liability for misleading conduct but structurally unsuited to AI design/performance issues.
Applicability	Broad: Can reach anyone in the chain (developers, operators, intermediaries, firm directors, or brokers).	CGA: Potentially applicable to qualifying consumers, but coverage is uncertain for bespoke or professional AI trading systems and depends on the nature and purpose of the acquisition and any valid contracting-out arrangement. FTA: Broad; applies to any "misleading conduct in trade," but does not address design defects or systemic risks.
Standard of Proof	High: Must prove the defendant failed to act as a "reasonable AI developer" would (establishing duty, breach, causation, and damage).	CGA: Lower evidentiary burden if applicable, but the threshold issue of consumer status is determinative. FTA: Claimant must prove a misleading representation; it does not address product performance absent misrepresentation.
Scope of Harm	Primarily personal or property damage. Recovery for trading losses (pure economic loss) is very difficult.	CGA: May allow recovery of consequential economic loss where the Act applies but does not provide a general remedy for losses suffered by third parties or for market-wide harm. FTA: Covers loss from misrepresentation but not design defects or performance failures.

Evidentiary Burden	High: Requires technical evidence of fault, often difficult with opaque AI systems.	CGA: Lower burden if applicable, but the consumer-status requirement is a significant hurdle. FTA: Moderate; must show objectively misleading conduct rather than proving negligence.
Human Agency	Central: Liability attaches to human actors who failed to exercise reasonable care.	CGA: Liability attaches to the supplier/manufacturer because the product failed (if the CGA applies), regardless of fault. FTA: Focuses on the conduct of the representor in trade.
Systemic Harm	Difficult: Liability for "flash crashes" or market-wide errors is often limited by policy.	CGA: Very restricted - does not apply to market-wide harms; FTA: Limited to misrepresentation claims; neither addresses systemic market instability.
Remedies	Damages for proven loss; very rare to see exemplary/punitive damages.	CGA: Statutory remedies, including repair, replacement, refund and damages for consequential loss, may be available where the Act applies, but coverage is uncertain for professional, institutional and bespoke AI trading systems. FTA: Regulatory or court-ordered remedies (injunctions, damages, contract variation), but these address only misleading conduct.

This comparative analysis reveals significant gaps in the current frameworks when applied to AI trading systems used for wealth accumulation. Negligence remains constrained by evidentiary burdens and doctrinal limits on pure economic loss. The CGA, whilst offering strict liability protections for consumers, is unlikely to apply to retail investors using AI

systems for active trading, as such use involves consuming the system in a commercial or productive process rather than for personal, domestic, or household purposes. The FTA provides a mechanism for addressing misleading representations but does not extend to product design, algorithmic performance, or systemic risks.

The result is uneven rather than categorically absent protection. Computer software falls within the statutory definition of "goods", and a retail purchaser may be a "consumer" where the system is of a kind ordinarily acquired for personal, domestic or household use and is not acquired for one of the excluded purposes in s 2. The fact that software is used to make investments or accumulate wealth does not, by itself, establish that it is being consumed in a process of production or manufacture. By contrast, bespoke institutional platforms and professional trading systems are less likely to satisfy the consumer definition, and acquisitions made in trade may also be subject to a valid written contracting-out provision under s 43. The more precise limitation is therefore that CGA coverage depends on the nature of the system, the ordinary purpose for which that kind of system is acquired, the status and purpose of the acquirer, and any valid contracting-out arrangement. Even where the CGA applies, its bilateral consumer remedies do not readily address losses suffered by third parties or market-wide disruption.

The gaps identified in this analysis point toward the need for regulatory reform that is specifically tailored to AI-driven financial products and would address both the unique characteristics of AI systems (complexity, opacity, continuous evolution) and the nature of the harm they can cause (predominantly pure economic loss). Chapter 4 addressed market manipulation reforms under the FMC Act; the product liability and governance gaps identified here suggest that further regulatory innovation may be required, as discussed in Chapter 7.

Ultimately, the choice of framework should be guided by the nature of the relationship between the parties, the type of harm suffered, and the pathways to accountability that the law currently recognises. Negligence focuses on who made a preventable mistake and requires proof of a breach of a duty of care owed to the plaintiff. Product liability (under the CGA/FTA) focuses on what went wrong with the product or how it was marketed, and imposes liability based on the outcome (the product's failure or the falsity of claims) rather than on proof of carelessness. However, as this analysis has shown, the CGA is unlikely to apply to AI trading

systems used for wealth accumulation, and the FTA addresses only a narrow subset of potential harms (misleading representations).

Each regime has its advantages and blind spots: negligence can reach a wider range of defendants (anyone in the supply or oversight chain) but may be harder to establish (especially for pure economic loss without a special relationship), whereas the CGA offers powerful protections in theory but is practically inaccessible to retail investors using AI trading systems for wealth accumulation, and the FTA addresses only misleading conduct, not design defects or performance failures.

Courts and policymakers must remain attentive to the evolving capabilities and risks of AI in trading, and ensure that liability frameworks remain responsive, coherent, and fair. As this analysis has demonstrated, existing frameworks do leave significant gaps: retail investors using AI trading systems for wealth accumulation fall outside the CGA's scope, the FTA does not address design defects or performance failures, and negligence law imposes restrictive barriers to recovery for pure economic loss. These gaps warrant serious consideration of law reform by recognising new forms of liability tailored to AI-related harms in financial markets. Whilst Chapter 4 examined market conduct reforms, Chapter 7 will consider complementary reforms addressing product liability and governance accountability for AI systems.

6.5. CONCLUSION

This chapter has examined the two principal doctrinal strategies through which New Zealand tort law may respond to harm caused by AI-driven trading systems: negligence and product liability. Rather than focusing on remedies, the analysis has centred on the conditions under which each liability framework applies, the evidentiary and conceptual challenges they face, and the extent to which they can accommodate the distinctive features of AI technologies, particularly opacity, autonomy, and distributed responsibility.

Negligence is a constrained basis of liability in this context. Its continued relevance lies in its capacity to assign responsibility to human actors: developers, operators, and firms who exercise meaningful control over the design, deployment, and governance of AI systems. The analysis has shown that negligence can operate effectively where harm arises from identifiable

failures to exercise reasonable care at specific control points, such as inadequate testing, poor oversight, or negligent system modification. However, the doctrinal limits of negligence become apparent when the harm alleged is pure economic loss suffered by a diffuse class of market participants.

In such cases, New Zealand courts will insist on a confined relational setting, usually grounded in an actual assumption of responsibility or a similarly proximate relationship, to justify the imposition of a duty of care. The law's traditional reluctance to recognise open-ended duties for economic loss, particularly in financial markets, reflects enduring concerns about indeterminate liability and about preserving coherence with contract and statutory frameworks. Indeed, as discussed, courts in overseas jurisdictions have increasingly affirmed that without a violation of a protected right or a clear undertaking of responsibility, there is no negligence liability for causing someone to lose financial expectancies. If New Zealand follows that path, as it arguably should to maintain consistency with fundamental principle, then negligence law will have little role to play in addressing AI-driven market losses except where a special factual matrix brings the case within a tested exception.

Product liability, through the CGA, may provide a statutory basis for liability where an AI trading system and its acquisition fall within the Act. The CGA's principal limitation is not that every investment-related use is necessarily non-consumer use. Its application turns on the statutory definition of "consumer" and the circumstances of the acquisition. Many professional and institutional deployments may fall outside that definition or may be subject to valid contracting out under s 43. Even where the Act applies, the claimant must establish a failure of an applicable statutory guarantee and recoverable loss, and the regime does not provide a general product liability remedy for third parties affected by an AI-driven market event.

The FTA provides a separate avenue by targeting misrepresentations about AI capabilities, which are common in the marketing of such products, but it does not address design defects, performance failures, or systemic risks. Moreover, the FTA was not designed with dynamic technologies like AI in mind and is structurally unsuited to addressing the complex challenges posed by AI-driven financial products.

Taken together, negligence and product liability offer limited and incomplete strategies for governing AI-driven trading systems. Negligence provides a fault-based mechanism for

assigning responsibility in professional and institutional contexts but is limited by evidentiary burdens and doctrinal constraints on economic loss. Product liability through the CGA may provide meaningful protection where its statutory requirements are satisfied, but its coverage is uncertain for professional, institutional and bespoke AI trading systems and it does not provide a general remedy for third-party or market-wide loss. The FTA addresses misleading representations but not design defects, performance failures, or systemic risks.

The chapter has therefore argued that the choice between these frameworks should not be seen as binary. Rather, each should be applied where its doctrinal structure and policy rationale are best aligned with the nature of the harm, the relationship between the parties, and the role of human agency in the deployment of AI. However, the analysis has also revealed significant gaps where AI trading systems fall outside the CGA, where negligence is constrained by the rules governing pure economic loss, or where harm extends beyond the immediate contractual or consumer relationship, creating a case for targeted regulatory reform.

Comparative developments, particularly in the European Union and the United States, suggest a growing recognition that existing liability frameworks must evolve to address the challenges posed by AI. The EU's revised Product Liability Directive 2024 offers models for burden-shifting, expanded definitions of “product” and “damage,” and enhanced access to evidence. These reforms aim to preserve the compensatory and deterrent functions of tort law in the face of technological opacity and complexity. New Zealand's legal system, while already equipped with flexible and technology-neutral statutes, would benefit from considering similar reforms.

Ultimately, the law’s response to AI-driven trading systems must balance innovation with accountability. Negligence and product liability, properly understood and applied, address certain categories of harm but leave significant gaps in protection, particularly where AI systems fall outside existing consumer protections or cause loss beyond the immediate relationship between supplier and user. As AI systems become more autonomous, opaquer, and more embedded in critical financial infrastructure, the adequacy of these frameworks is being tested. Ongoing doctrinal improvement, informed by comparative experience and guided by principles of fairness, proportionality, and legal coherence, will be essential to ensuring that tort law remains capable of responding to the risks and realities of AI in financial markets.

Whilst Chapter 4 addressed market manipulation reforms under the FMC Act, this chapter therefore suggests a more coherent foundation for a bespoke product liability regime tailored to AI financial products. Chapter 7 will consider how these gaps might be addressed through a combination of reforms, including the development of such a product liability regime, the introduction of enhanced governance obligations, and the recognition of a novel tort-like cause of action, alongside other complementary reforms designed to ensure accountability in AI-driven financial markets.

Chapter 7: Pathways Forward - Achieving an Effective Regulatory Framework

7.1. OVERVIEW

Chapters 4, 5, and 6 collectively demonstrated that neither existing public-law enforcement mechanisms nor orthodox private-law doctrines can, on their own, deliver a coherent or complete response to the risks posed by AI-driven trading. Chapter 4 showed that while New Zealand's anti-market manipulation regime under the Financial Markets Conduct Act 2013 (FMC Act) is not obsolete, it is increasingly strained by assumptions of human intent, discrete acts, and post hoc enforcement that sit uneasily with the speed, opacity, and autonomy of AI systems. Chapter 5 established that the courts have nevertheless developed effect-based and technology-neutral reasoning capable of attributing responsibility to those who design, deploy, or control automated trading systems, providing a doctrinal foundation upon which further regulatory and liability reform can build.

Chapter 6 then exposed the structural limits of private law. Negligence is significantly constrained where the harm involves pure economic loss or arises in the absence of a proximate relationship or assumption of responsibility. The Consumer Guarantees Act 1993 (CGA) may apply to an AI trading system where the statutory definition of "consumer" is satisfied. The use of such a system to pursue investment returns does not, by itself, exclude the Act; applicability instead turns on whether the goods or services are of a kind ordinarily acquired for personal, domestic, or household use or consumption, whether any of the excluded purposes in s 2(1) applies, and whether the parties have validly contracted out under s 43.[362] The Fair Trading Act 1986 (FTA) addresses misleading or deceptive conduct and false or misleading representations, but it is not a general remedy for design defects, performance failures, or systemic market harm in the absence of qualifying conduct and loss caused by that conduct.[363] Separately, the Financial Markets Conduct Act 2013 (FMC Act) retains a civil compensation route: a court may make a compensatory order where a civil liability provision has been contravened and an aggrieved person has suffered, or is likely to suffer, loss or damage because

[362] Consumer Guarantees Act 1993, s 2(1), s 43(1)–(2A).; *Nesbit v Porter* [2000] 2 NZLR 465 (CA) at 473.
[363] Fair Trading Act 1986, ss 9–13 and 43; *Red Eagle Corporation Ltd v Ellis* [2010] NZSC 20, [2010] 2 NZLR 492 at [26] – [31].

of the contravention.[364] Taken together, these chapters reveal remaining accountability and compensation gaps that may leave some victims of AI-related market harm without an effective remedy under existing frameworks.

Part 8 of the FMC Act provides the statutory benchmark for the reforms proposed in this chapter. It already permits declarations of contravention and compensatory orders for loss caused by specified civil liability contraventions, including the Part 5 market provisions. The case for reform is therefore not that statutory compensation is absent. Rather, Part 8 does not provide a general remedy for non-manipulative design, supervision or product failures, and it may be difficult to invoke where contravention and causation are obscured by opaque, high-speed trading. The reforms proposed in this chapter should accordingly be understood as supplementary to, rather than substitutes for, the existing statutory route.[365]

In response to this gap, Section 7.2 explains why a dedicated AI trading compliance and liability regime is required. It links the doctrinal limits identified in Chapters 4 and 6 to the practical accountability gap that arises in AI-driven markets and draws on successful regulatory models already in place in New Zealand. For example, the Anti-Money Laundering and Countering Financing of Terrorism (AML/CFT) Act 2009 provides a high-level legal mandate, while the Department of Internal Affairs (DIA) issues detailed, practical guidelines that operationalise compliance. Similarly, the relationship between the FMC Act and the NZX Rules illustrates how overarching legislation can be complemented by granular, operational rules. The proposed compliance regime developed in this chapter follows that model, offering practical, AI-specific guidance that builds upon the legislative adaptations proposed in Chapter 4.

Sections 7.3 to 7.8 set out the core elements of the proposed compliance architecture. Section 7.3 introduces the concepts of pre-market approval and conformity assessments, drawing inspiration from international regulatory models such as the United States Food and Drug Administration and proposing a New Zealand-specific agency, the NZX AI Regulatory Office (NZX AI-RO), to oversee high-risk AI trading systems. Section 7.4 focuses on compliance by design, including the use of circuit breakers, speed bumps, and embedded safeguards to prevent manipulative behaviour at the algorithmic level. Section 7.5 addresses

[364] Financial Markets Conduct Act 2013, ss 385(2) and (3)(b), 485 and 494–495.
[365] Financial Markets Conduct Act 2013, ss 385, 486–487 and 494–495.

the importance of training and awareness, proposing certification and continuing education requirements for both compliance officers and technical staff involved in AI trading.

Section 7.6 outlines a governance and human oversight framework that ensures responsibility for AI systems remains with identifiable individuals and is supported by robust documentation, intervention mechanisms, and board-level accountability. Section 7.7 sets out a risk management framework structured around pre-trade, real-time, and post-trade controls, complemented by simulation testing, stress testing, and business continuity planning. Section 7.8 turns to market surveillance and additional measures, including the use of Supervisory Technology (SupTech) and the potential for behaviour-based tools such as Pigovian cancellation taxes to deter harmful AI-driven trading patterns.

Recognising that compliance alone cannot eliminate residual risk, the chapter then moves to the design of complementary liability mechanisms. Section 7.9 introduces a novel cause of action, AI Governance Liability (AGL), specifically designed to address the supervisory and organisational failures that traditional negligence law cannot adequately capture. Grounded in a statutory duty of diligent governance, the AGL framework provides a fault-based mechanism for holding firms accountable for systemic oversight failures, particularly where harm arises from the deployment or mismanagement of autonomous trading systems. Section 7.10 proposes a bespoke statutory product liability regime under the FMC Act to address the limitations of the CGA and FTA in covering AI trading systems used for wealth accumulation. This strict liability framework treats AI financial systems as products and imposes liability for defects regardless of fault, providing a clear and accessible remedy for victims of design-related failures, while drawing on comparative developments such as the European Union's revised Product Liability Directive.

Section 7.11 examines complementary compensation mechanisms, including mandatory liability insurance and the potential development of an industry-funded disruption fund, designed to ensure that victims of AI-related harm can obtain redress even where enforcement or litigation proves difficult, delayed, or evidentially demanding. Section 7.12 concludes by demonstrating how the proposed compliance, liability, and compensation measures interact to form a coherent and proportionate regulatory framework. Each mechanism addresses a distinct legal and policy gap identified in earlier chapters: compliance measures reduce the likelihood and severity of harm; the AGL framework and statutory product liability regime ensure that responsibility attaches to those best placed to manage AI-related risks; and

compensation mechanisms ensure that victims are not left without practical redress. Together, these measures provide a comprehensive, principled, and forward-looking response to the risks and realities of AI in New Zealand's financial markets.

7.2. WHY A DEDICATED AI TRADING COMPLIANCE REGIME?

Existing market abuse regulations, such as the EU's Market Abuse Regulation (MAR) and the Criminal Sanctions Directive for Market Abuse (MAD), prohibit market manipulation, with violations potentially leading to administrative or criminal penalties depending on their severity. However, as discussed in Chapter 3, the speed, opacity and autonomy of AI-driven trading can complicate detection, attribution and proof under these technologically neutral prohibitions, even where the resulting conduct remains within their legal scope.

Similarly, as discussed in Chapter 4, ss 262 and 265 of the FMC Act are civil liability provisions to which the remedies in Part 8 apply. On the application of the FMA or any other person, the High Court may make a declaration of contravention under s 486. Under s 487, that declaration is conclusive evidence of the matters stated in it for the purposes of a later application for a civil liability order. Sections 494 and 495 further permit the FMA or any other person to seek any order the court considers just to compensate a person who has suffered, or is likely to suffer, loss or damage because of the contravention. Separately, market-manipulation contraventions may attract pecuniary penalties and, where the statutory elements are established, criminal liability under ss 264 and 269.[366]

Part 8 must therefore be treated as an existing compensation route and an important benchmark, rather than as a remedial vacuum. Its limitations lie in its scope and evidential requirements. A compensatory order depends on establishing a civil liability contravention and loss caused by that contravention; it is not a general remedy for defective design, negligent supervision or a non-manipulative operational failure. Further, opaque, high-speed and multi-actor trading may complicate proof of the relevant contravention and causal connection. The reforms proposed in this chapter are justified only to the extent that they address those residual cases and should be framed as complementary to Part 8.[367]

---

[366] Financial Markets Conduct Act 2013, ss 262, 265, 385, 486–487, 494–495, 264 and 269.
[367] Yesha Yadav "The Failure of Liability in Modern Markets" (2016) 102 *Virginia Law Review* 1031 at 1073–1096.

Additionally, the rights-based limits of the common law confirm that negligence alone cannot bridge this gap. As we saw in Chapter 6, courts are reluctant to recognise open-ended duties of care that would effectively insure all other market participants against pure economic loss absent a special relationship or undertaking of responsibility. These factors underscore that new tool, beyond the adaptation of existing law, are needed to comprehensively address AI-related market harms.

In May 2013, Germany enacted the High-Frequency Trading Act.[368] This Act explicitly regulates securities trading firms, basing its oversight on their technological infrastructure and the volume and nature of their trading activity.[369] The key elements of this pioneering Act include mandating licensing requirements for proprietary traders, establishing organisational requirements for firms engaged in algorithmic trading, and seeking to increase transparency for all market participants. The German HFT Act did not direct firms to minimise latency. Rather, infrastructure intended to minimise latency formed part of the statutory description of a high-frequency algorithmic trading technique. The Act imposed licensing, organisational, algorithm-identification, order-to-trade ratio and minimum tick-size requirements on relevant firms and trading venues.[370]

Article 3 of the German High-Frequency Trading Act inserted s 33(1a) into the German Securities Trading Act. That provision required trading firms that use algorithmic trading to establish robust systems and controls. These measures are designed to ensure that their automated trading systems are resilient, have capacity, and are subject to appropriate thresholds and limits. Firms must also implement strategies to prevent incorrect orders from being routed, avoid actions that could destabilise the market, and ensure their systems comply with both European and national market abuse regulations, as well as the specific rules of the trading venues to which they connect. It also mandates the creation of robust business continuity plans for system failures, the regular testing and monitoring of trading systems to identify and

[368] Gesetz zum Hochfrequenzhandel (High-Frequency Trading Act) (Germany), art 3 s 33 <https://www.bafin.de/SharedDocs/Veroeffentlichungen/EN/Aufsichtsrecht/Gesetz/hft_en.html>.
[369] Spindler, above n 92, at 210–211.
[370] Bundesanstalt für Finanzdienstleistungsaufsicht "Translation of the Main Provisions of the High Frequency Trading Act" (8 January 2014), arts 1–3, especially Stock Exchange Act ss 16(2), 26a and 26b, and Securities Trading Act s 33(1a).

mitigate risks, and the maintenance of detailed records of any modifications to trading algorithms.[371]

These requirements are now more directly reflected at European Union level in art 17 of MiFID II and Commission Delegated Regulation (EU) 2017/589 (RTS 6). That framework requires resilient systems, appropriate thresholds and limits, prevention of erroneous orders and disorderly trading, testing, monitoring, record-keeping, business continuity, kill functionality, automated surveillance, and pre-trade, real-time and post-trade controls.[372]

New Zealand's financial market, while sophisticated, is considerably smaller and less complex than Germany's, particularly in the realm of HFT and algo-trading. Therefore, directly implementing a broad, prescriptive act like the German HFT Act might be an overreach, potentially creating disproportionate compliance burdens for financial institutions and stifling innovation. A "light-touch, proportionate and risk-based approach" to AI regulation, often preferring to leverage and adapt existing legislation or develop targeted guidelines, would be more favourable to New Zealand rather than enacting entirely new, standalone laws. This approach recognises AI as a general-purpose technology and seeks to avoid overly broad regulations based on speculative harms, which could impede productivity and development.

The EU AI Act also adopts a risk-based framework, but algorithmic trading, market-making and high-frequency trading systems are not generally classified as "high-risk" merely because they operate in financial markets. Under art 6(2), the systems classified as high-risk are those listed in Annex III, which includes specified financial uses such as evaluating the creditworthiness of natural persons and certain risk assessments for life and health insurance, but does not list trading systems. A trading system would fall within art 6(1) only if it were a regulated product, or a safety component of such a product, covered by Annex I and required to undergo third-party conformity assessment. The AI Act's risk-management, documentation, logging, human-oversight, accuracy and cybersecurity requirements should therefore be used

[371] Gesetz zum Hochfrequenzhandel (High-Frequency Trading Act) (Germany), art 3 s 33.
[372] Directive 2014/65/EU [2014] OJ L 173/349, art 17; Commission Delegated Regulation (EU) 2017/589 [2017] OJ L 87/417, arts 5–17.

as comparative policy guidance for New Zealand, rather than described as obligations that ordinarily apply to AI trading systems.[373]

Again, while the EU AI Act provides a robust and forward-thinking framework, its immediate wholesale adoption might not be the most pragmatic approach for New Zealand's financial markets. New Zealand's market is significantly smaller and less complex than the EU's, requiring a more agile and tailored regulatory response. The FMA and NZX possess deep expertise in financial market dynamics and are better positioned to develop nuanced rules that directly address the specific risks of AI in trading, including market manipulation. A broad, overarching AI Act could introduce unnecessary regulatory burden or stifle innovation in a rapidly evolving technological space, especially if its provisions are not perfectly aligned with the unique characteristics of New Zealand's financial ecosystem.

Focusing on immediate and pressing issues, including market manipulation through targeted amendments as proposed in Chapter 4, allows for quicker implementation and adaptation to emerging threats, rather than waiting for a comprehensive, cross-sectoral framework to be fully developed and enacted. However, the core principles underlying the German Act and the EU AI Act remain profoundly relevant and serve as an excellent source of inspiration for New Zealand's regulatory evolution. These principles, namely, transparency, robust risk management, and the unwavering commitment to market integrity, are universal to ensuring fair and orderly markets, regardless of size or technological advancement. For New Zealand, this means figuring out how to use AI's potential while reducing its specific risks.

Drawing on the EU experience, a nuanced approach for New Zealand would involve several key elements forming the core of a robust compliance programme for AI in trading. Firstly, a proactive stance necessitates pre-market approval, conformity assessments, transparency and explainability requirements. However, as will be discussed further in Section 7.3, it is important to recognise that even transparency and explainability measures have inherent limitations when dealing with complex self-learning AI systems. Transparency, in this context, would need to extend beyond surface-level disclosures to include audit trails, comprehensive documentation of model design and objectives, and structured information

[373] Regulation (EU) 2024/1689 [2024] OJ L 1689/1, arts 6(1)–(2), 9–15 and Annex III; Alessio Azzutti "AI Governance in Algorithmic Trading: Some Regulatory Insights from the EU AI Act" [2024] SSRN Electronic Journal at 268–272.

about training data and testing processes. While such measures would not eliminate the inherent opacity of complex or “black-box” AI systems, they could nonetheless provide regulators with a meaningful and practicable level of insight into how AI systems operate, sufficient to assess risk, governance, and compliance prior to deployment.

Secondly, the programme would require compliance by design and speed bumps or circuit breakers. This critical element requires that market conduct rules not be merely external guidelines but, be intrinsically embedded in the AI’s architecture and programming. This involves hard-coding compliance mechanisms and integrating speed bumps or circuit breakers to prevent manipulative actions and ensure proper market behaviour. Thirdly, increasing awareness about the risks associated with AI and providing sufficient training at all levels is paramount. This requires a comprehensive approach to educating all stakeholders on potential hazards.

Fourthly, establishing an effective human oversight, governance, and responsibility framework is essential. This measure ensures that, despite increasing automation, human accountability remains firmly in place, supported by clear lines of responsibility for AI system performance and compliance, as well as robust technical documentation. While various responsibilities should be allocated among senior managers, risk and compliance managers, and trading staff, the board of directors at trading firms should ultimately be responsible for assessing market manipulation risks and formulating a monitoring strategy.

Fifthly, to safely adopt AI in the New Zealand securities market, a risk-based regulatory framework is proposed that mandates a three-part "pre-trade, real-time, and post-trade" risk management system, continuous human oversight, and rigorous testing in isolated simulation environments. This holistic approach, informed by EU regulations like MiFID II and the German HFT Act, aims to ensure market integrity while accommodating the complexity of self-learning AI algorithms. Sixthly, effective market surveillance requires regulators to be equipped with sufficient resources, such as Supervisory Technology (SupTech), to combat modern market manipulation.

Seventhly, to address the inevitable residual risks in an AI-driven market, the programme must establish a vital and comprehensive liability and enforcement framework, which will serve as an indispensable legal safety net. This framework needs to move beyond

mere regulatory fines and sanctions to ensure investors and market participants can achieve substantive compensatory relief against unique modern harms. Finally, the feasibility of employing insurance schemes, market-based funds, and disruption funds will be considered to ensure sufficient victim compensation. The subsequent sections will detail the implementation plan for these measures to deliver a robust yet balanced framework for responsible AI innovation.

7.3. PRE-MARKET APPROVAL: CONFORMITY ASSESSMENTS AND COMPLIANCE AUDITS

Tutt highlights the crucial need to deter harms caused by increasingly complex algorithms. It is suggested that, instead of relying on a reactive criminal or tort regulatory regime, a proactive pre-market review is a more effective way of ensuring that unacceptably dangerous algorithms are not introduced into the market without unduly stifling innovation. In the context of trading, a fully AI-enabled trading algorithm is likely to fall under such categorisation. A centralised government agency, similar to the United States Food and Drug Administration ("FDA"), is proposed.[374]

Pre-market Approval is the most rigorous process the FDA uses to ensure that high-risk medical devices, like those that sustain life or pose significant health risks, are safe and effective before they can be sold. Unlike less risky devices, these "Class III" devices undergo extensive review and require robust scientific evidence from both laboratory and human studies. The company developing the device must submit a comprehensive application detailing everything from its design and purpose to its performance and potential risks, including how it was tested. This strict review process ensures that only devices with proven safety and effectiveness reach the market, and that ongoing monitoring continues even after approval.

Tutt's proposed agency, referred to as "FDA for Algorithms", would evaluate complex and potentially dangerous algorithms similarly to how the FDA evaluates high-risk medical devices. This rigorous evaluation would aim to determine whether algorithms are safe and effective for their intended use, and to identify satisfactory measures to prevent harmful

[374] Tutt, above n 15, at 92–96.

misuse.[375] The analogy is useful at the level of ex ante screening, but it should not be applied mechanically. Medical devices normally have a comparatively stable intended use, whereas trading algorithms interact continuously with changing market conditions and may be materially updated after deployment. A point-in-time approval must therefore be supplemented by material-change controls, continued monitoring, periodic validation and a power to suspend or withdraw approval.[376]

For New Zealand's exchange market, a similar agency, to be referred to as the "NZX AI Regulatory Office" ("NZX AI-RO"), is proposed. It is suggested that the NZX AI-RO be established as a specialised unit or division within NZX RegCo, maintaining strong reporting lines and collaborative mechanisms with the FMA. Such an agency should operate as a proportionate and risk-based regulator, providing clear guidance and practical engagement with applicants while retaining binding powers in relation to designated high-risk systems. If pre-deployment approval is intended to be compulsory, the source of that power and the consequences of non-compliance should be stated in legislation or enforceable market rules rather than left to non-binding guidance.[377]

Furthermore, such an agency, by consolidating specialised technical and regulatory expertise under one roof, would be uniquely positioned to strike the proper equilibrium between fostering innovation and ensuring robust regulation. Its comprehensive understanding of both the potential benefits and inherent risks of AI would enable it to develop proportionate and effective rules that encourage technological advancement while simultaneously safeguarding market integrity. A significant benefit of this centralised approach is its inherent capacity to promote collaboration among diverse stakeholders. This would include industry regulators, the scientific and software development communities, customers, victims, and commercial firms. The agency would act as a central forum for dialogue and knowledge exchange, facilitating the co-development of best practices and ensuring that varied perspectives and expertise inform regulatory frameworks. This multi-stakeholder engagement would be particularly valuable during the pre-market review process, ensuring that all concerns, from technical feasibility to consumer protection and market stability, are addressed comprehensively before potentially harmful algorithms are deployed.

---

[375] 21 USC § 360e (premarket approval of medical devices).
[376] Tutt, above n 15; Commission Delegated Regulation (EU) 2017/589, arts 9, 11 and 16–17.
[377] See, by analogy, Financial Markets Conduct Act 2013, pt 5 subpt 7, especially ss 327–333.

The agency could also pursue algorithmic safety by imposing transparency requirements, including openness and disclosure. Traditional algorithms are relatively straightforward to understand by reviewing their code, but advanced AI-driven trading algorithms have an inherent "black-box" nature that can make their decision-making opaque even to their creators. There is growing agreement that one of the foundational tools for regulating AI algorithmic safety is mandatory transparency and explainability. At the same time, it must be acknowledged that these tools have clear limits in the context of complex, self-learning systems.[378]

Adopting explainability in the context of AI trading systems means implementing measures and techniques that enable humans to clearly understand the AI's decision-making process, including how complex machine learning algorithms arrive at their conclusions and predictions. There are two primary approaches to achieving explainable AI. The first is an exogenous approach that aims to explain how the algorithmic model operates under specific circumstances. This includes detailing its underlying principles and assumptions, providing an overview of the developer's or programmer's intentions, the dataset and parameters used to train the model, and the testing methods employed to evaluate the model's performance and identify potential biases or errors.

The second approach focuses on replicating the decision-making process of a complex AI algorithm itself. While disclosing the underlying code can be helpful, it may not be sufficient to grasp the intricacies of advanced AI trading models fully. An alternative is to create a simpler, more interpretable model, referred to as a "surrogate model," that approximates the behaviour of the original complex model. The surrogate model can then provide valuable insights such as the key factors affecting the algorithm's decision-making process, the decision boundaries, and how changes in input features affect the algorithm's output.[379]

Under the compliance regime, the proposal is to adopt explainability for all AI-driven trading systems used in New Zealand's financial markets. This does not necessarily mean exclusively allowing AI models that are inherently transparent from inception. Instead, it

[378] Tutt, above n 15, at 109–110; Fletcher, above n 167, at 313.
[379] Fletcher, above n 167, at 314–315.

requires that even complex AI models, including those with “black-box” characteristics, must be accompanied by robust explainability features. The proposed centralised agency would be responsible for establishing the specific standards and techniques for achieving this required level of explainability, ensuring that all AI decisions can be adequately understood by human oversight, regulators, and, for audit purposes, thereby making their operations sufficiently comprehensible for effective governance and accountability within the financial market.

Arguments suggest that mandating certain levels of AI explainability risks compelling developers to disclose proprietary or sensitive information, potentially leading to its misuse. However, these challenges do not render the proposal to mandate explainability unviable; instead, they highlight the critical role of the proposed centralised agency in carefully navigating the inherent trade-offs. The agency's task would be to define a proportionate level of explainability that maximises regulatory insight and public trust without unduly compromising performance or revealing sensitive intellectual property.

This would involve establishing nuanced standards, where not all aspects of an AI need to be fully transparent, but specific, actionable explanations for key decisions or patterns would be required. Furthermore, the agency would be responsible for developing secure disclosure protocols to protect proprietary information while allowing necessary regulatory access, and for mandating stringent cybersecurity measures to ensure any required explainability insights are protected from unauthorised access or malicious misuse. This systematic approach, leveraging a dedicated regulatory body, is vital to effectively implement explainability, ensuring its benefits are realised while actively mitigating its potential drawbacks in the financial market context.

Industry standards can establish benchmarks for generating AI explanations by defining the parameters for developing explanations, ensuring accuracy, and determining when explanations should be generated to ensure efficient testing. This would mean that, in cases of suspected manipulation, the focus would shift from understanding the algorithm's specific decision-making process to assessing whether the developer adhered to industry standards for testing and oversight as part of the market review process. Industry standards can significantly influence industry practices. These standards can encourage responsible AI development and deployment by setting expectations and providing guidance.

The agency would also be responsible for ensuring that transparency requirements are satisfied. Transparency, in the context of advanced AI trading algorithms, can indeed manifest in diverse forms, from a lighter touch to the most stringent, with significant implications for the depth, scope, and timing of disclosures. A lighter touch is generally anticipated where the imperative to preserve a firm's intellectual property, specifically the proprietary nature of their AI-enabled or machine-learning algorithms, codes, and training data, legitimately outweighs the need for complete public transparency. To address this, a sensible approach would require specific components of these high-risk algorithms to be certified by independent, accredited third-party organisations. This method safeguards the proprietary nature of the algorithms while still providing essential assurance. For such scenarios, the depth of disclosure to the regulator and certified third parties would not necessitate revealing the intricate, secret sauce of how algorithms precisely function or their complete source code.

Instead, firms would be required to provide meaningful information about the algorithm's high-level functionality, its demonstrated effectiveness through rigorous testing, and its potential limitations, including identified biases or vulnerabilities. Disclosure would typically be limited to a small group of highly credentialed, external experts operating under strict non-disclosure agreements. It could be subject to carefully managed delays to protect commercially sensitive information, perhaps for a period consistent with industry norms for safeguarding trade secrets post-launch, rather than for "years or decades", which would indeed be an exaggeration and render the oversight ineffective. This nuanced approach acknowledges that the goal is not to reverse-engineer proprietary systems, but to gain sufficient assurance of safety, fairness, and market integrity. [380]

On the more stringent end, where the need to provide transparency outweighs the need to preserve secrecy, the agency could mandate that all information about an algorithm, including the technical details, be disclosed for public safety. For high-risk algorithms, the depth of disclosure is expected to be thorough, with the information made generally available to the public and the timing of disclosure immediate. Without examining the specific pros and cons of any approach, it is essential to emphasise that striking a balance between innovation and safety will require careful, extensive consideration. The agency could navigate this delicate

[380] Tutt, above n 15, at 110.

balance by engaging various stakeholders and mandating only disclosures most relevant to specific algorithms.[381]

While this approach might appear "vague" at first glance, it is intentionally flexible, recognising the rapid evolution of AI technology and the need to avoid overly prescriptive rules that could stifle innovation or become quickly outdated. The intention is not to abdicate responsibility for defining the specifics, but rather to establish a robust framework within which those specifics can be collaboratively developed and adapted. This is a common regulatory model in which principles are set, and detailed implementation is left to industry best practices, technical standards, and certified third-party evaluations, under the regulator's ultimate oversight. The regulator's role, therefore, shifts from micro-managing code to overseeing the robustness of the certification process and having the ultimate authority to greenlight or halt deployment based on certified findings. This ensures that while proprietary information remains protected, the core objectives of market integrity, transparency in outcomes, and risk mitigation are comprehensively addressed through an auditable and accountable process.

The suggested transparency and explainability requirements are indeed crucial and are highly endorsed by leading global regulatory principles for AI, particularly in high-stakes environments such as financial markets. They align closely with international best practices emerging from frameworks such as the EU AI Act. Therefore, it is strongly recommended that such transparency and explainability requirements form a fundamental part of the proposed pre-market approval process within New Zealand's exchange market. In its frontline role, the proposed NZX AI-RO would conduct pre-deployment conformity assessments for high-risk AI trading algorithms. This would entail leveraging its technical expertise in AI and trading systems to review submissions from trading firms thoroughly.

At the international level, ISO has developed safety standards for industrial robots, including standards addressing the design of industrial robots and their integration into industrial applications and robot cells.[382] Although those standards do not apply to trading algorithms, they provide a useful analogy for the use of technical standards to manage risks

[381] Tutt, above n 15, at 110.
[382] ISO 10218-1:2025 *Robotics — Safety requirements — Part 1: Industrial robots* (International Organization for Standardization, 3rd ed, 2025); ISO 10218-2:2025 *Robotics — Safety requirements — Part 2: Industrial robot applications and robot cells* (International Organization for Standardization, 2nd ed, 2025).

arising from complex automated systems. By analogy, the Agency could issue guidance on the design, testing and performance of AI-driven trading algorithms, with the aim of building in adequate safety margins. Consistently with Tutt's proposal, that guidance could take account of an algorithm's intended use, the distinction between critical and acceptable errors, and the legal standard likely to apply if its operation causes harm. As an illustration, Tutt proposes a possible performance standard under which an autonomous stock-trading algorithm's average return volatility would be predicted with 95 per cent confidence from historical data and disclosed to investors.[383]

The pre-market approval process would allow firms to assess whether their trading algorithms meet applicable market conduct and performance standards and, where necessary, to modify or redesign them. The agency could work with an applicant to develop appropriate studies or testing protocols through which the applicant could demonstrate, to the agency's satisfaction, that its algorithm meets those standards. Where unrestricted approval is not justified, the agency could grant conditional approval subject to clearly defined limits on the algorithm's permitted use, with use outside the approved scope or deployment without approval potentially attracting legal sanctions.[384] In the proposed New Zealand regime, conditional approval could also require enhanced monitoring and the completion of specified remediation milestones, while the enabling framework should authorise the suspension or withdrawal of approval if those conditions are breached.

Moreover, as part of the authorisation process, a trading algorithm must be certified to demonstrate "good conduct" by design, meaning it is resilient against known forms of abuse. For instance, the algorithm is both not programmed to engage in manipulative practices and is difficult for bad actors to exploit for manipulation. Fletcher suggests that this certification should involve obtaining personal attestations from the algorithm's designers and users, similar to the certifications required from a CEO or CFO for a company's financial reports. These attestations would verify that the individuals have personally reviewed their algorithms, confirming that the code is not designed to breach applicable laws and regulations. Such certification requires rigorous testing in simulated environments under various market conditions to assess the algorithm's potential vulnerabilities, measure its robustness, and gauge

[383] Tutt, above n 15, at 108.
[384] Tutt, above n 15, at 111.

its impact on market stability. The effectiveness of this pre-approval regime relies on the ability to detect whether an AI system could autonomously develop or execute strategies that exploit market rules. By subjecting AI systems to this in-depth testing, potential risks of market manipulation can be identified and mitigated before they enter the live market.[385]

Beyond the initial pre-market approval, the proposed NZX AI-RO would implement a robust system of continuous oversight and regular compliance audits to maintain a dynamic and adaptive regulatory environment for AI in financial markets. Its surveillance teams would continuously monitor live AI trading activity on the NZX, ensuring ongoing adherence to approved parameters and actively detecting any anomalous or manipulative behaviour stemming from AI systems. This real-time monitoring would involve leveraging advanced analytical tools to spot deviations in trading patterns, unexpected changes to the order book, or unusual message traffic that could indicate a system malfunction, disorderly trading or possible manipulative conduct. The aim is to identify and alert to potential issues as they occur, allowing for rapid intervention before they escalate into significant market disruptions.

Complementing this continuous oversight, regular compliance audits and inspections of firms' use of AI trading systems would be a cornerstone of the post-deployment regulatory framework. These audits would not be mere paper exercises; they would be comprehensive, risk-based assessments designed to verify that AI systems are operating as intended, that internal controls remain effective, and that firms continuously manage the risks associated with their AI models. The NZX AI-RO's audit procedures would likely involve several key components. Firstly, they would entail a data audit, meticulously examining the quality, integrity, and ongoing relevance of the data used by the AI systems, including checks for accuracy, completeness, and potential biases that might emerge over time.

Secondly, an algorithm review would assess the ongoing functioning of the AI algorithms, ensuring they remain free from unintended errors or drift from their approved specifications, and that the governance structure around their continuous development and maintenance is robust. Thirdly, outcome evaluation would assess the AI-generated results against expected performance metrics, identifying any deviations or anomalies that could

[385] Azzutti, Ringe and Stiehl, above n 12, at 130–131.

indicate issues with the model or its environment. Finally, these audits would rigorously verify compliance with all relevant regulations and the specific pre-approval conditions, critically assessing the effectiveness of human oversight mechanisms and internal risk management frameworks. This multi-faceted approach, combining real-time surveillance with periodic, in-depth audits, aims to ensure the long-term safety, reliability, and integrity of AI-driven trading on the NZX.

7.4. COMPLIANCE BY DESIGN AND SPEED BUMPS OR CIRCUIT BREAKERS

Azzutti noted that, in addition to adapting existing anti-market manipulation provisions, implementing rules that mandate that trading algorithms "self-regulate" can serve as an effective deterrent. This "self-regulation" refers to the AI's internal design, preventing it from engaging in manipulative practices. It constitutes "deterrence" ьecause the architecture of the AI directs its actions away from illicit behaviour by embedding specific market conduct rules directly into its programming code and algorithmic design. For example, programming can ensure "good behaviour" by either hard-coding compliance mechanisms that make manipulative actions computationally impossible for the AI to execute, or by structuring the AI's internal reward and loss functions such that engaging in illicit practices results in a direct computational penalty, making non-compliant behaviour inherently suboptimal for the AI's performance objectives. This internal "cost" serves as a built-in "deterrent."[386]

The principle that AI algorithms should be designed to support traceability and market stability, as reflected in European and German regulatory practice, is relevant to reducing the risk of AI-enabled market manipulation in New Zealand. New Zealand's "light-touch" AI strategy, which prioritises responsible AI adoption and the leveraging of existing regulatory frameworks over creating entirely new legislation, means these well-established principles are not just relevant but also highly adaptable for practical implementation. This approach allows New Zealand to proactively address potential manipulative practices directly at their algorithmic source, before they can manifest in market behaviour. It offers a more robust and immediate form of protection than traditional, after-the-fact legal sanctions.[387]

---

[386] Azzutti, Ringe and Stiehl, above n 12, at 130.
[387] Spindler, above n 92, at 212; Commission Delegated Regulation (EU) 2017/589 of 19 July 2016 supplementing Directive 2014/65/EU [2017] OJ L 87/417, art 8.

For AI trading in New Zealand, this integrated approach is proposed to be adopted as a fundamental component of both the pre-market approval process and ongoing post-deployment oversight, under the purview of the proposed NZX AI-RO. Compliance by design would be a mandatory tenet, meaning that AI trading algorithms would be inherently constructed with regulatory requirements and risk mitigation at their core from conception. This includes mandating full auditability and traceability for every AI-driven order and transaction. Each AI-generated order should be recorded using identifiers and records sufficient to identify the responsible algorithm, its version and the relevant deployment settings. This granular traceability is essential for the NZX AI-RO's surveillance teams to swiftly identify any "rogue" algorithm or pinpoint the source of poor performance, facilitating rapid intervention before minor issues escalate into significant market disruptions.

Furthermore, AI algorithms would be designed to integrate proactive guardrails to reduce the risk of market disruption and support liquidity provision on a regular and predictable basis. Drawing inspiration from European requirements, these would encompass defined pre-trade limits embedded within, or applied to, the AI system's operational framework. These controls would constrain the AI system's ability to exceed predefined thresholds for trading characteristics such as the price, volume or value of orders, message volumes, strategy positions and market or credit exposures. Where an AI system pursues a market-making strategy, its quoting obligations and the exceptional circumstances in which those obligations may be reduced or suspended should be defined in advance.

An AI system used for directional trading should be subject to proportionate position, exposure, price and message limits designed to reduce the risk of undue market impact. These firm-level pre-trade controls should operate alongside, but not be confused with, trading-venue circuit breakers. Firm-level controls constrain orders and exposures, whereas venue-level circuit breakers halt or constrain trading when specified market conditions arise.[388] It is therefore strongly suggested that such 'compliance by design' principles be adopted as a key requirement for AI trading algorithms deployed in New Zealand. As part of pre-market approval and ongoing supervision, it is suggested that the NZX AI-RO verify that such embedded controls are in place, rigorously review them, and monitor their continuing

[388] Directive 2014/65/EU, arts 17(3)–(4) and 48(5); Commission Delegated Regulation (EU) 2017/578, arts 2–4; Commission Delegated Regulation (EU) 2017/589, arts 8, 15 and 16.

effectiveness through advanced real-time surveillance, thereby supporting the safety, reliability and integrity of AI-driven trading within New Zealand's financial markets.

Trading venues that allow or enable algorithmic trading should be required to maintain resilient systems with sufficient capacity, effective controls to reject orders that exceed predetermined volume and price thresholds or are clearly erroneous, and arrangements to halt or constrain trading where significant price movements occur. They should also maintain effective business continuity and physical and electronic security arrangements. These venue-level obligations are distinct from the pre-trade; real-time and post-trade controls imposed on investment firms. Those provisions do not, however, impose a general obligation on trading venues to provide algorithmic traders with "seamless access" or unspecified trading tools.[389]

Beyond these foundational requirements, a critical consideration, particularly for the ultra-fast nature of AI trading, is the implementation of "speed bumps" or other controlled trading delays on exchange platforms. As suggested by Scopino and demonstrated by the CFTC's 2019 approval of Intercontinental Exchange (ICE) to slow trading in gold and silver futures, these measures aim to increase the time between order placement and execution. Although the growing use of speed bumps supports considering them as a regulatory option, the evidence does not establish that they are invariably beneficial. Any NZX model should be tested against its likely effects on liquidity, spreads, price discovery and migration to other venues, and should be targeted to identified risks rather than treated as a universal control.[390]

Drawing on these established international regulatory approaches, New Zealand is well-positioned to adapt and implement similar measures for its own financial markets. Given the FMA's ongoing engagement with AI in the financial sector, it is strongly proposed that the NZX, as a critical trading venue, adopts these robust operational and design requirements. Specifically, the proposed NZX AI-RO should mandate stringent requirements for trading

---

[389] Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments [2014] OJ L 173/349, art 48(1), (4) and (5); Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector [2022] OJ L 333/1, arts 5–12; Commission Delegated Regulation (EU) 2017/584 supplementing Directive 2014/65/EU with regard to regulatory technical standards specifying organisational requirements of trading venues [2017] OJ L 87/350, arts 11, 15–20 and 23; Commission Delegated Regulation (EU) 2017/589 supplementing Directive 2014/65/EU with regard to regulatory technical standards specifying the organisational requirements of investment firms engaged in algorithmic trading [2017] OJ L 87/417, arts 15–17.

[390] Scopino, above n 183, at 388–389.

venues to ensure their operational resilience against system failures and cyberattacks, implement dynamic circuit breakers to maintain market stability during volatility, and provide advanced simulation environments for firms to test AI strategies. Crucially, the NZX should also consider incorporating "speed bumps" or controlled trading delays on its platforms, a measure particularly relevant for mitigating the potential disruptive impact and ensuring a level playing field against the ultra-fast execution capabilities of AI trading algorithms.

7.5. TRAINING AND AWARENESS

As demonstrated in Chapters 4 and 5, the effective application of the market manipulation provisions depends not only on the statutory framework but also on the ability of firms and regulators to identify, explain and attribute potentially manipulative trading behaviour. These challenges become more significant where trading decisions are generated or influenced by complex AI systems. The training and awareness measures proposed in this section respond to these evidential and supervisory difficulties by ensuring that those responsible for developing, operating and overseeing AI trading systems possess the knowledge necessary to identify and manage relevant risks.

The effective oversight of AI trading algorithms hinges significantly on ensuring that the human element, particularly the compliance and technical personnel, possesses the requisite knowledge and skills. This principle resonates with Article 2 of the Commission Delegated Regulation (EU) 2017/589, which mandates that investment firms employ compliance personnel with a general understanding of how algorithms function, fostering continuous communication with technical staff. While compliance officers are not expected to be AI specialists, they must have sufficient knowledge to effectively oversee algorithmic trading, identify red flags, and ensure compliance with regulatory frameworks.[391]

Conversely, the technical staff responsible for developing, operating, and managing these algorithms must demonstrate a detailed, in-depth knowledge of their functionality, the specific risks they pose, and their potential for market impact. This comprehensive understanding is crucial for them to effectively manage, monitor, and rigorously test algorithms

[391] Commission Delegated Regulation (EU) 2017/589, above n 387, art 2.

to ensure proper functioning, compliance with regulations, and the prevention of unintended consequences or manipulative behaviours.[392]

Building on this, Fletcher contends that the escalating prevalence and impact of algorithmic trading necessitate more stringent oversight, particularly on the human actors. She advocates for a comprehensive registration framework for programmers and traders involved with these sophisticated systems. Within this framework, a pivotal element would be the requirement for programmers not only to comply with specific regulations, such as embedding safeguards against manipulative algorithmic behaviour, but also to demonstrate their expertise through rigorous qualifying exams. These examinations would serve to ascertain that programmers possess sufficient, in-depth knowledge of complex market dynamics, the relevant regulatory landscape, and, crucially, the inherent risks associated with both algorithmic trading and advanced AI. By upholding such stringent proficiency standards for programmers, Fletcher contends that the likelihood of algorithmic errors and various forms of misconduct would be significantly reduced, thereby bolstering market integrity.[393]

To implement such a robust training and awareness framework for AI trading in New Zealand, drawing on these insights, a multi-tiered approach is suggested. Firstly, for compliance personnel, dedicated training modules should be developed by the proposed NZX AI Regulatory Office (NZX AI-RO) or an accredited third party. These modules would cover foundational AI concepts, common AI trading strategies, the types of data AI models consume, typical failure modes and the regulatory obligations surrounding AI use, with a particular focus on explainability and transparency requirements. This training could be delivered through online platforms like NZX's existing issuer training modules, incorporating quizzes and scenarios to test understanding, and regular refresher courses would be mandatory to keep pace with evolving technology. The goal is not to make compliance officers coders, but to equip them with the "AI lens" necessary to effectively challenge, question, and oversee AI-driven operations and their governance.

Secondly, for technical staff (AI developers, data scientists, quantitative analysts, etc.), the requirements would be more stringent. A certification or licensing scheme, similar to

[392] Commission Delegated Regulation (EU) 2017/589, above n 387, art 2
[393] Fletcher, above n 167, at 322.

Fletcher's suggestion,[394] could be established by the proposed NZX AI-RO or in partnership with relevant academic and industry bodies. This would involve mandatory advanced training covering topics such as algorithmic finance, machine-learning ethics, secure coding practices for financial algorithms, rigorous testing methodologies for high-frequency AI, the design of embedded compliance controls within algorithms, and post-trade analysis of AI-driven anomalies.

Certification would be contingent on passing rigorous examinations (potentially including practical coding challenges or case-study analyses related to preventing manipulation and maintaining liquidity), and regular re-certification or continuing professional development would ensure technical staff remain current with the latest advancements. This fosters a culture of continuous learning and responsible innovation in New Zealand's AI trading community. Importantly, developers and operators of AI trading systems in New Zealand will need such expertise regardless of whether the underlying AI technology is created domestically or overseas; this requirement is aimed at ensuring that local professionals can effectively supervise and manage even foreign-developed AI platforms.

It is proposed that developers of AI trading algorithms should adhere to standards and principles similar to those followed by accountants under Generally Accepted Accounting Principles ("GAAP"), which ensure consistency and transparency in financial reporting. To achieve the safety and reliability of their AI systems, this would involve adopting industry best practices for secure coding, implementing rigorous testing and validation protocols, adhering to ethical AI guidelines, and meticulously documenting every stage of the algorithm's lifecycle—from data selection and model training to deployment and ongoing monitoring. Just as GAAP provides a common framework for financial professionals, a set of evolving standards and internal policies would guide AI developers in creating algorithms that are robust, transparent, and operate within acceptable risk parameters.[395]

---

[394] Fletcher, above n 167, at 322.
[395] Caitlin V Slemmer "Crashed by Algorithms: Trader Liability in a New Era of Securities Regulation" (2017) 12 Brooklyn Journal of Corporate, Financial & Commercial Law 169 at 170–171.

7.6. HUMAN OVERSIGHT, GOVERNANCE AND RESPONSIBILITY FRAMEWORK

The analysis in Chapters 4 and 5 also demonstrated the importance of reliable evidence concerning trading conduct, including order patterns, the circumstances surrounding trading decisions and the attribution of responsibility for potentially manipulative activity. These evidential challenges are heightened where AI systems operate autonomously or opaquely. The documentation, model-versioning and algorithm-inventory requirements proposed in this section are intended to improve traceability and make relevant evidence more readily available to firms, regulators and, where necessary, courts.

Technical documentation is crucial for responsible AI development and deployment in securities trading, serving as a comprehensive record of an AI's design, functionality, limitations, and testing processes. This detailed record facilitates clear communication among stakeholders, demonstrates regulatory compliance, supports audits, and helps identify and mitigate risks, ultimately guiding users toward informed and responsible algorithm use. Drawing on international best practices, such as Annex IV of the EU AI Act, New Zealand should adopt a similarly thorough approach, requiring documentation of the AI system's intended purpose, functionality, components, user interface, instructions, design specifications, risk controls, and systems for performance evaluation and post-market monitoring.[396] To enhance transparency and accountability, New Zealand's framework, perhaps through the proposed NZX AI-RO, should also mandate detailed algorithm inventories documenting algorithm types, trading strategies, owners, and internal approvers, thereby serving as a cornerstone for effective audits and the integrity of AI trading in New Zealand's financial markets.[397]

To ensure robust AI governance in New Zealand's securities market, a tiered framework for technical documentation is proposed. For in-house developed AI trading algorithms, firms would maintain detailed internal documentation covering design principles, data sets, validation, risk controls, and ongoing monitoring, subject to NZX AI-RO inspection. For AI systems procured from third-party providers, the NZX AI-RO would mandate stringent contractual requirements for comprehensive technical documentation, encompassing system

[396] Azzutti, above n 75, at 24.
[397] Azzutti, above n 75, at 24.

functionality, intended purpose, design specifications, embedded risk controls, performance evaluation, and post-market monitoring. This enhanced disclosure from third parties is vital for preventing risks, including market manipulation, by ensuring transparency regarding the system's design and operational parameters.[398]

To further enhance accountability and traceability and inspired by Section 80(3) of the German Securities Trading Act, the proposed NZX AI-RO could also require detailed algorithm inventories from all firms, meticulously documenting algorithm types, trading strategies, internal owners, and approvers. This comprehensive approach, incorporating both technical specifications and transparent accountability, would form a cornerstone for effective audits and the overall integrity and stability of AI trading in New Zealand's financial markets.[399]

New Zealand should implement a robust framework for human oversight of AI in its securities markets, drawing on principles from the EU AI Act and MiFID II. This oversight must be substantial, designed to mitigate risks associated with both the use and misuse of AI, with its intensity scaled according to the AI application's complexity, nature, and inherent risks. Control systems should provide human experts with tools to monitor AI operations, interpret outputs, and intervene when necessary. Reflecting the "human-in-the-loop" approach foundational to MiFID II, investment firms should integrate human experts into AI decision-making processes, including automated surveillance to detect market manipulation, with human oversight validating and responding to anomalies.[400]

Human oversight and intervention are paramount for monitoring trading algorithms, with humans retaining ultimate control over pre-trade checks. Firms must have procedures for manual review and approval of blocked orders, ensuring human traders can verify legitimacy and authorise submission under specific circumstances. Likewise, investment firms engaged in algorithmic trading should maintain functionality capable of cancelling unexecuted orders immediately where necessary. Under RTS 6, this "kill functionality" forms part of the organisational requirements applicable to algorithmic trading systems.[401] A comparable requirement should form part of the proposed New Zealand framework, ensuring that firms can

[398] Azzutti, above n 75, at 24–25.
[399] Spindler, above n 92, at 212; Wertpapierhandelsgesetz (Germany), § 80(3).
[400] MiFID II, above n 49, art 17(7)(a) and (d).
[401] Commission Delegated Regulation (EU) 2017/589, art 12.

intervene promptly where an AI trading system malfunctions, behaves unexpectedly or creates a risk of disorderly trading.

Furthermore, human overseers, including risk or compliance managers, are expected to act directly and independently of their trading teams to stop or adjust a trade when algorithms malfunction or their behaviour violates good trading practices. These oversight systems are designed to identify suspicious trading patterns, generate alerts and reports for human notification, offer clear insights into detected irregularities, facilitate in-depth investigations into suspicious activity, and enable the implementation of appropriate corrective actions.[402] Additionally, developers would be required to periodically log and archive versions of their AI models to allow for retrospective analysis of the AI model's behaviour and decision-making processes. Notably, in legal proceedings, this could enable access to earlier model versions and obtain further explanations that the parties could have withheld.[403]

The board of directors at trading firms should ultimately bear responsibility for assessing market manipulation risks and formulating an appropriate monitoring strategy. Firms should implement effective real-time monitoring systems, with personnel responsible for monitoring sufficiently independent from those responsible for the trading activity to ensure effective oversight. Firms should also maintain appropriate business continuity arrangements to address system failures, outages and other disruptions. These requirements draw on the organisational framework established by MiFID II and the more detailed requirements contained in RTS 6 concerning real-time monitoring and business continuity.[404]

Direct Electronic Access ("DEA") permits clients to transmit orders electronically to a trading venue using the trading code of an investment firm. Because this arrangement may expose both the provider and the wider market to risks arising from client trading systems, the DEA provider must maintain appropriate controls over orders submitted through its access arrangements. RTS 6 contains detailed organisational requirements governing DEA, including requirements concerning prospective client assessment, controls, monitoring and the provider's responsibility for trading conducted through its systems.[405] A comparable approach should be

[402] Azzutti, above n 75, at 26.
[403] Slemmer, above n 395, at 170–171.
[404] Directive 2014/65/EU (MiFID II), art 17; Commission Delegated Regulation (EU) 2017/589, arts 13–14 and 16.
[405] Lee and Schu, above n 76, at 11–12; Commission Delegated Regulation (EU) 2017/589, arts 19–23.

adopted in New Zealand so that firms providing direct market access remain responsible for maintaining effective risk controls capable of identifying and mitigating risks arising from malfunctioning or inadequately controlled client algorithms.

7.7. RISK MANAGEMENT FRAMEWORK

The analysis in Chapters 4 and 5 identified an important limitation of relying principally on enforcement after potentially manipulative trading has occurred. Although the FMC Act prohibits specified manipulative market effects, the application of those provisions necessarily occurs after or in response to relevant trading conduct. The preventive controls proposed in this section therefore complement the existing statutory regime by seeking to reduce the likelihood that AI trading systems generate the harmful or misleading market effects examined in Chapters 4 and 5.

A risk-based regulatory approach is proposed to promote the trustworthy adoption of AI within the securities market, specifically designed to address potential market failures, including market manipulation. This strategy aligns directly with the EU AI Act's underlying rationale, which prioritises a risk-tiering system to foster responsible innovation while mitigating the significant harms posed by AI.[406] The German HFT framework similarly emphasises the need for adequately resourced risk-management systems for algorithmic trading. At the EU level, Article 17 of MiFID II establishes the overarching organisational and risk-control obligations applicable to algorithmic trading, while RTS 6 provides more detailed requirements concerning pre-trade controls, real-time monitoring and post-trade controls.[407] Together, these measures provide a useful comparative model for structuring risk management across the lifecycle of an AI trading system.

This framework requires firms to subject algorithmic trading systems to appropriate testing before deployment and to maintain effective controls once those systems are operating in live markets. RTS 6 distinguishes between pre-trade controls, real-time monitoring and post-trade controls, thereby providing a structured framework for managing risks throughout the trading lifecycle.[408]

---

[406] EU AI Act, above n 373.
[407] Directive 2014/65/EU (MiFID II), art 17; Commission Delegated Regulation (EU) 2017/589, arts 15–17.
[408] Commission Delegated Regulation (EU) 2017/589, arts 5–7 and 15–17.

AI-based trading systems that can retrain, adapting or otherwise changing their operation create challenges for existing testing requirements. Rather than generally prohibiting adaptive or self-learning algorithms, the EU framework requires firms to assess whether changes to an algorithm or trading system are material and therefore require further testing. Recent ESMA supervisory guidance confirms that material changes or substantial updates to algorithmic trading systems should trigger appropriate testing and identifies changes to machine-learning components or adaptive capabilities as potentially relevant to that assessment.[409] This approach reflects the broader regulatory challenge of ensuring that systems which evolve after initial deployment remain subject to effective testing, monitoring and governance.

Regulatory guidance also emphasises that increasing automation does not remove the need for effective human oversight and robust testing. Firms should maintain appropriate controls over algorithmic outputs and test material changes in environments that minimise the risk of unintended effects on live markets.[410] These requirements are particularly important for AI-enabled systems whose behaviour may change following retraining, modification or changes in underlying data.

Building on this, it is understood that, before deploying any trading strategies, trading firms must conduct thorough performance testing. This testing is essential to ensure that their algorithms fully comply with the specific requirements of the trading venue or direct market access provider, including perfect integration with the venue's matching engine and accurate processing of data flows. Furthermore, if a trading venue undertakes significant changes to its systems or infrastructure, it is crucial to recognise that this may impact how algorithmic trading systems interact with that venue, necessitating re-testing. Within this framework, market participants should be granted the autonomy to choose their preferred testing environments and to self-certify the performance of their algorithms, while still being subject to regulatory oversight and validation.[411]

[409] European Securities and Markets Authority Supervisory Briefing on Algorithmic Trading in the EU (ESMA74-1505669079-10311, 26 February 2026) at [30] – [31].
[410] Lee and Schu, above n 76, at 8.
[411] Lee and Schu, above n 76, at 8–9.

Critically, the operators of these venues must collaborate closely with investment firms, ensuring that their algorithmic trading practices, especially those driven by AI, fully comply with established market conduct rules. Additionally, this oversight extends to requiring venues to provide comprehensive simulation environments for investment firms. These environments would enable firms to rigorously test their trading algorithms, particularly complex AI-enabled strategies, in a controlled setting that accurately simulates real-market conditions. Such pre-deployment testing is invaluable for identifying potential issues, fine-tuning algorithms, and making necessary adjustments before live deployment, thereby contributing significantly to overall market stability and integrity.[412]

A robust risk management system requires a comprehensive business continuity plan to minimise the impact of disruptions and ensure operational resilience in the face of unforeseen challenges. Key elements of such a plan include implementing redundant systems to eliminate single points of failure, establishing detailed recovery procedures for system outages, workplace disruptions, staff shortages, and failures of external suppliers or data centres, and developing strategies for data loss or corruption. Furthermore, regular staff training on emergency protocols and recovery plans, along with periodic testing of the entire business continuity framework, is essential to maintaining preparedness.[413]

It is therefore proposed that the NZX, overseen by the FMA, mandate that algorithmic traders implement a comprehensive risk management system, structured around “pre-trade, real-time, and post-trade” mechanisms, similar to the German HFT Act and MiFID II. This involves rigorous testing and validation of AI algorithms before deployment, continuous real-time monitoring of their operations on the NZX, and meticulous post-trade reconciliation of all positions, with organisational requirements guided by ESMA's RTS 6. Crucially, recognising the adaptive nature of self-learning AI, New Zealand's framework must also prioritise continuous human oversight and robust checks on algorithmic outputs, learning from authorities such as BaFin, while requiring investment firms to conduct rigorous testing in isolated simulation environments.

[412] Lee and Schu, above n 76, at 9–10.
[413] Commission Delegated Regulation (EU) 2017/589, art 14; Spindler, above n 92, at 218–219.

This holistic approach, encompassing both a structured risk management lifecycle and adaptable oversight for evolving AI, will ensure the integrity and stability of New Zealand's financial markets. Furthermore, New Zealand should mandate comprehensive business continuity plans as a critical component of this robust risk management system. These plans would be essential for minimising disruptions, ensuring operational resilience in the face of unforeseen challenges, and incorporating measures such as redundant systems, detailed recovery procedures, and regular staff training to maintain continuous market operations.

Effective oversight of AI trading algorithms in New Zealand will require a robust system of regular audits, implemented both before and during their deployment. The primary goal of these audits is to proactively identify and swiftly address any vulnerabilities, thereby ensuring the continuous and stable operation of trading systems. A critical component of this audit exercise is that trading algorithms should undergo rigorous stress testing, ideally within a simulated but highly realistic live environment.[414] This is crucial to ensure that they do not behave in unintended ways, that they consistently comply with the firm's obligations and the specific rules and systems of the trading venues, and, most importantly, that they do not contribute to disorderly trading conditions. This level of testing is particularly vital for AI, given its adaptive nature and potential for emergent behaviours.[415]

In the New Zealand context, these audit requirements would be a core mandate of the proposed NZX AI-RO. Beyond initial pre-market approval, the NZX AI-RO would enforce a regime requiring broking firms to regularly reassess the risks associated with AI trading and continuously review their overall risk management programs. A key expectation would be for these firms to promptly flag any AI trading algorithms that fail due to bugs or system failures. Crucially, any such algorithms would be immediately disallowed from operation until all identified bugs or errors are thoroughly fixed and re-validated.

Furthermore, to ensure comprehensive oversight, trading firms should require their Direct Market Access clients to explicitly disclose their use of any AI trading algorithms in their trading activities and to fully cooperate in providing details of their algorithm usage upon request. This multi-layered approach to auditing, encompassing pre-deployment stress testing,

---

[414] Commission Delegated Regulation (EU) 2017/589, above n 379, arts 5(4), 6 and 7; Spindler, above n 92, at 217.
[415] Spindler, above n 92, at 213.

continuous risk reassessment, mandatory bug reporting, and client disclosure, would significantly enhance the integrity and stability of New Zealand's securities market in the face of advanced AI trading.[416]

7.8. MARKET SURVEILLANCE AND ADDITIONAL MEASURES

Chapter 5 demonstrated the evidential and enforcement demands involved in establishing market manipulation after trading has occurred, including the need to analyse trading patterns, market effects and the circumstances surrounding particular orders. These difficulties are amplified where AI systems operate at speeds and volumes beyond ordinary human monitoring. The enhanced surveillance and additional market measures proposed in this section therefore seek to improve the earlier identification and, where possible, prevention of potentially harmful trading patterns.

A significant challenge for the NZX mirrors that of its EU counterparts: a critical technological resource gap that risks leaving the public regulator behind the sophisticated private financial entities they supervise. This means that while financial institutions increasingly use advanced AI-driven algorithms for complex, high-speed trading, regulators often lack the equivalent state-of-the-art market surveillance and analytical tools to effectively monitor, detect, and investigate potential misconduct generated by this technology. The sophistication of private-sector algorithms, including their speed and adaptability, can easily outpace the NZX's existing, slower, or less-integrated monitoring systems. For this reason, similar to the EU approach, granting the NZX more resources to develop, test, and implement advanced Supervisory Technology ("SupTech") to counter sophisticated AI methods is urgent and crucial. However, this advanced SupTech requires the regulator to overcome the dual challenge of lacking the necessary technical resources and, most critically, limited access to proprietary AI systems and training data held by private entities for effective auditing.[417]

To address these technological deficiencies and enhance the FMA's SupTech capabilities, a mandatory AI System Data Access and SupTech collaboration framework, adapted from the EU model, is proposed for the NZ context. The FMC Act should be amended

[416] Spindler, above n 92, at 213.
[417] Azzutti, Ringe and Stiehl, above n 12, at 18.

to mandate a framework requiring private entities utilising high-risk trading AI systems to provide the NZX with read-only access to relevant, anonymised trade data, algorithm specifications, and a SupTech "sandbox" environment for testing. This resolves the critical issues of data access and a lack of a training environment, allowing the NZX to rapidly develop and deploy its own AI monitoring tools (Supervisory AI) trained specifically to detect subtle, evolving patterns of manipulation or systemic risk in real time.[418]

On another front, Beylin and Scopino propose the introduction of a targeted Pigovian tax on order messages or cancellations as a regulatory response to market manipulation in electronic markets .[419] The proposal proceeds from the premise that excessive order submission and rapid cancellation impose real and measurable costs on financial markets that are not presently internalised by the traders or algorithms that generate them.[420] Beylin argues that electronic order books constitute a shared informational and infrastructural resource, and that the overwhelming volume of unexecuted orders, which in some markets exceeds ninety per cent of all order activity, strains exchange systems, distorts price signals, and degrades the quality of information available to other participants.[421] These effects occur regardless of whether the trader intends to manipulate the market, and yet current regulatory regimes generally penalise such conduct only where manipulative intent can be established.[422]

The proposed tax would operate by attaching a small financial cost to excessive order messaging, calculated by reference to objective and readily observable metrics such as an abnormally high order-to-cancellation ratio or order-to-trade ratio measured over short time frames.[423] Where a trader or algorithm exceeds a defined threshold, each additional cancellation beyond that point would automatically incur a nominal fee.[424] The design is deliberately behavioural rather than fault based. Liability would not depend on proving that an order was placed with the intention of cancellation or manipulation but would arise solely from

[418] Azzutti, Ringe and Stiehl, above n 12, at 18.
[419] Ilya Beylin "Taxing Fictive Orders: How an Information-Forcing Tax Can Reduce Manipulation and Distortion in Financial Product Markets" (2018) 85 University of Cincinnati Law Review 1 at 5–10; Scopino, above n 183, at 389–390.
[420] Beylin, above n 419, at 30–31.
[421] Beylin, above n 419, at 15–20; CFTC and SEC, above n 36, at 11–12.
[422] Beylin, above n 419, at 25–28.
[423] Beylin, above n 419, at 30–35; Scopino, above n 183, at 389–390.
[424] Scopino, above n 183, at 389–390.

the trader's observable pattern of activity, so that the deterrent is self-executing and triggered by conduct that exceeds what is reasonably necessary for legitimate trading.[425]

This feature is particularly significant in the context of AI-driven trading. Algorithmic and high-frequency trading systems can generate thousands of orders per second, far beyond human capacity to monitor or interpret in real time, and strategies such as spoofing and quote stuffing rely on this speed and scale to create false impressions of supply and demand or to overwhelm other traders' ability to process market data.[426] As Scopino observes, modern spoofing is often carried out by software algorithms operating at superhuman speed, making traditional intent based enforcement both practically and evidentially difficult.[427] A Pigovian cancellation tax responds directly to this problem by regulating the economic effects of the behaviour rather than its mental element, thereby reducing the enforcement gap created by autonomous or opaque decision-making systems.[428]

The tax is not intended to replace existing remedies including enforcement under the FMC Act, but to operate alongside them as a complementary regulatory layer. Traditional enforcement mechanisms remain essential for addressing serious and deliberate manipulation, particularly where there is clear evidence of intent and significant market harm.[429] However, such enforcement is necessarily reactive, resource intensive, and dependent on complex investigations and litigation. By contrast, a cancellation tax provides continuous, real time market discipline, imposing an immediate cost on distortive behaviour without the need for regulatory intervention or adjudication, so that enforcement resources can be reserved for the most egregious cases while routine or ambiguous instances of excessive activity are addressed automatically through pricing mechanisms.[430]

The regulatory focus is therefore shifted from proving why a trader acted to simply observing what they did. This shift is particularly important where AI systems may generate harmful trading patterns without clear human intent at the point of execution.[431] By imposing an immediate financial consequence on excessive cancellations, the tax acts as a proactive

[425] Beylin, above n 419, at 30–35; Scopino, above n 183, at 389–390.
[426] Scopino, above n 183, at 411, 431 and 436–438.
[427] Scopino, above n 183, at 431 and 436–438.
[428] Scopino, above n 183, at 431 and 436–438.
[429] Financial Markets Conduct Act 2013, pt 5 subpt 3; Beylin, above n 419, at 35–40.
[430] Beylin, above n 419, at 30–35; Scopino, above n 183, at 389–390.
[431] Scopino, above n 183, at 431 and 436–438.

disincentive, reducing the attractiveness of strategies that depend on high volumes of non-executable orders and encouraging firms to design algorithms that are more conservative in their use of shared market infrastructure.[432] In this way, the tax operates not only as a deterrent to manipulation, but also as a governance tool influencing ex ante algorithmic design choices by pricing conduct that would otherwise be "free" to the trader while costly to the market. [433]

Comparable regulatory approaches in other jurisdictions demonstrate that mechanisms targeting excessive order messaging are technically feasible and capable of shaping trading behaviour through objective, behaviour-based constraints, even where they are not framed as taxes in the strict sense. In the European Union, MiFID II requires trading venues to maintain systems and procedures to ensure that algorithmic trading does not create or contribute to disorderly trading conditions, including monitoring and, where appropriate, limiting the ratio of unexecuted orders to transactions.[434] These obligations are given concrete effect through Commission Delegated Regulation (EU) 2017/566, which sets out regulatory technical standards for calculating the ratio of unexecuted orders to transactions at the level of members and financial instruments using trading-venue monitoring.[435] The EU framework is significant not because it imposes a charge, but because it establishes that excessive order submission can be regulated ex ante through quantitative thresholds without reference to intent, thereby recognising excessive messaging as a source of market disorder in its own right.[436]

More explicit pricing mechanisms have been adopted in some markets, most notably in Italy. Italy's financial transaction tax includes a high-frequency trading component that applies where algorithm generated orders are cancelled or modified at an excessive rate within very short time intervals. Under this regime, a tax is triggered where the relevant ratio exceeds 60 per cent, calculated on a daily basis for each financial instrument and applying only to cancellations or modifications occurring within less than half a second.[437] Administrative guidance also clarifies that the calculation is confined to orders generated by computer

---

[432]Beylin, above n 419, at 30–35; Scopino, above n 183, at 389–390.
[433] Beylin, above n 419, at 30–35.
[434] MiFID II, above n 49, art 48(12)(b).
[435] Commission Delegated Regulation (EU) 2017/566 of 18 May 2016 supplementing Directive 2014/65/EU [2017] OJ L 87/84.
[436] Commission Delegated Regulation (EU) 2017/566, above n 435, recital 1.
[437] Law No 228 of 24 December 2012 (Italy), art 1(495) (Italian financial transaction tax); Decree of the Minister of Economy and Finance (Italy), 21 February 2013, arts 12–13.

algorithms meeting the specified characteristics and addresses exclusions for certain order types and platform-driven cancellations.[438] While the Italian model has been criticised for its breadth and potential liquidity effects, it nonetheless demonstrates that exchanges can monitor cancellation behaviour at sub second granularity and impose automated charges by reference to observable conduct alone, without any inquiry into manipulative intent.[439]

A more targeted illustration of how such mechanisms can be calibrated is provided by the approach adopted by the Oslo Stock Exchange. Oslo introduced a fee payable by market participants whose order to trade ratios exceeded 70:1, accompanied by exemptions designed to protect genuine liquidity provision, and empirical analysis reports that following introduction of the fee traders reduced message traffic while standard measures of market quality, including spreads, depth and realised volatility, remained largely unchanged.[440] The authors contrast this outcome with experiences in other markets where less carefully designed interventions were associated with liquidity deterioration, underscoring the importance of calibration.[441] Notably, the Oslo fee functioned primarily as a behavioural deterrent rather than as a revenue raising instrument, with compliance achieved largely through the credible prospect of a charge rather than its routine imposition.[442]

Taken together, these comparative examples support the conceptual foundations of a Pigovian cancellation tax. They show that modern trading venues can monitor order level behaviour in real time and constrain excessive messaging through quantitative thresholds, and that automated constraints or charges tied to message traffic can meaningfully alter trader incentives without necessarily undermining market quality.[443] This experience aligns closely with Beylin's argument that order book information should be understood as a shared resource vulnerable to distortion by excessive cancellations, and that intent based anti-manipulation regimes are poorly suited to the scale and speed of AI-driven trading.[444] Viewed in this light, a cancellation fee represents an incremental extension of existing regulatory practice rather than

---

[438] Agenzia delle Entrate (Italy) *Financial Transaction Tax: High-Frequency Trading — Explanatory Notes* (Agenzia delle Entrate, Rome, 2013).
[439] Agenzia delle Entrate, above n 438.
[440] Kjell Jørgensen, Johannes Skjeltorp and Bernt Arne Ødegaard "Throttling Hyperactive Robots — Order-to-Trade Ratios at the Oslo Stock Exchange" (2018) 37 Journal of Financial Markets 1 at 5–11.
[441] Jørgensen, Skjeltorp and Ødegaard, above n 440, at 11–17.
[442] Jørgensen, Skjeltorp and Ødegaard, above n 440, at 17–22.
[443] Commission Delegated Regulation (EU) 2017/566, above n 435; Agenzia delle Entrate, above n 438; Jørgensen, Skjeltorp and Ødegaard, above n 440.
[444] Beylin, above n 419, at 30–35.

a radical departure from it, operating as an automated complement to conventional enforcement rather than a substitute.[445]

7.9. AI GOVERNANCE LIABILITY: A STATUTORY CAUSE OF ACTION FOR SUPERVISORY FAILURES

The analysis in Chapters 4 and 5 demonstrated the difficulties that AI-driven trading may create for the application and enforcement of existing market manipulation provisions, while Chapter 6 examined whether private law could provide an alternative or complementary basis for liability. As discussed in Chapter 6, negligence law is a doctrinally constrained basis for liability in the context of AI-driven trading. Its limitations are particularly acute where harm takes the form of pure economic loss suffered by diffuse market participants, or where the causal chain is obscured by the opacity and autonomy of AI systems. Courts have traditionally required a confined relational setting or an assumption of responsibility to impose a duty of care, and they remain reluctant to recognise open-ended duties in financial markets. These constraints may leave victims without an effective common law remedy.

However, the FMC Act already provides a statutory route to compensation where the conduct contravenes a civil liability provision. A declaration of contravention may support an application by the FMA or an affected person for a compensatory order under Part 8. The practical gap addressed by the proposed AGL therefore concerns cases in which harmful AI governance failures cannot be established as a contravention of an existing civil liability provision, or where causation, identification of the responsible actor, or proof of loss makes the existing route ineffective.

In this context, the New Zealand Supreme Court's decision in *Smith v Fonterra Co-operative Group Ltd* is relevant, although its significance should be understood in a limited and carefully qualified way. The Court did not assess the merits of the claims or endorse any expansion of the duty of care. Instead, it declined to strike out novel causes of action, including public nuisance, negligence, and a proposed climate duty, allowing them to proceed to trial. This procedural decision did not indicate a likelihood of success but demonstrated the Court's

[445] Beylin, above n 419, at 30–35; Scopino, above n 183, at 389–390.

willingness to allow novel legal claims to be tested through evidence.[446] The subsequent legislative response to the decision further demonstrates both aspects of the relationship between the common law and legislation that are relevant to the proposed AGL framework. Courts may permit novel causes of action to be tested where existing legislation has not clearly displaced the common law, but Parliament retains the ability to define, restrict or exclude tort liability where it considers legislative intervention appropriate. Accordingly, *Smith v Fonterra* demonstrates procedural openness to the development of novel claims, but it does not establish that novel tort liability will ultimately be recognised or remain operative. This reinforces the approach adopted in this thesis of grounding the proposed AGL framework expressly in legislation rather than relying solely on judicial development of the common law.

Scopino proposes a liability framework based on supervisory negligence, directly addressing harmful trading practices by holding accountable the individuals tasked with controlling and overseeing these systems. This pragmatic approach imposes clear supervisory duties to ensure AI trading algorithms are used responsibly and ethically, allowing for legal action against a broad spectrum of misconduct, including both unintended algorithmic errors and systemic organisational failures. This concept aligns with existing financial regulation, specifically the enforcement actions of the Commodity Futures Trading Commission ("CFTC"), which routinely pursues failure-to-supervise claims against firms that neglect their duty to oversee employees managing automated trading systems adequately.[447]

U.S. Federal regulations impose a non-delegable duty of diligent supervision on CFTC registrants, requiring them to actively oversee their personnel and agents to protect customers from fraud or manipulation. This robust regulatory power enables the CFTC to pursue enforcement actions for supervisory failures even when no specific rule was broken. Evidence of such failure is often found in systemic deficiencies, such as a failure to enforce internal compliance policies or to adequately monitor for risk, which allow repeated or high-risk violations to go undetected, proving the firm's supervisory system itself was faulty. However, establishing liability for a supervisory failure is not automatic, as a single, isolated wrongful

---

[446] *Smith v Fonterra Co-operative Group Ltd [2024] NZSC 5* at [167] – [170]. See also the Climate Change Response (Tort Liability) Amendment Act 2026, which subsequently introduced a statutory restriction on tort liability for specified emissions-related climate change effects and applies to proceedings that had not been finally determined at its commencement.

[447] Scopino, above n 183, at 391–392.

act by an employee, without additional proof of the firm's overall deficient oversight, is typically insufficient to prove a systemic breach of the regulatory duty.[448]

Scopino clarifies that to prove a failure-to-supervise claim, regulators must evaluate the adequacy of the entire supervisory system, examining the specific roles and responsibilities of employees and whether those individuals diligently carried out their assigned duties. As a necessary first step, registrants must establish robust procedures designed to detect and prevent misconduct, including training employees on relevant regulations and strictly enforcing internal compliance policies. The CFTC's supervisory requirement does not impose strict liability (automatic fault). Instead, it applies a reasonableness standard, holding registrants accountable only for failures to implement reasonable supervisory systems or for neglecting to take appropriate action after noticing red flags that indicate potential misconduct. Ultimately, this analysis focuses on whether, given the available information and the firm's system quality, the registrant should reasonably have been able to detect the specific instances of misconduct.[449]

Scopino further emphasises that an adequate supervisory system must feature automated alerts for suspicious patterns and other computerised controls to respond quickly to potential misconduct. The CFTC explicitly mandates that firms diligently supervise all officers, employees, and agents responsible for establishing and maintaining electronic trading platforms, including monitoring trading accounts. As a more specific regulatory analogy, NFA Interpretive Notice 9060 applies to Forex Dealer Members operating electronic trading systems and requires written procedures addressing matters including security, capacity, credit and risk-management controls, recordkeeping and trade integrity.[450] Although this framework concerns electronic foreign exchange trading rather than AI-driven securities trading generally, it illustrates how detailed supervisory and systems-control obligations can supplement a broader duty of diligent supervision. Crucially, when proving a supervisory violation, one need not prove the intent behind the underlying improper conduct (e.g., wash sales or price manipulation). Instead, one must demonstrate that the trading firm either lacked an adequate supervisory system or failed to diligently oversee its employees and agents.[451]

---

[448] Scopino, above n 183, at 392–394.
[449] Scopino, above n 183, at 395–399.
[450] National Futures Association, Interpretive Notice 9060.
[451] Scopino, above n 183, at 400–401.

A firm that allowed its customers to use their accounts for trades resembling wash trades, market manipulation, or other disruptive practices could face liability for supervisory failure. To be exonerated, the firm must demonstrate reasonable diligence, including providing evidence that it promptly noticed suspicious activity, thoroughly investigated the trading patterns, and determined that the trading did not distort or disrupt prices. To strengthen oversight of algorithmic trading, Scopino suggests that employees responsible for programming and operating these systems must adhere to strict standards to actively prevent manipulative or disruptive practices such as wash trading or market manipulation.[452]

Building on Gregory Scopino's proposal of a supervisory negligence-based approach to algorithmic trading misconduct, this thesis introduces AI Governance Liability ("AGL") as a statutory cause of action for systemic failures in the governance and supervision of AI trading systems. AGL is framed as a tort claim with the familiar four elements of duty, breach, causation, and damage, but it explicitly recognises that the duty at its core, the Duty of Diligent Governance, would likely be defined by statute or regulation. This reflects the reality that the interest protected by AGL, the integrity of financial markets in the face of AI-driven trading, is not a traditional natural right typically protected by common law torts. As Edelman J observed in *Mallonland Pty Ltd v Advanta Seeds Pty Ltd*, tort law generally protects pre-existing rights, while novel obligations responding to emerging risks may require legislative creation.[453] AGL would therefore not depend on judicial recognition of a new common law duty of care. Instead, Parliament would define the relevant governance obligation and provide a corresponding statutory right of action where breach of that obligation causes recoverable harm.

The first element of the proposed AGL claim is the existence of a Duty of Diligent Governance. The primary duty would be imposed on the trading firm that deploys or operates the relevant AI trading system. It would require the firm to take reasonable steps to supervise and control the AI system throughout its operational lifecycle. In practical terms, the duty would require the firm to establish and maintain an adequate governance framework, including appropriate design and testing standards, continuous monitoring mechanisms, escalation procedures and internal controls designed to ensure that autonomous trading systems operate

[452] Scopino, above n 183, at 402–405.
[453] *Mallonland*, above n 274, at [102] – [103] per Edelman J.

within applicable legal and market conduct requirements. The proposed AGL framework would not automatically impose personal civil liability on directors, senior officers or individual employees merely because they hold supervisory responsibilities. Any separate personal liability would require express statutory provision and an appropriately defined threshold of responsibility. This approach keeps the proposed cause of action focused on systemic organisational governance failures rather than individual errors.

The CFTC's Regulation 166.3 provides a useful regulatory analogy by requiring relevant registrants to exercise diligent supervision over specified activities undertaken by their partners, officers, employees and agents.[454] The regulation does not itself impose a private damages regime for failures involving AI systems, nor does it treat algorithms as employees or agents. Its relevance to AGL is therefore analogical rather than precedential. It demonstrates that financial regulation can impose an affirmative obligation on firms to maintain adequate supervisory systems rather than focusing exclusively on the underlying misconduct. The proposed AGL framework would adapt that supervisory principle to AI-driven trading by imposing a legislatively defined governance obligation on firms responsible for deploying and controlling AI trading systems.

The second element would require the claimant to establish breach of the statutory Duty of Diligent Governance by demonstrating a material governance or supervisory failure. The inquiry would focus on the adequacy of the defendant firm's overall governance system rather than solely on the particular trading incident that caused the loss. The central question would be whether the firm failed to take the reasonable governance measures required by the statutory standard in designing, deploying, monitoring or controlling its AI trading system.

Evidence of breach could include systemic shortcomings in the firm's governance framework. These could include inadequate design and testing, such as deploying an AI system without appropriate testing against foreseeable extreme market conditions or manipulative scenarios. They could also include inadequate monitoring, such as failing to maintain effective real-time surveillance or alert systems capable of identifying erratic or potentially abusive trading behaviour or failing to respond appropriately to identified warning signs. A further

[454] 17 CFR § 166.3.

example would be a failure to implement or enforce internal controls, such as maintaining written policies governing algorithmic trading but failing to apply them in practice.

What matters under AGL is whether the firm's overall governance and risk-control framework fell below the standard required by the statutory Duty of Diligent Governance, having regard to the nature and risks of the AI trading system. This approach is consistent with the systemic focus of supervisory negligence in regulatory contexts. The CFTC, for example, does not impose supervisory liability merely because an isolated unauthorised act occurs, but may find a supervisory failure where deficiencies in the firm's oversight system contributed to the misconduct.[455] Similarly, breach under AGL would depend on whether the firm failed to take the governance measures reasonably required by the statutory standard, taking into account relevant regulatory requirements, industry standards and the firm's knowledge of the capabilities and risks of the AI system.[456]

Third, the claimant would be required to establish causation between the firm's breach of the Duty of Diligent Governance and the claimed loss. This would involve both factual causation and the statutory scope of liability. The claimant would first need to establish that, but for the relevant governance failure, the harmful AI-driven market conduct would not have occurred, or that the governance failure materially contributed to that conduct. The claimant would then need to establish that the resulting loss fell within the class of harm against which the statutory duty was intended to protect.

Causation may be particularly difficult where several algorithms, firms or market events contribute to the same market disruption. The proposed statutory framework should therefore permit the court to consider multiple contributing causes rather than requiring the claimant to identify a single algorithm as the exclusive cause of the loss. Intervening market conduct would remain relevant. Liability should not extend to losses that are too remote from the governance failure or that result predominantly from an independent intervening event outside the risks that the Duty of Diligent Governance was intended to address.

[455] Scopino, above n 183, at 391–399.
[456] Scopino, above n 183, at 400–405.

The evidential difficulties created by opaque AI systems should also be recognised. Where relevant technical information is predominantly within the defendant firm's control, the statutory framework could provide appropriate disclosure obligations and, in defined circumstances, rebuttable evidential presumptions. Such mechanisms would assist claimants in establishing causation without removing the requirement to demonstrate a sufficient connection between the governance failure, the AI-driven market conduct and the claimed loss.

The proposed legislation would also need to define the class of persons entitled to bring an AGL claim. Because AI-driven market misconduct may affect numerous market participants who have no direct relationship with the defendant firm, the cause of action should not depend on contractual proximity or an assumption of responsibility. Instead, standing should extend to persons who can demonstrate that they suffered loss of a kind protected by the statutory Duty of Diligent Governance because of the defendant firm's breach. This could include investors or other market participants who traded at prices materially affected by the relevant AI-driven conduct. However, the statutory framework should impose appropriate limits on recoverable loss to avoid indeterminate liability to an unlimited class of market participants.

Finally, the claimant must establish recoverable damage resulting from the breach. In an AGL claim, this would typically comprise quantifiable financial loss directly attributable to the relevant AI-driven market event, such as losses incurred through transactions executed at materially distorted prices or reasonable costs incurred in mitigating the consequences of the event. The proposed legislation would need to define the categories and limits of recoverable economic loss to provide meaningful compensation while avoiding indeterminate liability for general market movements or remote losses.

AGL would complement, rather than replace, existing remedies under the FMC Act. As discussed in Chapter 4, the market manipulation provisions are civil liability provisions and Part 8 provides a statutory route through which the FMA or an affected person may seek compensation following a contravention.[457] However, that route depends on establishing a contravention of an existing civil liability provision. AGL would address the distinct situation in which compensable harm results from a firm's breach of its statutory Duty of Diligent Governance, including circumstances where the relevant governance failure cannot readily be

[457] Financial Markets Conduct Act 2013, ss 485 – 487 and 494 – 495.

established as market manipulation or another existing statutory contravention. The prospect of civil liability under AGL would also reinforce deterrence by creating an additional financial incentive for firms to invest in effective governance, compliance and risk mitigation.

In summary, the proposed AGL would establish a statutory cause of action for breach of a legislatively defined Duty of Diligent Governance. Its structure would draw on familiar concepts of breach, causation and damage, while its scope would be expressly defined by legislation to address the distinctive risks created by AI-driven trading. The proposal draws on regulatory analogies such as the CFTC's duty of diligent supervision but does not treat those regimes as direct precedents for private civil liability. By focusing on systemic failures in the human governance and supervisory structures surrounding AI trading systems, AGL would complement the existing FMC Act civil liability framework and provide a distinct avenue for redress where harm results from deficient AI governance that does not readily fall within an existing civil liability provision. Its statutory foundation would also provide clearer limits concerning the persons subject to liability, the persons entitled to claim, recoverable loss, causation, defences and remedies, thereby reducing the uncertainty that would arise from relying solely on judicial development of a novel common law duty.

7.10. NOVEL PRODUCT LIABILITY FRAMEWORK

7.10.1. Rationale

As demonstrated in Chapter 6, New Zealand's existing liability frameworks may leave a significant accountability gap when applied to AI-driven trading systems. This gap must, however, be considered alongside the existing civil liability and compensation provisions in Part 8 of the FMC Act. As discussed in Chapter 4, those provisions provide a statutory route to compensation where loss results from a contravention of an applicable civil liability provision, including relevant Part 5 market provisions.[458] They do not, however, establish a general strict liability regime merely because an AI trading product is defective, malfunctions, or causes financial loss. The CGA may also have limited application where AI trading systems are acquired or used for commercial or productive purposes, while the FTA may address misleading or deceptive conduct but does not establish a general strict product liability regime

[458] Financial Markets Conduct Act 2013, ss 385, 485–487 and 494–495.

for defective AI trading systems. The accountability gap identified in this thesis therefore concerns circumstances in which loss results from a defective AI trading product but cannot readily be addressed through an existing FMC Act civil liability provision, consumer protection legislation, or negligence.

To address this accountability gap, this chapter proposes the introduction of a statutory product liability regime under the FMC Act, specifically tailored to AI-driven financial products. This regime would impose strict liability on the providers, developers, and distributors of AI trading systems for defects causing foreseeable harm to users, regardless of fault. The proposal is grounded in comparative developments, particularly the European Union's revised Product Liability Directive 2024, which has extended strict liability to software and AI systems in recognition of the unique risks posed by digital and autonomous technologies.[459]

By treating AI trading systems as "products" subject to no-fault liability, the FMC Act-based regime would ensure that retail investors and other market participants can obtain compensation when AI systems fail, without the need to satisfy the demanding prerequisites of negligence or consumer law. This approach aligns with New Zealand's principles of proportionate, technology-neutral regulation and ensures that AI-driven financial innovation proceeds with appropriate safeguards and accountability mechanisms in place.

Importantly, this product liability regime would complement the pre-market approval and conformity assessment regime proposed in Section 7.3, as well as the compliance by design measures in Section 7.4. Whereas those provisions establish ex ante standards to prevent defects from arising, the FMC Act product liability regime provides ex post redress when prevention fails. Together, these mechanisms create a comprehensive framework that incentivises safe AI design, deployment, and operation, while ensuring that victims of AI-related harm are not left without effective legal recourse.

[459] Directive (EU) 2024/2853, above n 354.

7.10.2. Comparative Inspiration: The EU Product Liability Directive 2024

The proposal for an FMC Act-based product liability regime draws on the European Union's recent reform of its long-standing product liability framework. In November 2024, the EU adopted Directive (EU) 2024/2853 (the "Revised PLD"), which modernises the previous product liability regime to address developments including digital technologies, software, artificial intelligence, and increasingly complex supply chains.[460] The Revised PLD is relevant to the New Zealand context because it provides a recent legislative model for extending strict liability principles to software and other digital products. However, the proposed New Zealand regime would not simply replicate the Revised PLD. As discussed below, important differences would be necessary to address the distinctive nature of AI-driven financial trading, particularly because the principal harm in this context is likely to be standalone financial loss.

One of the most significant changes introduced by the Revised PLD is the express inclusion of software within the definition of a "product".[461] This removes much of the uncertainty surrounding the treatment of intangible software under the previous regime and means that software used in or supplied as part of an AI system may fall within the product liability framework. The Directive therefore provides an important comparative model for considering liability arising from software-based AI trading systems. However, the Directive does not establish that every adverse output, inaccurate prediction, or poor investment performance produced by an AI system constitutes a product defect. Liability continues to depend on whether the product satisfies the Directive's test of defectiveness, which is concerned with the safety that a person is entitled to expect.[462]

The Revised PLD also adopts a concept of defectiveness capable of accommodating the characteristics of software and AI. A product is defective where it does not provide the safety that a person is entitled to expect or that is required under Union or national law, taking into account all relevant circumstances.[463] Those circumstances include matters such as the product's presentation and reasonably foreseeable use, its ability to continue learning or acquire new features after being placed on the market, relevant product safety requirements, and the

[460] Directive (EU) 2024/2853, above n 354, recitals 1–5.
[461] Directive (EU) 2024/2853, above n 354, art 4(1).
[462] Directive (EU) 2024/2853, above n 354, art 7.
[463] Directive (EU) 2024/2853, above n 354, art 7(1).

reasonably foreseeable effect of other products on it.[464] The relevant inquiry is therefore one of safety rather than whether the product achieved a particular commercial or investment outcome. In the AI trading context, poor investment performance or an inaccurate prediction would not, without more, necessarily establish defectiveness.

The Revised PLD also recognises that the safety of software and related digital products may be affected after they have been placed on the market, including through updates, upgrades, modifications and the capacity of certain systems to continue learning.[465] Where a product remains within a manufacturer's control, relevant circumstances may therefore include the effect of software updates or upgrades and the failure to provide updates or upgrades necessary to maintain safety.[466] These features are particularly relevant to AI trading systems because their operation may change over time. However, any proposed New Zealand regime would need to distinguish between an AI system that becomes unsafe or defective and one that merely produces an unsuccessful investment outcome.

To address the evidential difficulties created by technically complex products, the Revised PLD introduces mechanisms concerning both disclosure and the burden of proof. Article 9 provides for disclosure of relevant evidence in specified circumstances, while Article 10 establishes rebuttable presumptions concerning defectiveness and causation.[467] These presumptions may apply, for example, where a defendant fails to comply with an obligation to disclose relevant evidence, where the product does not comply with mandatory product safety requirements intended to protect against the relevant risk, or where the claimant faces excessive difficulties in proving defectiveness or causation because of the technical or scientific complexity of the case.

These evidential mechanisms are particularly significant in the AI context, where the technical complexity and opacity of some systems may make it difficult for claimants to identify why a system failed or to obtain the information necessary to establish defectiveness and causation. By providing disclosure mechanisms and rebuttable presumptions in defined

[464] Directive (EU) 2024/2853, above n 354, art 7(2).
[465] Directive (EU) 2024/2853, above n 354, arts 4(5) and 7.
[466] Directive (EU) 2024/2853, above n 354, art 11(2).
[467] Directive (EU) 2024/2853, above n 354, arts 9–10.

circumstances, the Revised PLD seeks to reduce evidential barriers without automatically imposing liability whenever an AI or software-based product causes harm.

The Revised PLD does not create a general right to recover pure economic loss. Article 6 identifies the categories of compensable damage, including death or personal injury, specified damage to or destruction of property, and destruction or corruption of data that is not used for professional purposes, together with material losses resulting from those forms of damage.[468] Standalone financial loss, such as a reduction in investment value caused by an unsuccessful or defective trading decision, is not itself included as a general category of compensable damage under the Directive.

This limitation is particularly important in the AI trading context because the principal harm caused by a defective trading system may be purely financial. The proposed New Zealand regime therefore deliberately departs from the Revised PLD by permitting recovery of defined and reasonably foreseeable economic loss caused by a defective AI financial product. The EU model is consequently used as a comparative foundation for matters such as the treatment of software, defectiveness, disclosure and evidential presumptions, rather than being adopted wholesale.

Finally, the Revised PLD addresses the difficulties created by complex and international supply chains by identifying a range of economic operators that may be liable for defective products, including manufacturers, importers and authorised representatives and, in specified circumstances, other economic operators and online platforms.[81] This approach is relevant to New Zealand because AI trading systems may be developed or supplied from outside the jurisdiction. The proposed New Zealand regime similarly seeks to ensure that an appropriate defendant can be identified within the jurisdiction where an AI financial product supplied in New Zealand causes compensable harm.

A further distinction is necessary concerning the persons protected by the respective regimes. The Revised PLD is principally directed towards compensation for natural persons suffering the categories of damage recognised by the Directive. The regime proposed in this thesis would be broader because its purpose is to address defective AI financial products

[468] Directive (EU) 2024/2853, above n 354, art 6.

operating within financial markets, where relevant claimants may include both individual investors and other market participants. This represents a deliberate departure from, rather than a direct transplantation of, the EU model.

Taken together, these features of the Revised PLD provide a useful comparative framework for considering how product liability principles may respond to software and AI-related risks. Its treatment of software as a product, its safety-based approach to defectiveness, its disclosure mechanisms and evidential presumptions provide useful guidance for the proposed New Zealand regime. However, the Revised PLD cannot simply be transplanted into the AI trading context. Its limitations concerning recoverable damage and the persons protected are particularly significant where the principal harm is standalone financial loss suffered by investors or other market participants. The proposed FMC Act regime therefore draws selectively on the Revised PLD while departing from it where necessary to address the distinctive risks and losses associated with AI-driven financial trading.

7.10.3. Proposed FMC Act Amendments

To operationalise a product liability regime for AI financial products in New Zealand, this thesis proposes amendments to the FMC Act. Rather than enacting a standalone statute or attempting to retrofit the CGA which is ill-suited to commercial or productive use cases, the FMC Act offers the most suitable legislative framework for implementing this reform. While it is primarily concerned with regulating market conduct, its overarching purpose is to promote confident and informed participation in financial markets and to ensure that those markets are fair, efficient, and transparent. Its existing regulatory architecture, particularly the provisions on disclosure, licensing, and fair dealing provides a coherent foundation for integrating product liability obligations without the need for a separate legislative instrument. Moreover, the FMC Act already regulates financial products and services and imposes a range of obligations on market participants.[469] Extending it to include product liability for AI trading systems would therefore integrate seamlessly with the current framework, ensuring consistency and minimising regulatory fragmentation.

[469] Financial Markets Conduct Act 2013, ss 3, 6 and 431.

The FMC Act is technology-neutral and principles-based, which makes it well-suited to accommodate emerging technologies such as AI without requiring frequent legislative amendments. Moreover, the FMA already has expertise in supervising AI-related risks in financial markets, and it is well-positioned to oversee and enforce a product liability regime for AI financial products. Incorporating product liability within the FMC Act ensures close alignment with the pre-market approval regime proposed in Section 7.3, enabling regulators to treat compliance with conformity assessments as relevant evidence of non-defectiveness, while making clear that regulatory approval does not immunise providers from liability.

The proposed amendments would introduce a new Part within the FMC Act, potentially titled "Product Liability for AI-Driven Financial Products." The core elements of this regime would be as follows. First, the amendments would define "AI financial product" broadly to include any software, algorithm, or automated system that is designed, marketed, or used for the purpose of making, recommending, or executing financial decisions, including trading, portfolio management, risk assessment, and market analysis.

This definition would encompass standalone AI trading platforms, algorithmic trading tools, robo-advisers, and any other AI-enabled system that operates in financial markets. The definition would be technology-neutral, ensuring that it applies regardless of the specific AI technique used, whether rule-based algorithms, machine learning models, or hybrid systems. Importantly, the definition would not be limited to products supplied to "consumers" in the CGA sense; rather, it would apply to any AI financial product supplied in trade, thereby ensuring that retail investors using AI systems for wealth accumulation are covered.

Second, the regime would impose strict liability on the "provider" of an AI financial product for harm caused by a defect in the product. The term "provider" would be defined to include the developer, manufacturer, importer, distributor, or any person who places the AI financial product on the market or puts it into service in New Zealand. This broad definition ensures that liability attaches to parties operating within New Zealand's jurisdiction, even where the AI system was developed overseas or by a third party. The regime would also impose joint and several liability on multiple providers where appropriate, ensuring that victims can recover from any solvent party in the supply chain. Providers would, in turn, have rights of contribution and indemnity against upstream parties, preserving incentives for safe design and discouraging the externalisation of risk onto end users.

Third, the amendments would define "defect" in a manner tailored to the characteristics of AI financial products. A product would be deemed defective if it does not provide the level of safety and reliability that persons are entitled to expect, taking into account all circumstances, including: (a) the product's presentation and marketing, including any representations made about its performance, accuracy, or suitability for particular purposes; (b) the reasonably expected use of the product, including use by retail investors for wealth accumulation; (c) the time the product was put into service, recognising that AI systems may degrade, drift, or become obsolete over time; and (d) any applicable industry standards, regulatory guidance, or conformity assessments conducted under the pre-market approval regime proposed in Section 7.3 of this chapter.

Importantly, compliance with applicable regulatory requirements would not necessarily determine whether an AI financial product is defective, although such compliance would remain relevant to the assessment of defectiveness. A product could nevertheless be defective because of matters such as deficiencies in its design, inadequate testing, failure to provide necessary updates or maintenance, or inadequate warnings or instructions. The proposed regime should also recognise an appropriate defence where the alleged defect results directly from compliance with mandatory legal requirements.[470] This approach draws on the Revised PLD while ensuring that regulatory approval or conformity assessment does not, by itself, provide complete immunity from liability.

Fourth, the regime would introduce disclosure mechanisms and rebuttable presumptions to address the evidential difficulties faced by claimants, drawing on Articles 9 and 10 of the Revised PLD. Article 9 addresses disclosure of evidence, while Article 10 provides rebuttable presumptions concerning defectiveness and causation in specified circumstances. Drawing on these mechanisms, the proposed New Zealand regime would provide for rebuttable presumptions of defectiveness or causation in defined circumstances, including: (a) where the provider has failed to comply with an obligation to disclose relevant evidence under the proposed regime, thereby materially impairing the claimant's ability to establish defectiveness or causation; (b) where the AI financial product has clearly malfunctioned when used as intended, in circumstances supporting an inference that the product was defective; or (c) where the claimant demonstrates that it would be excessively

[470] Directive (EU) 2024/2853, above n 354, arts 7 and 11(1)(d).

difficult to prove defectiveness or causation because of the technical or scientific complexity or opacity of the AI system, provided that the claimant has produced sufficient evidence to support the application of the relevant presumption.

These presumptions would be rebuttable, meaning that the provider could avoid liability by proving that no defect existed, that the defect did not cause the harm, or that the harm was due to misuse or intervening causes. However, the burden of rebuttal would rest with the provider, reflecting the principle that the party with superior knowledge and control over the product should bear the evidential burden.

Fifth, the regime would provide for a broad definition of compensable damage, including pure economic loss. Recognising that the primary harm caused by defective AI trading systems is financial loss, the amendments would permit recovery for: (a) direct financial loss arising from defective trading decisions, including losses on executed trades, opportunity costs, and transaction costs; (b) consequential economic loss, including loss of portfolio value and costs of remediation or mitigation; and (c) loss or corruption of data, including trading data, account information, and personal financial data. Damages would be limited to reasonably foreseeable harm, and claimants would be required to mitigate their losses where reasonable. However, the regime would not impose arbitrary caps on recovery, ensuring that victims of large-scale AI failures are not left under-compensated.

Sixth, the regime would impose ongoing obligations on providers to monitor, maintain, and update AI financial products throughout their lifecycle. Providers would be required to implement post-market monitoring systems to detect defects, performance degradation, or emerging risks, and to provide timely updates, patches, or warnings to users. Failure to comply with these obligations would be treated as evidence of defectiveness, and providers could be held liable for harm caused by known defects that were not remedied within a reasonable time. This approach aligns with the governance and risk management obligations proposed in Sections 7.6 and 7.7, and it ensures that product liability is not a one-off assessment at the point of deployment but a continuing responsibility throughout the AI system's operational life.

Finally, the regime would include procedural provisions to facilitate access to evidence and to ensure fair and efficient dispute resolution. Providers would be required to retain technical documentation, audit logs, and version histories for AI financial products for a

specified period, and claimants would have a right to request disclosure of relevant information, subject to appropriate safeguards for confidential or proprietary data. Courts would have the power to appoint independent technical experts to assess defectiveness or causation, and to allocate the costs of such assessments equitably between the parties. Additionally, the regime would permit representative or class actions where multiple claimants have suffered harm from the same defective product, ensuring that victims of large-scale AI failures can achieve collective redress without duplicative litigation.

These proposed amendments would establish a comprehensive and coherent product liability regime for AI financial products within the FMC Act. The regime would be strict (no-fault), technology-neutral, and focused on outcomes rather than conduct. It would ensure that retail investors and other users of AI trading systems can obtain compensation when those systems fail, without needing to satisfy the demanding prerequisites of negligence or consumer law. At the same time, the regime would preserve incentives for innovation by limiting liability to foreseeable harm, by providing clear standards for defectiveness, and by ensuring that providers can rebut presumptions and defences where appropriate. In this way, the regime would strike an appropriate balance between accountability and innovation, consistent with New Zealand's regulatory philosophy.

7.10.4. Policy Justification and Legal Coherence

The proposed FMC Act-based product liability regime is not only doctrinally sound but also normatively justified on policy grounds. It addresses a clear and pressing accountability gap, aligns with New Zealand's regulatory philosophy, and reflects international trends toward adapting product liability principles to the digital age. First, the regime is necessary to address the remedial gap left by existing law. As Chapter 6 demonstrated, the CGA does not protect retail investors who use AI trading systems for wealth accumulation, negligence law imposes prohibitive barriers to recovery for pure economic loss, and the FTA addresses only misleading representations, not design defects or performance failures.

This gap is not merely theoretical: it leaves a large and growing class of users, retail investors who rely on AI systems to manage their financial affairs without effective legal recourse when those systems fail. The proposed regime fills this gap by providing a clear,

accessible, and predictable remedy for harm caused by defective AI financial products, without requiring claimants to satisfy the demanding prerequisites of fault-based liability or consumer law. In doing so, it ensures that the law keeps pace with technological change and that victims of AI-related harm are not denied redress simply because their case does not fit neatly into traditional doctrinal categories.

Second, the regime aligns with New Zealand's principles of proportionate, technology-neutral regulation. New Zealand has consistently favoured regulatory approaches that are flexible, adaptive, and focused on outcomes rather than prescriptive rules. The proposed product liability regime reflects this philosophy. It is technology-neutral, applying to any AI financial product regardless of the specific AI technique used. It is outcome-focused, imposing liability based on defectiveness and harm rather than on compliance with detailed technical standards. It is proportionate, limiting liability to foreseeable harm and providing defences and rebuttals where appropriate. And it is integrated with existing regulatory frameworks, building on the FMC Act's existing structure and the FMA's expertise rather than creating a new, standalone regime. This approach minimises regulatory fragmentation, reduces compliance costs, and ensures that the regime can adapt to future technological developments without requiring constant legislative amendments.

Third, the regime reflects an emerging international trend towards reconsidering how product liability principles should apply to digital and AI-based products. The EU's revised Product Liability Directive is the most prominent example, but it is not the only one. Developments in the United States also illustrate emerging attempts to apply product and design-based liability principles to harms associated with software and algorithmic systems, although the legal position remains unsettled and does not establish a general rule that AI software constitutes a product for strict liability purposes.[471] Commentators have nevertheless argued that product liability principles may provide useful incentives for safer AI design and may address some of the difficulties associated with proving negligence in cases involving complex and opaque systems.[472]

---

[471] *Garcia v Character Technologies Inc*, above n 358; *Lemmon v Snap Inc*, above n 359.
[472] Matthew U Scherer "Of Wild Beasts and Digital Analogues: The Legal Status of Autonomous Systems" (2018) 19 Nevada Law Journal 259 at 280–290; Bryan Casey and Mark A Lemley "You Might Be a Robot" (2019) 105 Cornell Law Review 287 at 300–310.

This trend reflects a recognition that AI and software are not fundamentally different from physical products when it comes to safety and accountability: both can be defective, both can cause harm, and both should be subject to strict liability principles that incentivise responsible innovation and ensure fair compensation for victims. New Zealand's adoption of a similar approach would therefore align the country with international best practice and ensure that its legal framework remains competitive and attractive to responsible AI developers while providing robust protection for users.

Finally, the regime is justified by the unique characteristics of AI financial products and the distinctive harms they can cause. AI trading systems are complex, opaque, and autonomous. They operate at speeds and scales beyond human capacity to monitor in real time. They can degrade, drift, or malfunction in ways that are difficult to predict or explain. And they can cause rapid, widespread, and purely financial harm to large numbers of users. These characteristics make AI financial products particularly well-suited to strict liability: the risks are foreseeable, the harm is often difficult to trace to a specific negligent act, and the parties best placed to manage the risks are the developers and providers of the systems, not the users. Strict liability ensures that these parties internalise the costs of AI-related harm, creating incentives to invest in safety, testing, and ongoing monitoring. It also ensures that victims are not left without recourse simply because they cannot prove fault or because the AI system's decision-making is too opaque to penetrate. In this way, the regime aligns legal responsibility with practical control and superior knowledge, consistent with the foundational principles of product liability law.

7.11. VICTIM COMPENSATION

High-speed, autonomous trading programs can cause rapid, widespread damage, for instance, a sudden “flash crash” or a cascade of manipulative orders that affect many market participants simultaneously.[473] These challenges have prompted calls to supplement the traditional enforcement toolkit with specialised compensation mechanisms that can promptly make injured parties whole without diluting firms’ incentives to manage their technologies responsibly. Two such mechanisms have gained particular attention: (1) mandatory liability

[473] Alessio Azzutti "AI Trading and the Limits of EU Law Enforcement in Deterring Market Manipulation" (2022) 45 Computer Law & Security Review 105690 at 2–6.

insurance for AI-focused trading firms, and (2) an industry-funded compensation fund for market-wide disruptions. Each can be designed to complement, rather than replace, the proposed AGL cause of action and public enforcement under the FMC Act, filling gaps in victim protection while preserving overall deterrence and accountability.

These proposed mechanisms would also operate alongside the existing civil liability and compensation provisions in Part 8 of the FMC Act. As discussed in Chapter 4, the existing statutory framework provides a route through which compensation may be sought where loss results from a contravention of an applicable civil liability provision.[474] The justification for mandatory insurance and a potential compensation fund is therefore not that the FMC Act provides no avenue for compensation. Rather, these mechanisms are intended to strengthen the practical availability of compensation by addressing matters such as the collectability and timely payment of established claims and, in the case of a carefully defined compensation fund, providing a potential mechanism for responding to widespread losses arising from qualifying systemic AI-related market events.

One way to reinforce the AGL regime is by requiring firms that deploy AI trading systems to carry insurance coverage for losses caused by those systems. This proposal is analogous to the professional indemnity insurance that many financial service providers must maintain, and it aligns with expert suggestions that advanced algorithms be registered and insured (for example, Judge Karnow's notion of a "Turing Registry" for AI systems).[475] Under such a mandate, if an AI-driven trading incident leads to legal liability; for instance, a firm is found negligent under the AGL cause of action for failing to supervise its algorithm, the firm's insurance policy would cover the damages up to the policy limits. It is important to clarify that this does not create a no-fault or strict liability scheme: the injured party would still bear the burden of proving the firm's liability, by showing a breach of the proposed duty of diligent governance and causation of harm, just as they would without insurance.

The insurance simply guarantees that, once liability is established, the victim can recover the award from a solvent source. In other words, insurance ensures that AGL's promise

[474] Financial Markets Conduct Act 2013, ss 485–487 and 494–495.
[475] Curtis EA Karnow "Liability for Distributed Artificial Intelligences" (1996) 11 Berkeley Technology Law Journal 147 at 170–173.

of compensation is not hollowed out by a defendant's insolvency or by protracted enforcement delays. This mechanism directly addresses scenarios where a major AI-caused loss might otherwise force a trading firm into bankruptcy (potentially leaving victims with uncollectable judgments), a particularly salient concern given that AI failures could produce losses both large in magnitude and dispersed among many claimants. By spreading the risk of such catastrophic losses across the insurance market, individual firms are less likely to be financially crippled by a single adverse event, and victims are more likely to be paid promptly.

Requiring AI liability insurance also carries behavioural benefits that reinforce the regulatory framework. Insurers, in pricing and underwriting these policies, will effectively conduct a private oversight role: firms with stronger controls, rigorous testing, and good compliance records should enjoy lower premiums, whereas those with poor governance or a history of close calls may face higher costs or difficulty obtaining coverage. This creates a market-based incentive for firms to invest in robust AI risk management, complementing the direct legal incentive provided by AGL itself. If a firm's management knows that better training of its AI, more diligent human oversight, or independent algorithm audits will reduce its insurance premiums, it has an added financial motivation to take those precautionary steps. At the same time, policymakers and insurers must guard against moral hazard, the tendency for insurance to weaken the insured's vigilance by insulating them from consequences. This can be mitigated by structuring the insurance intelligently: for example, policies can include substantial deductibles or coinsurance ensuring the firm always bears part of any loss out of pocket, exclude coverage for intentional misconduct or flagrant regulatory breaches, and require firms to implement certain minimum safety standards as a condition of coverage.

Overall, a mandatory insurance requirement would operate as a supportive backstop to the AGL cause of action. It does not change the standard of care or make lawsuits any easier or harder to win, but it guarantees that when a firm is held liable after due process under the existing rules, the firm's financial capacity to compensate victims is bolstered by an insurer's funds. This not only protects victims but also benefits the market as a whole: by capping a firm's direct exposure to an AI related catastrophe through insurance payouts, it reduces the risk that firms will be overly deterred from innovating or participating in the market due to fear of an unmanageable, ruinous liability. In short, insurance reinforces AGL by ensuring that accountability leads to actual compensation and by stabilising the risk environment so that conscientious firms are not driven out of the market by single disasters.

A second mechanism to enhance the resilience of the compensation framework is the establishment of an AI Market Disruption Fund, financed by the industry, to provide swift payouts in the wake of major AI-induced market events. This proposal takes inspiration from existing investor protection schemes. For example, Part 6 of the Deposit Takers Act 2023 establishes the Depositor Compensation Scheme, which protects eligible deposits up to a maximum entitlement of $100,000 per depositor per deposit taker and is funded through levies imposed on deposit takers.[476] Similarly, one could envision a capital markets compensation fund that would step in when an algorithmic trading debacle causes widespread damage across the market. All firms engaging in algorithmic or high-frequency trading or perhaps all NZX participants above a certain size might be required to contribute to the fund.

If a qualifying event occurs like a flash crash or a runaway algorithm that severely disrupts pricing and causes losses to many market users, the fund will immediately compensate affected investors for a substantial portion of their losses, without each investor needing to prove a case in court. Such a scheme could be modelled on overseas ideas: Yadav, for instance, has argued for a U.S. "National Market System" fund to cover losses from algorithmic crashes, precisely because traditional legal remedies are too slow and uncertain in these scenarios.[477] The key is to balance quick relief with accountability. Participation rules would likely limit payouts to truly significant, systemic incidents (to avoid depleting the fund for minor issues) and impose payout caps per incident or per investor (for example, covering a fixed percentage of each loss up to a maximum amount). These caps ensure the fund remains solvent and that claimants still have incentive to pursue additional recovery through normal legal channels if their losses are only partially covered.

Most importantly, the fund must work in tandem with the proposed liability framework by giving the fund strong subrogation rights. In practice, after the fund compensates victims, it would acquire those victims' rights to sue any parties responsible for the harm. The fund or the regulator managing it can then bring an action. For example, an AGL claim or an FMC Act proceeding against the at-fault firm or individuals to recover the money that was paid out. Subrogation ensures that the ultimate cost of the event is borne by the wrongdoers, not by the industry collectively. It also preserves deterrence: a firm cannot escape liability simply because

[476] Deposit Takers Act 2023, pt 6, ss 190, 196, 202–203 and 235–238.
[477] Yadav, above n 193, at 1093–1095.

a fund covered its victims' losses, since the fund or its administrator will seek reimbursement from that firm.

In essence, the fund would provide immediate liquidity and certainty to victims, while the legal system, through subrogation, continues to adjudicate fault and assign financial responsibility. This approach mirrors how other compensation schemes maintain discipline. A comparable principle operates under the Depositor Compensation Scheme, under which statutory subrogation provisions allow the scheme administrator to assume relevant rights following compensation.[478] The significance of this analogy is that a compensation mechanism can provide prompt relief while preserving the ability to pursue recovery from responsible parties. Likewise, an AI disruption fund could compensate eligible investors promptly while retaining statutory rights of recovery against persons legally responsible for the relevant event.

Both mandatory insurance and a compensation fund offer valuable ways to support the AGL framework, but they address different facets of the risk. Insurance works on the individual firm level to ensure that each liable party can make good on its obligations, directly reinforcing AGL's deterrent and compensatory functions. A compensation fund operates at the systemic level to address multi-firm or diffuse harms; it can boost market confidence by guaranteeing that even in a complex crisis, victims will not be left waiting for relief. However, a fund would require significant setup and oversight, and it might be less crucial until algorithmic trading is more ubiquitous and capable of causing truly market-wide events.

In the New Zealand context, mandatory liability insurance emerges as the more immediately practical and closely aligned tool to support the proposed AGL framework. By directly tying into the AGL cause of action's outcomes, insurance bolsters the enforceability of that tort and ensures that firms are not unduly deterred from using AI due to fears of unbounded liability, they can manage their risk through insurance while still being held to account for any negligence. Insurance also resonates with New Zealand's focus on proportionate, adaptive regulation: it leverages market mechanisms like the insurance industry's risk assessment to improve AI oversight without heavy-handed new bureaucracy.

[478] Deposit Takers Act 2023, s 227.

An industry fund could be introduced later if AI-driven trading becomes so widespread that a collective safety net is justified for rare but significant disasters. For now, focusing on mandatory insurance, overseen by regulators to ensure policies are adequate and exclusions are fair, is likely the best fit. This approach neatly aligns with New Zealand's balanced, light-touch regulatory philosophy: it fortifies the legal framework including the proposed AGL framework and FMC Act duties and encourages responsible innovation by mitigating financial uncertainties for both firms and investors.

The feasibility of implementing mandatory liability insurance for AI trading firms will ultimately depend on how the insurance market evolves in response to the growing complexity and scale of AI deployment. While the concept of mandatory insurance aligns well with New Zealand's regulatory philosophy, the current global insurance landscape for AI-related risks remains in a developmental stage. Internationally, insurers are beginning to develop and pilot specialised products that address the unique liabilities posed by AI systems, including algorithmic trading. For example, Lloyd's of London has backed policies developed by Armilla AI that cover losses resulting from AI malfunctions, such as chatbot errors or performance degradation. These policies are underwritten by syndicates like Chaucer and are designed to respond when an AI system performs below its expected accuracy.[479]

These developments suggest that the insurance industry is increasingly viewing AI risk as quantifiable and insurable, albeit with caution. Underwriting practices are evolving to include rigorous assessments of AI governance, including documentation of model testing, bias audits, and oversight mechanisms. This trend indicates that the insurance market is not only responding to regulatory and litigation pressures but is also shaping AI governance through its own risk assessment criteria.

However, challenges remain. Pricing AI risk is inherently difficult due to the lack of historical claims data and the dynamic nature of AI systems, which can change behaviour over time through retraining or adaptation. As a result, firms may find that their current insurance arrangements do not adequately cover AI-specific liabilities. This has led to calls from brokers and industry analysts for AI to be treated as a distinct risk class, similar to the evolution of

[479] Lockton Re and Armilla AI *Ready or Not: The Impact of Artificial Intelligence on Insurance Risks* (Lockton Re, 4 February 2026).

cyber insurance.[480] The emergence of standalone "affirmative AI" policies is a step in this direction, but widespread availability and affordability remain limited.

In the New Zealand context, no dedicated algorithmic trading liability insurance product currently exists. Firms seeking coverage would likely need to access bespoke policies through international brokers or participate in pilot schemes offered by global insurers.[481] Therefore, while mandatory insurance is conceptually sound and aligns with the goals of the proposed AGL framework, its practical implementation may need to be phased or conditional. Regulators could consider a transitional approach, encouraging voluntary uptake of AI liability insurance in the short term while monitoring market developments and engaging with insurers to build capacity. Over time, as the insurance market matures and more tailored products become available, a mandatory scheme could be introduced with greater confidence in its feasibility and effectiveness.

7.12. CONCLUSION

The regulatory framework proposed in this chapter is designed to embed accountability at every stage of the AI trading lifecycle, from design and deployment through to real-time operation and post-incident redress. At its core, the framework rests on three interlocking pillars. The first is proactive compliance mechanisms, including pre-market approval, explainability standards, and embedded compliance by design features. The second is robust oversight and enforcement, including continuous monitoring, algorithm inventories, and a proposed statutory cause of action for AI Governance Liability (AGL), designed to hold firms accountable for systemic supervisory and governance failures. The third is a layered compensation architecture, combining mandatory liability insurance and, potentially, an industry-funded disruption fund to ensure that victims of AI-related harm are made whole.

These remedies are not mutually exclusive. Rather, they are designed to interact and co-exist, each addressing a distinct doctrinal and regulatory gap identified in earlier chapters.

---

[480] Daniel Wood "AI is the next cyber: Why insurers may carve out standalone cover" *Insurance Business* (24 February 2026) <https://www.insurancebusinessmag.com/au/news/technology/ai-is-the-next-cyber-why-insurers-may-carve-out-standalone-cover-566199.aspx>.

[481] Lockton Re and Armilla AI, above n 479.

The compliance regime, developed in Sections 7.3 and 7.4, responds to the challenges of AI-driven market manipulation discussed in Chapter 4 by adapting the FMC Act market conduct provisions to the realities of autonomous and opaque trading systems. It establishes ex ante obligations such as pre-market approval, conformity assessments, and embedded safeguards to prevent harm before it occurs.

These measures principally perform preventive and public enforcement functions. However, the existing FMC Act also provides a statutory route to compensation through the civil liability provisions in Part 8, including compensatory orders sought by the FMA or affected persons. The additional mechanisms proposed in this chapter are therefore intended to address residual gaps, including defective products or systemic governance failures that do not fall readily within an existing civil liability provision, and practical difficulties concerning proof, causation and recovery. The proposed NZX AI Regulatory Office (NZX AI-RO) would serve as the institutional cornerstone of this regime. It would be responsible for administering pre-market conformity assessments, issuing technical guidance, overseeing compliance audits, and facilitating stakeholder engagement. By consolidating technical and regulatory expertise, the NZX AI-RO would ensure that AI trading systems are not only compliant with legal standards but also aligned with evolving best practices in safety, transparency, and accountability.

To address the limitations of existing private law remedies, the chapter proposes two complementary liability mechanisms. The first is the FMC Act-based product liability regime, set out in Section 7.10, which addresses limitations in the protection available under existing consumer and financial markets legislation. As discussed in Chapter 6, the application of the Consumer Guarantees Act may be limited where AI trading systems are acquired or used for commercial or productive purposes, while the Fair Trading Act may address misleading or deceptive conduct but does not establish a general strict product liability regime for defective AI trading systems. The proposed regime would impose strict, no-fault liability on providers of AI financial products for harm caused by defects, subject to the requirements and limitations of the proposed statutory framework.

The proposed regime would create a statutory cause of action imposing strict liability for harm caused by defective AI financial products. Its statutory foundation would define the relevant standard of defectiveness, the persons subject to liability, the categories of recoverable

loss, applicable evidential presumptions and available defences. This approach would preserve doctrinal coherence by avoiding the need to extend common law negligence beyond its established limits, while providing a specifically designed remedy for defective AI financial products. As discussed in Chapter 6, this is particularly important where conventional negligence principles create difficulties concerning duties of care, pure economic loss, causation and the opacity of AI decision-making.

The second mechanism is the proposed AGL statutory cause of action, introduced in Section 7.9, which addresses supervisory and governance failures that negligence law may be ill-equipped to capture. As explained in Chapter 6, negligence is constrained by the difficulties associated with recognising duties of care for pure economic loss, particularly where claimants are diffuse market participants and there is no sufficiently proximate relationship or assumption of responsibility. AGL would respond to these limitations by establishing a statutory duty of diligent governance directed specifically at systemic failures in the supervision and control of AI trading systems. It would remain fault-based, with liability depending on whether the firm failed to meet the required standard of reasonable governance and whether that failure caused recoverable harm. AGL would therefore complement the existing FMC Act civil liability framework by addressing a distinct category of organisational and supervisory failure that may not readily fall within an existing civil liability provision.

Together, these three proposed mechanisms, comprising enhanced regulatory compliance, product liability and governance liability, would supplement the existing FMC Act framework to create a more comprehensive and proportionate framework for governing AI in New Zealand's financial markets. Their interaction would create complementary incentives and remedies supporting responsible AI innovation. The compliance regime would establish minimum standards for safe AI deployment and operation, reducing the likelihood of defects and governance failures. Product liability and AGL would, in turn, provide distinct ex post avenues for redress. Product liability would address defects in the AI system itself, while AGL would address failures in the human governance and supervisory systems surrounding its deployment and operation. Depending on the circumstances, the same incident could potentially engage more than one part of the framework.

This layered approach would reduce, rather than eliminate, gaps in accountability and compensation. It would create additional incentives for firms to invest in robust design, testing,

monitoring and oversight, while operating alongside the existing enforcement and compensation mechanisms under the FMC Act. It also draws on relevant international developments, particularly European approaches to AI governance, algorithmic trading regulation and product liability, while adapting those approaches to the New Zealand financial markets context. Doctrinal coherence would be maintained by treating regulatory compliance as public regulation, the proposed product liability regime as a statutory strict liability cause of action, and AGL as a separate fault-based statutory cause of action directed at governance and supervisory failures.

Ultimately, this framework does not seek to stifle innovation. On the contrary, it is designed to enable responsible AI adoption by providing clarity, certainty, and confidence to market participants. By embedding compliance into the architecture of AI systems, aligning incentives through liability and insurance, and improving both the prevention of harm and the availability of compensation when harm occurs, the framework seeks to support responsible innovation rather than treating regulation solely as a constraint. It offers a forward-looking and principled response to the risks and realities of AI in modern financial markets.

Chapter 8: Overall Conclusion

"By far, the greatest danger of Artificial Intelligence is that people conclude too early that they understand it." – Eliezer Yudkowsky[482]

Financial markets in New Zealand, as elsewhere, have undergone a fundamental transformation. Trading activity that was once conducted by human actors exercising judgement in real time is now dominated by automated, algorithmic, and increasingly artificial intelligence-driven systems operating at extraordinary speed and scale. Decisions that once took minutes or seconds are now executed in microseconds by machines that learn from data, adapt their strategies, and interact continuously with other automated systems. This transformation has delivered significant benefits, including increased liquidity, narrower spreads, and improved execution efficiency. At the same time, it has exposed structural vulnerabilities in the legal framework governing market integrity, particularly in relation to market manipulation.

This thesis has been concerned with a central question: whether New Zealand's legal system, and in particular the market manipulation regime under the Financial Markets Conduct Act 2013, is capable of effectively regulating AI-driven market manipulation in the New Zealand stock exchange market, and if not, how it should evolve. The answer developed through this research is that New Zealand law is not conceptually incapable of responding to AI-driven manipulation, but it is structurally strained by assumptions embedded within it about human agency, intent, and observable conduct. Without doctrinal recalibration and governance-focused reform, those assumptions risk allowing technologically mediated misconduct to evade effective regulation.

The starting point for this enquiry lies in the nature of the risks introduced by automation and artificial intelligence. Market manipulation has always been concerned with conduct that creates false or misleading appearances of trading activity, supply, demand, or price. The core harm has remained constant: distortion of price discovery, erosion of investor

[482] Eliezer Yudkowsky "Artificial Intelligence as a Positive and Negative Factor in Global Risk" in Nick Bostrom and Milan M Ćirković (eds) *Global Catastrophic Risks* (Oxford University Press, Oxford, 2008) 308 at 308.

confidence, and undermining of market integrity. What has changed is the mechanism through which those harms can occur. Algorithmic and AI-driven systems can generate manipulative effects without discrete human acts that resemble traditional misconduct. Patterns of distortion may emerge from system architecture, optimisation objectives, training data, and interactions between multiple automated agents rather than from a single deceptive trade or statement.

Speed is a defining feature of this transformation. High-frequency and AI-enabled trading systems operate at temporal scales that exceed human comprehension, let alone intervention. At such speeds, market impact can occur before compliance teams, regulators, or even system designers are aware that a problem has arisen. Scale compounds this risk. Automated systems can place and cancel thousands of orders per second across multiple venues, amplifying even minor design flaws into market-wide disruptions. Opacity further intensifies the challenge. Machine learning models, particularly those based on deep learning, often lack transparent decision pathways, making it difficult to explain why a system behaved as it did or to reconstruct causal chains after harm has occurred. Finally, emergence introduces a qualitative shift. Harmful outcomes may arise not from any single rule or parameter, but from the interaction of multiple systems responding to each other in complex and unpredictable ways.

These features do not create new categories of harm, but they fundamentally alter how familiar harms arise and how they must be regulated. False liquidity can be created through rapid order placement and cancellation. Prices can be influenced through feedback-driven momentum strategies. Information asymmetries can be deepened through co-location, privileged data feeds, and superior latency. AI systems can autonomously generate trading signals, market commentary, or sentiment analysis that influence market behaviour at scale. In each case, the resulting market distortion may be indistinguishable in economic effect from traditional manipulation, yet radically different in form.

Against this backdrop, the Financial Markets Conduct Act 2013 represents both a strength and a constraint. The Act was designed to promote confident and informed participation in New Zealand's financial markets by prohibiting conduct that undermines genuine price formation. Sections 262 and 265 contain the principal market manipulation prohibitions, addressing false or misleading statements or information and conduct that creates, or is likely to create, a false or misleading appearance of trading. Sections 266 to 268

supplement those prohibitions by providing an exception for specified short selling and crossings, deemed contraventions involving particular trading conduct, and a defence based on accepted market practices and proper purpose. This approach has provided clarity and certainty, and it has enabled effective enforcement against a wide range of human-driven misconduct. However, the enumerated structure of the Act also reflects an implicit assumption that manipulation will resemble recognisable human behaviour.

Sections 262 and 265 are framed around concepts of statements, information, trades, and orders that can be linked to human actors exercising contemporaneous control. Whilst the Act does not expressly require proof of subjective intent, its language presupposes a form of agency that becomes increasingly diffuse in AI-driven markets. Knowledge is distributed across developers, traders, compliance officers, and corporate governance structures. Decisions are embedded in code rather than made in the moment. Harm may arise through optimisation rather than deception.

This thesis has demonstrated that the existing market manipulation framework under the FMC Act is not inherently obsolete but is increasingly misaligned with the operational realities of AI-driven trading. Whilst sections 262, 265, 267, and 268 have been interpreted flexibly by New Zealand courts and regulators, their underlying structure continues to assume manipulation will take the form of discrete human acts, identifiable decision points, and observable deception. This assumption creates pressure points when manipulation arises instead through system design, parameter setting, optimisation logic, and emergent algorithmic behaviour rather than through a single misleading trade or statement.

Definitional ambiguity is a key source of regulatory strain. The absence of an overarching definition of market manipulation, combined with the statutory emphasis on “false or misleading appearances”, risks directing attention away from economically distortive conduct that operates through amplification of genuine signals or exploitation of market microstructure. Many AI-driven strategies undermine price formation without producing information that is obviously false or trades that are individually deceptive. As a result, harmful conduct may evade enforcement not because it is consistent with market integrity, but because it does not resemble the forms of manipulation contemplated when the provisions were drafted.

Attribution of responsibility, rather than proof of harm, is the principal doctrinal challenge posed by AI-driven manipulation. Whilst the “knows or ought reasonably to know” standard offers conceptual flexibility; its application becomes increasingly complex in contexts of distributed agency and algorithmic opacity. Responsibility in AI-mediated markets is often upstream and structural, residing in choices about system architecture, data inputs, risk controls, and monitoring rather than in moment-to-moment trading decisions. Without explicit recognition of these governance dimensions, liability risks being either under inclusive, allowing foreseeable system-level misconduct to escape regulation, or over inclusive, discouraging legitimate algorithmic trading through uncertainty.

The thesis does not argue for wholesale replacement of the existing regime. Instead, it demonstrates how targeted doctrinal recalibration can preserve the core objectives of the FMC Act whilst restoring regulatory coherence. Proposed amendments to sections 262, 265, 266, 267, and 268 exemplify a governance-oriented approach that aligns liability with control and foreseeability at the system level. By explicitly recognising automated trading systems as vectors of manipulation, extending responsibility to those who design, deploy, and oversee them, and refining defences to reflect accepted algorithmic practices, the framework can remain both effective and proportionate. Under this approach, liability does not depend on predicting each algorithmic output. Instead, it turns on whether misleading outcomes were reasonably foreseeable given the system's configuration and whether reasonable steps were taken to manage those risks. Responsibility tracks control rather than proximity to execution. This preserves space for legitimate algorithmic trading whilst ensuring that complexity cannot be used as a shield.

A core contribution of this research lies in demonstrating that existing case law already contains the doctrinal tools necessary to address AI-driven market manipulation, even in the absence of cases involving fully autonomous systems. Through detailed analysis of New Zealand and international authorities, this thesis has shown that courts have consistently prioritised market effect, objective knowledge, and responsibility for deployment over subjective intent or the mechanics of execution. This jurisprudence provides a strong foundation for regulating technologically mediated trading without requiring radical doctrinal innovation.

The analysis reveals a clear and consistent judicial approach. Courts have repeatedly held that market manipulation may arise from orders alone, even where no trade is executed, that small or uneconomic trades can distort market integrity, and that liability does not depend on proof of malicious intent. Instead, responsibility turns on whether a trader knew or ought reasonably to have known that their conduct was likely to create a false or misleading appearance of trading activity, supply, demand, or price. These principles are inherently adaptable to AI-driven markets, where manipulation often manifests through patterns, speed, and scale rather than discrete human decisions.

Courts are willing to infer knowledge and purpose from trading patterns, surrounding circumstances, and the absence of genuine commercial justification. This evidentiary approach is particularly significant in the AI context, where direct evidence of intent may be unavailable and decision-making is embedded in code. By relying on constructive knowledge and objective assessment, the case law avoids the need to attribute mental states to machines whilst still holding human actors accountable for foreseeable market effects arising from systems they deploy or control.

International jurisprudence reinforces this conclusion. The authorities examined in this thesis treat algorithmic trading systems as tools through which identifiable human or corporate actors act, rather than as independent legal actors. Depending on the applicable statutory framework, liability may turn on the intent embedded in the design or deployment of the algorithm, or on the objective market effect of the trading conduct together with the responsibility of the relevant human or corporate actor. The so-called “black box” problem of AI is therefore primarily an evidentiary challenge rather than a doctrinal one. Technological complexity or automation does not, by itself, displace responsibility. Instead, responsibility remains with those who design, deploy, control, or oversee trading systems where the applicable legal requirements for liability are established.

The analysis in this thesis demonstrates that New Zealand law already possesses several of the conceptual tools required to address AI-driven market manipulation. By focusing on market impact, constructive knowledge, and responsibility for the deployment and oversight of trading systems, courts can regulate AI-driven trading without anthropomorphising algorithms or requiring proof of machine intent. These principles are particularly well suited to AI-driven markets, where responsibility must remain attributable to human or corporate

actors. Under s 265 of the FMC Act, the relevant inquiry is whether the person knew or ought reasonably to have known that the conduct would, or was likely to, create the prohibited false or misleading appearance. In the AI context, foreseeable risks arising from system design, deployment, and oversight may therefore be relevant to determining whether that constructive knowledge standard is satisfied.

Public enforcement is only part of the regulatory response. One contribution of this thesis lies in its emphasis on private law and complementary compensation mechanisms in AI-driven markets. Part 8 of the Financial Markets Conduct Act 2013 already provides a statutory route to compensation: ss 485–487 establish the relevant civil liability and declaration framework, while ss 494–495 permit the court to make compensatory orders where a contravention has caused, or is likely to cause, loss or damage. Pecuniary and criminal sanctions do not themselves compensate affected persons, but the statutory compensation route means that it would be inaccurate to suggest that market manipulation law is solely punitive. The private law and compensation reforms proposed in this thesis should therefore supplement, rather than replace, the existing statutory route, particularly where attribution, causation, access to evidence, or market-wide loss makes that route difficult to use in practice.

From this analysis emerges the normative core of the thesis. Artificial intelligence does not necessitate a complete rethinking of market regulation, but it does demand a new architecture of governance. Traditional enforcement models, which rely on retrospective investigation and response, are ill-suited to the realities of AI-driven markets, where harm can materialise in microseconds and propagate across systems before detection is possible. What is required is a regulatory framework that anticipates risk, maintains continuous oversight, ensures accountability, and provides effective remedies when harm occurs.

This thesis proposes a governance-based framework that integrates public and private law mechanisms with institutional oversight. It builds on existing legal foundations and regulatory models in New Zealand, adapting them to the specific challenges posed by AI in financial markets. The proposed compliance regime includes pre-market approval for high-risk AI trading systems, explainability standards, and embedded compliance features. These measures are designed to ensure that AI systems are safe, transparent, and aligned with market integrity before they are deployed.

The framework is tailored to New Zealand's market structure and regulatory culture. It avoids overregulation by adopting a proportionate, risk-based approach that leverages existing institutions and regulatory expertise. Rather than replicating large-scale foreign models, it draws on domestic precedents such as the AML/CFT regime and the NZX Rules to create a system that is both practical and enforceable.

Central to this approach is the principle that AI systems must be explainable and auditable. Firms should be able to demonstrate how their systems function, what data they rely on, and how they mitigate foreseeable risks. This includes maintaining documentation, audit trails, and internal controls that allow for meaningful regulatory scrutiny. Systems that cannot be explained or monitored should not be permitted to operate in critical market functions.

The framework also emphasises the importance of embedding compliance into the design of AI systems. Rather than treating regulation as an external constraint, firms should integrate safeguards into the architecture of their algorithms. This includes automated controls to prevent manipulative behaviour, real-time monitoring tools, and escalation protocols that enable human intervention when anomalies arise. Responsibility for these systems must be clearly allocated within firms, with senior management and boards held accountable for governance failures.

Furthermore, the research has demonstrated that private law has an important but carefully bounded role to play in responding to harm caused by AI-driven trading systems. Negligence is constrained where harm takes the form of pure economic loss, particularly by requirements of proximity and the courts' reluctance to impose duties of care in the absence of a confined relational setting or an assumption of responsibility. Product liability under the Consumer Guarantees Act 1993 is also limited in this context. The CGA is unlikely to apply to AI trading systems used for active trading and wealth generation where the purchaser does not fall within the statutory definition of a consumer. The Fair Trading Act 1986 provides a separate avenue where misleading or deceptive conduct or false or misleading representations are involved, but it does not provide a comprehensive product liability framework for defects in the design or performance of AI trading systems.

These limitations are particularly significant where harm arises from systemic failures in design, deployment, or oversight, or is suffered by market participants who did not acquire

the AI system themselves. Existing private law may therefore leave some forms of AI-related financial harm without an effective remedy, particularly where the loss is market-wide or cannot be linked to misleading conduct, a consumer transaction, or a sufficiently proximate relationship. This gap in redress underscores the need for targeted legal reform to ensure that accountability mechanisms remain effective in the face of technological change.

To address these gaps, this thesis proposes two targeted legal reforms designed to extend accountability and redress beyond the limitations of existing private law. The first is a statutory product liability regime under the Financial Markets Conduct Act. This regime would impose strict liability on providers, developers, and distributors of AI financial products for harm caused by defects, irrespective of fault. It would treat AI trading systems as financial products and establish a new statutory right to be free from harm caused by defective algorithmic systems. This approach is particularly important for retail investors and other market participants who fall outside the scope of the CGA and the FTA, and who are currently left without effective remedies when AI systems malfunction or behave unpredictably. By focusing on the safety and reliability of the product itself, rather than the conduct of the provider, this regime would enable victims to seek compensation without needing to prove negligence or misrepresentation, although defect, causation, qualifying loss, and the statutory scope of the regime would remain relevant.

The second reform is the introduction of the tort of AI Governance Liability, grounded in a statutory duty of diligent governance. This novel cause of action is designed to address systemic supervisory failures in the oversight of AI trading systems. It provides a fault-based remedy where harm arises not from a technical defect in the AI system, but from organisational shortcomings such as inadequate monitoring, poor risk controls, or failure to respond to known risks. The AGL framework shifts the focus from the behaviour of the AI system to the quality of the governance structures surrounding it, holding firms accountable for failing to exercise reasonable oversight over high-risk technologies. These reforms are not intended to displace existing remedies but to complement them. Together, they form part of a broader governance-based framework that integrates public enforcement, private law, and institutional oversight. This integrated approach ensures that accountability and compensation remain viable even in the face of technological complexity, distributed agency, and the evolving nature of AI in financial markets.

The framework also includes proposals for compensation mechanisms that reflect the scale and complexity of AI-related harm. These include mandatory insurance requirements and the potential establishment of a market disruption fund to provide timely compensation in the event of widespread harm. These measures ensure that victims are not left without remedy, even where litigation is impractical or enforcement is delayed.

This approach redefines responsibility in AI-driven markets. It recognises that AI systems are tools, not actors. Harm does not result from machine intent, but from human decisions about how systems are designed, trained, and deployed. Legal responsibility must therefore remain with those who control and benefit from these systems. Courts need not engage in speculative inquiries into algorithmic intent. Instead, they can assess whether firms took reasonable steps to prevent foreseeable harm, based on the quality of their governance and oversight.

The proposed framework delivers four essential outcomes for market integrity in the age of AI. It promotes prevention by requiring firms to demonstrate safety and compliance before deployment. It ensures control through continuous monitoring and intervention mechanisms. It reinforces accountability by holding firms and their leadership responsible for governance failures. And it guarantees restoration by providing victims with access to compensation when harm occurs.

This is not a call to resist innovation. Rather, it is a call to ensure that innovation does not come at the expense of trust. Investors do not need to understand the technical workings of AI to participate in financial markets, but they must be confident that the systems in use are safe, that misconduct will be detected and addressed, and that the rules of the market remain meaningful. That confidence is essential to the functioning of the market and must be preserved.

This research is not without limitations. It is primarily doctrinal in nature and does not include empirical analysis of trading data or direct observation of AI systems in operation. The findings are based on legal sources, regulatory materials, and reported enforcement actions, rather than on quantitative evidence. The field of AI is also rapidly evolving, and the risks identified in this thesis may change as technologies develop. The proposed framework is

designed to be adaptable, with an emphasis on continuous monitoring and regulatory flexibility.

The thesis is focused on the New Zealand context, and while the proposals are tailored to local institutions and market characteristics, financial markets are increasingly global. Cross-border coordination and regulatory cooperation will be essential to address the international dimensions of AI-driven trading. The thesis also does not explore broader ethical or distributive concerns, such as the impact of AI on market fairness or the burden of compliance on smaller firms. These issues warrant further investigation.

Future research could explore the empirical dimensions of AI trading in New Zealand, including the prevalence of potentially manipulative strategies and the effectiveness of current surveillance tools. Comparative studies could examine how other jurisdictions are adapting their legal frameworks to address similar challenges. Further work is also needed on evidentiary standards in litigation involving opaque technologies, and on the allocation of proof burdens where claimants lack access to technical information. Longitudinal research could assess whether governance frameworks remain effective as AI capabilities continue to evolve.

Despite these limitations, this thesis makes an original and meaningful contribution to the field. It demonstrates that New Zealand already possesses the legal foundations necessary to regulate AI-driven market manipulation, provided that responsibility is re-oriented from narrow conceptions of intent to broader principles of governance and control. It shows how public enforcement and private law can work together to prevent harm, ensure accountability, and provide redress. It proposes a governance framework that is both principled and practical, grounded in existing institutions and tailored to the realities of AI in financial markets.

The law need not stand still in the face of technological transformation. By recognising that AI is a tool whose risks are foreseeable and governable, by aligning responsibility with control, and by embedding governance into system design, New Zealand can ensure that as machines execute, humans remain accountable. The research presented in this thesis offers the legal and institutional foundations necessary to achieve this outcome. The challenge is urgent, but the tools are available. What remains is the will to act before delay renders effective regulation more difficult and the costs of inaction more severe.

Bibliography

A. Legislation, Codes and Regulations

New Zealand

Contract and Commercial Law Act 2017.

Consumer Guarantees Act 1993.

Deposit Takers Act 2023.

Fair Trading Act 1986.

Financial Markets Conduct Act 2013.

Australia

Corporations Act 2001 (Cth).

European Union

Commission Delegated Regulation (EU) 2017/566 of 18 May 2016 supplementing Directive 2014/65/EU [2017] OJ L 87/84.

Commission Delegated Regulation (EU) 2017/589 of 19 July 2016 supplementing Directive 2014/65/EU [2017] OJ L 87/417.

Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments [2014] OJ L 173/349 (MiFID II).

Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products [2024] OJ L.

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L 1689/1.

France

Code monétaire et financier (France).

Germany

Gesetz zum Hochfrequenzhandel (High Frequency Trading Act).

Wertpapierhandelsgesetz (German Securities Trading Act).

Italy

Decree of the Minister of Economy and Finance (Italy) 21 February 2013.

Italian Financial Transaction Tax, Law No 228 of 24 December 2012.

International

ISO 10218-1:2025 Robotics — Safety requirements — Part 1: Industrial robots.

ISO 10218-2:2025 Robotics — Safety requirements — Part 2: Industrial robot applications and robot cells.

United States

21 USC § 360e (Premarket Approval).

Securities Exchange Act of 1934, 15 USC § 78c(a)(1).

Commodity Futures Trading Commission Regulation 17 CFR § 166.3.

B. Cases

New Zealand

Attorney-General v Carter [2003] 2 NZLR 160 (CA).

Body Corporate Number DPS 91535 v 3A Composites GmbH [2023] NZCA 647.

Body Corporate No 286562 v Stanley [2012] NZHC 2860.

Financial Markets Authority v Brian Peter Henry [2014] NZHC 1853.

Financial Markets Authority v Matthew Geoffrey Hill [2024] NZHC 1353.

Financial Markets Authority v Kok Ding Cheng [2025] NZHC 1017.

Financial Markets Authority v Warminger (No 3) [2017] NZHC 327.

Financial Markets Authority v Wei Zhong [2022] NZHC 480.

Financial Markets Authority v Wei Zhong [2023] NZHC 766.

Grange v Attorney-General [2012] 3 NZLR 53.

Kaori Ltd v Shrinkforce Shrink Wrap Services Ltd (in rec) [2012] NZHC 3204.

North Shore City Council v Attorney-General [2012] NZSC 49, [2012] 3 NZLR 341.

Nesbit v Porter [2000] 2 NZLR 465 (CA).

R v Hansen [2007] NZSC 7, [2007] 3 NZLR 1.

Scott Group Ltd v McFarlane [1978] 1 NZLR 553 (CA).

Sleight v Beckia Holdings Ltd [2020] NZHC 2851.

Smith v Fonterra Co-operative Group Ltd [2024] NZSC 5.

South Pacific Manufacturing Co Ltd v New Zealand Security Consultants & Investigations Ltd [1992] 2 NZLR 282 (CA).

Australia

Mallonland Pty Ltd v Advanta Seeds Pty Ltd [2024] HCA 25.

Canada

1688782 Ontario Inc v Maple Leaf Foods Inc 2020 SCC 35, [2020] 3 SCR 504.

Deloitte & Touche v Livent Inc (Receiver of) 2017 SCC 63, [2017] 2 SCR 855.

United Kingdom

Anns v Merton London Borough Council [1978] AC 728 (HL).

Ayinde v London Borough of Haringey [2025] EWHC 1383 (Admin).

Al-Haroun v Qatar National Bank [2025] EWHC 1383 (Admin).

Caparo Industries plc v Dickman [1990] 2 AC 605 (HL).

Donoghue v Stevenson [1932] AC 562 (HL).

Drummond v Van Ingen (1887) 12 App Cas 284 (HL).

Fairchild v Glenhaven Funeral Services Ltd [2002] UKHL 22, [2003] 1 AC 32.

Financial Conduct Authority v Da Vinci Invest Ltd [2015] EWHC 2401 (Ch).

Hedley Byrne & Co Ltd v Heller & Partners Ltd [1964] AC 465 (HL).

Sarao v Government of the United States of America [2016] EWHC 2737 (Admin).

Spartan Steel & Alloys Ltd v Martin & Co (Contractors) Ltd [1973] QB 27 (CA).

European Union

Boston Scientific Medizintechnik GmbH v AOK Sachsen-Anhalt (C-503/13) EU:C: 2015:148, [2015] 3 CMLR 6.

United States

Commodity Futures Trading Commission v Thakkar No 1:18-cv-00619 (ND Ill, 26 January 2018).

Department of Market Regulation v Lek Securities Corporation and Samuel Frederik Lek FINRA Disciplinary Proceeding No 20110297130-04.

Garcia v Character Technologies Inc No 6:24-cv-01903 (MD Fla, 21 May 2025).

In re Panther Energy Trading LLC and Michael J Coscia CFTC Docket No 13-26 (22 July

2013).

In the Matter of Athena Capital Research LLC SEC Administrative Proceeding File No 3-16199 (16 October 2014).

Joiner v Mitchell County Hospital Authority 186 SE 2d 307 (Ga Ct App 1971).

Kociemba v GD Searle & Co 683 F Supp 1579 (D Minn 1988).

Lemmon v Snap Inc 995 F 3d 1085 (9th Cir 2021).

Mata v Avianca Inc No 1:22-cv-01461 (SDNY, 22 June 2023).

United States v Sarao No 1:15-cr-00075 (ND Ill, 28 January 2020).

U.S. Commodity Futures Trading Commission v Nav Sarao Futures Ltd PLC No 1:15-cv-03398 (ND Ill, 17 November 2016) (Consent Order).

C. Official Publications and Regulatory Reports

Agenzia delle Entrate (Italy) *Financial Transaction Tax: High-Frequency Trading — Explanatory Notes* (Agenzia delle Entrate, Rome, 2013).

Artificial Intelligence Researchers Association *Aotearoa New Zealand Artificial Intelligence: A Strategic Approach* (November 2021).

Australian Securities and Investments Commission *ASIC Supervision of Markets and Participants: July to December 2014* (Report 425, 2015).

Australian Securities and Investments Commission *Regulatory Guide 241: Electronic Trading* (August 2022).

Australian Securities and Investments Commission *High-Frequency Trading in Australian Equities and the Australian–US Dollar Cross Rate* (Report 597, November 2018).

Australian Securities and Investments Commission "Gabriel Govinda Sentenced to Two and a Half Years' Imprisonment and Fined for Market Manipulation and Finfluencer Conduct" (Media Release 23-116, 5 May 2023).

Commerce Commission *The Fair Trading Act and New Technologies* (Discussion Paper, March 2020).

Commissione Nazionale per le Società e la Borsa (CONSOB) *Algotrading and HFT: Mapping of the Regulatory Framework and Reporting Obligations to Consob* (CONSOB, 2023).

European Commission *White Paper on Artificial Intelligence: A European Approach to Excellence and Trust* COM (2020) 65 final (19 February 2020).

European Commission *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final (21 April 2021).

European Securities and Markets Authority *Guidelines on Systems and Controls in a Highly Automated Trading Environment for Trading Platforms, Investment Firms and Competent Authorities* (Consultation Paper, 2011).

European Securities and Markets Authority *Supervisory Briefing on Algorithmic Trading in the EU* ESMA74-1505669079-10311 (26 February 2026).

Financial Conduct Authority *Final Notice: Michael Coscia* (3 July 2013).

Financial Markets Authority *Market Manipulation* (Information Page, FMA, Wellington, 2020).

Financial Markets Authority *Insider Conduct: Information Sheet* (FMA, Wellington, March 2021).

Financial Markets Authority *Market Cleanliness in the New Zealand Equity Market* (Occasional Paper Series, October 2024).

Financial Stability Board *Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications* (1 November 2017).

FINRA "FINRA Sanctions Trillium Brokerage Services LLC" (Press Release, 13 September 2010).

International Organization of Securities Commissions Technical Committee *Investigating and Prosecuting Market Manipulation* (IOSCO, May 2000).

International Organization of Securities Commissions *Regulatory Issues Raised by the Impact of Technological Changes on Market Integrity and Efficiency* (Final Report, October 2011).

Lockton Re and Armilla AI *Ready or Not: The Impact of Artificial Intelligence on Insurance Risks* (Lockton Re, 4 February 2026).

Morgan Lewis *Algorithmic Trading: Hong Kong Monetary Authority Discusses Supervisory Expectations and Evolving Market Manipulation Risk* (Client Alert, March 2020).

National Futures Association *Compliance Rule 2-9 and Interpretive Notice No 9060* (NFA, 2025).

NZX *Algorithmic Trading: Thematic Review* (NZX, Wellington, December 2018).

OECD *OECD Handbook on Competition Policy in the Digital Age* (OECD Publishing, Paris, 2022).

Securities and Futures Commission (Hong Kong) *Code of Conduct for Persons Licensed by or Registered with the Securities and Futures Commission* (March 2023).

US Commodity Futures Trading Commission and US Securities and Exchange Commission *Findings Regarding the Market Events of May 6, 2010: Report of the Staffs of the CFTC and SEC to the Joint Advisory Committee on Emerging Regulatory Issues* (30 September 2010).

US Securities and Exchange Commission *Concept Release on Equity Market Structure* 75 Fed Reg 3594 (21 January 2010).

US Securities and Exchange Commission *Staff Report on Algorithmic Trading in US Capital Markets* (August 2020).

US Securities and Exchange Commission "SEC Charges New York-Based High-Frequency Trading Firm with Fraudulent Trading to Manipulate Closing Prices" (Press Release 2014-229, 16 October 2014).

US Securities and Exchange Commission "SEC Charges Former Day Trader with Market Manipulation Scheme" (Litigation Release No 24989, 23 December 2020).

US Securities and Exchange Commission "SEC Charges Webcast Host for Role in Market Manipulation Scheme" (Press Release 2021-206, 25 August 2021).

D. Books

Abbott R *The Reasonable Robot: Artificial Intelligence and the Law* (Cambridge University Press, Cambridge, 2020).

Beever A *Rediscovering the Law of Negligence* (Hart Publishing, Oxford, 2007).

Cecconi F (ed) *AI in the Financial Markets: New Algorithms and Solutions* (Springer, Cham, 2023).

Fox MB, Glosten LR and Rauterberg GV *The New Stock Market: Law, Economics, and Policy* (Columbia University Press, New York, 2019).

Goudkamp J *Tort Law Defences* (Hart Publishing, Oxford, 2013).

Hildebrandt M *Law for Computer Scientists and Other Folk* (Oxford University Press, Oxford, 2020).

Hilpisch Y *Artificial Intelligence in Finance: A Python-Based Guide* (O'Reilly Media, Sebastopol (CA), 2020).

Jafar SH, ElChaarani H and Moturi S *Artificial Intelligence for Capital Markets* (CRC Press, Boca Raton (FL), 2023).

Kuhlman D *A Python Book: Beginning Python, Advanced Python, and Python Exercises* (Platypus Global Media, 2011).

Nilsson NJ *Introduction to Machine Learning* (Stanford University, draft text, 1998).

Pasquale F *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press, Cambridge (MA), 2015).

Scopino G *Algo Bots and the Law: Technology, Automation, and the Regulation of Futures and Derivatives* (Cambridge University Press, Cambridge, 2020).

Smith M *A History of the Global Stock Market: From Ancient Rome to Silicon Valley* (University of Chicago Press, Chicago, 2004).

Stevens R *Torts and Rights* (Oxford University Press, Oxford, 2007).

Todd S (ed) *The Law of Torts in New Zealand* (8th ed, Thomson Reuters, Wellington, 2023).

Turner J *Robot Rules: Regulating Artificial Intelligence* (Palgrave Macmillan, London, 2019).

E. Chapters in Edited Books

Aksoy PÇ "AI as Agents" in DiMatteo LA, Poncibò C and Cannarsa M (eds) *The Cambridge Handbook of Artificial Intelligence: Global Perspectives on Law and Ethics* (Cambridge University Press, Cambridge, 2022) 153.

Bagshaw R "The Edges of Tort Law's Rights" in Nolan D and Robertson A (eds) *Rights and Private Law* (Hart Publishing, Oxford, 2012) 405.

Ben-Shahar O "Causation and Foreseeability" in Arlen J (ed) *Research Handbook on the Economics of Torts* (Edward Elgar, Cheltenham, 2013) 83.

Ebers M "Regulating AI and Robotics: Ethical and Legal Challenges" in Ebers M and Navas S (eds) *Algorithms and Law* (Cambridge University Press, Cambridge, 2020) 37.

Kern S and Loiacono G "High-Frequency Trading and Circuit Breakers in the EU" in Mattli W (ed) *Global Algorithmic Capital Markets: High-Frequency Trading, Dark Pools, and Regulatory Challenges* (Oxford University Press, Oxford, 2018) 308.

Sharkey CM "Product Liability in the Digital Age: Liability for Algorithmic Harms" in McCausland K (ed) *Research Handbook on Law and Artificial Intelligence* (Edward Elgar, Cheltenham, 2024) 211.

Spindler G "Control of Algorithms in Financial Markets: The Example of High-Frequency Trading" in Ebers M and Navas S (eds) *Algorithms and Law* (Cambridge University Press, Cambridge, 2020) 207.

Stapleton J "Duty of Care Factors: A Selection from the Judicial Menus" in Cane P and Stapleton J (eds) *The Law of Obligations: Essays in Celebration of John Fleming* (Oxford University Press, Oxford, 1998) 59.

Van Hoecke M "Legal Doctrine: Which Method(s) for What Kind of Discipline?" in Van Hoecke M (ed) *Methodologies of Legal Research* (Hart Publishing, Oxford, 2011) 1.

Yadav Y "Algorithmic Trading and Market Regulation" in Mattli W (ed) *Global Algorithmic Capital Markets: High-Frequency Trading, Dark Pools, and Regulatory Challenges* (Oxford University Press, Oxford, 2019) 232.

F. Journal Articles

Abdussalam M "Redefining Legal Responsibility for Pure Economic Loss in the Innovation Economy" (2020) 26 *Auckland University Law Review* 33.

Armour J and Eidenmüller H "Self-Driving Corporations?" (2020) 10 *Harvard Business Law Review* 87.

Azzutti A "AI Trading and the Limits of EU Law Enforcement in Deterring Market Manipulation" (2022) 45 *Computer Law & Security Review* 105690.

Azzutti A "AI Governance in Algorithmic Trading: Some Regulatory Insights from the EU AI Act" [2024] *SSRN Electronic Journal*.

Azzutti A, Ringe W-G and Stiehl HS "Machine Learning, Market Manipulation, and Collusion on Capital Markets: Why the 'Black Box' Matters" (2021) 43 *University of Pennsylvania Journal of International Law* 79.

Bambauer DE "'Negligence and AI's Human Users' — A Reply" (2021) 101 *Boston University Law Review Online* 99.

Beck A "The Financial Markets Conduct Act and Market Manipulation" [2014] *NZLJ* 354.

Beever A "A Rights-Based Approach to the Recovery of Economic Loss in Negligence" (2004) 4(1) *Oxford University Commonwealth Law Journal* 25.

Bennett Moses L "How to Think about Law, Regulation and Technology: Problems with 'Technology' as a Regulatory Target" (2013) 5 *Law, Innovation & Technology* 1.

Beylin I "Taxing Fictive Orders: How an Information-Forcing Tax Can Reduce Manipulation and Distortion in Financial Product Markets" (2018) 85 *University of Cincinnati Law Review* 1.

Bhardwaj A and others "Risk Assessment and Management in Stock Trading Using Artificial Intelligence" in *2024 3rd International Conference on Sentiment Analysis and Deep Learning (ICSADL)* (IEEE, 2024) 138.

Casey B and Lemley MA "You Might Be a Robot" (2019) 105 *Cornell Law Review* 287.

Cetorelli N, Hirtle B, Morgan D, Peristiani S and Santos J "Trends in Financial Market Concentration and Their Implications for Market Stability" (2007) 13(1) *Economic Policy Review* 33.

Chaboud AP, Chiquoine B, Hjalmarsson E and Vega C "Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market" (2014) 69 *Journal of Finance* 2045.

Choi BH "Software as a Profession" (2020) 33 *Harvard Journal of Law & Technology* 557.

Choi BH "Negligence Liability for AI Developers" *Lawfare* (online, 26 September 2024).

Dakalbab F and others "Artificial Intelligence Techniques in Financial Trading: A Systematic Literature Review" (2024) 36 *Journal of King Saud University — Computer and Information Sciences* 102015.

Enriques L and Zetzsche D "Corporate Technologies and the Tech Nirvana Fallacy" (2020) 72 *Hastings Law Journal* 55.

Fischel DR and Ross DJ "Should the Law Prohibit 'Manipulation' in Financial Markets?" (1991) 105 *Harvard Law Review* 503.

Fletcher G-GS "Deterring Algorithmic Manipulation" (2021) 74 *Vanderbilt Law Review* 259.

Gomber P and others "High-Frequency Trading" [2011] *SSRN Electronic Journal*.

Hacker P "The European AI Liability Directives — Critique of a Half-Hearted Approach and Lessons for the Future" (2023) 51 *Computer Law & Security Review* 105871.

Haferkorn M and Zimmermann K "The German High-Frequency Trading Act" [2014] *SSRN Electronic Journal*.

Hutchinson T "The Doctrinal Method: Incorporating Interdisciplinary Methods in Reforming the Law" (2015) 8 *Erasmus Law Review* 130.

Jørgensen K, Skjeltorp J and Ødegaard BA "Throttling Hyperactive Robots — Order-to-Trade Ratios at the Oslo Stock Exchange" (2018) 37 *Journal of Financial Markets* 1.

Kaminski ME "Binary Governance: Lessons from the GDPR's Approach to Algorithmic Accountability" (2019) 92 *Southern California Law Review* 1529.

Karnow CEA "Liability for Distributed Artificial Intelligences" (1996) 11 *Berkeley Technology Law Journal* 147.

Kirilenko AA, Kyle AS, Samadi M and Tuzun T "The Flash Crash: High-Frequency Trading in an Electronic Market" (2017) 72 *Journal of Finance* 967.

Kirilenko A and Lo AW "Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents" (2013) 27 *Journal of Economic Perspectives* 51.

Kitchin R "Thinking Critically About and Researching Algorithms" (2017) 20(1) *Information, Communication & Society* 14.

Korsmo CR "High-Frequency Trading: A Regulatory Strategy" (2013) 48 *University of Richmond Law Review* 523.

Lai JC "The Consumer Guarantees Act and Digital Products" (2019) 50 *Victoria University of Wellington Law Review* 385.

Lee B-S, Rui OM and Wang SS "Information Transmission Between the NASDAQ and Asian Second Board Markets" (2004) 28 *Journal of Banking & Finance* 1637.

Lee J and Schu L "Algorithmic Trading Regulation: The Frameworks for Human Supervision and Direct Market Interventions" [2021] *SSRN Electronic Journal*.

Lemley M and Casey B "Remedies for Robots" (2019) 86 *University of Chicago Law Review* 1311.

Lin TCW "The New Investor" (2013) 60 *UCLA Law Review* 678.

Lin TCW "The New Market Manipulation" (2017) 66 *Emory Law Journal* 1253.

Lipton ZC "The Mythos of Model Interpretability" (2018) 16(3) *ACM Queue* 31.

MacKenzie D "How Algorithms Interact: Goffman's 'Interaction Order' in Automated Trading" (2019) 36(2) *Theory, Culture & Society* 39.

Mandes A "Algorithmic and High-Frequency Trading Strategies: A Literature Review" (MAGKS Joint Discussion Paper Series No 2, 2016).

McGowan M "The Rise of Computerized High Frequency Trading: Use and Controversy" (2010) 9 *Duke Law & Technology Review* 1.

Menkveld AJ "The Economics of High-Frequency Trading: Taking Stock" (2016) 8 *Annual Review of Financial Economics* 1.

Nishimura KG "Electronic Trading and Financial Markets" (paper presented to International Paris-Europlace Financial Forum, Tokyo, 29 November 2010).

O'Hara M "High Frequency Market Microstructure" (2015) 116 *Journal of Financial Economics* 257.

Oyedokun O, Ewim SE and Oyeyemi OP "Leveraging Advanced Financial Analytics for Predictive Risk Management and Strategic Decision-Making in Global Markets" (2024) 2(2) *Global Journal of Research in Multidisciplinary Studies* 16.

Prewitt M "High-Frequency Trading: Should Regulators Do More?" (2012) 19 *Michigan Telecommunications and Technology Law Review* 131.

Qiu Y, Song Z and Chen Z "Short-Term Stock Trends Prediction Based on Sentiment Analysis and Machine Learning" (2022) 26(5) *Soft Computing* 2209.

Reimann M "The Progress and Failure of Comparative Law in the Second Half of the Twentieth Century" (2002) 50 *American Journal of Comparative Law* 671.

Ren R, Wu DD and Liu T "Forecasting Stock Market Movement Direction Using Sentiment Analysis and Support Vector Machine" (2019) 13(1) *IEEE Systems Journal* 760.

Rjumohan *Stock Markets: An Overview and a Literature Review* (Munich Personal RePEc Archive Paper No 101855, 2019).

Samuel G "Is Legal Knowledge Cumulative?" (2012) 32 *Legal Studies* 448.

Sardar H and Sood C "Cryptocurrency Regulation: Need of the Hour" (2024) 2(1) *GLS KALP: Journal of Multidisciplinary Studies* 23.

Scherer MU "Of Wild Beasts and Digital Analogues: The Legal Status of Autonomous Systems" (2018) 19 *Nevada Law Journal* 259.

Selbst AD "Negligence and AI's Human Users" (2020) 100 *Boston University Law Review* 1315.

Slemmer CV "Crashed by Algorithms: Trader Liability in a New Era of Securities Regulation" (2017) 12 *Brooklyn Journal of Corporate, Financial & Commercial Law* 169.

Stiglitz JE "Tapping the Brakes: Are Less Active Markets Safer and Better for the Economy?" (paper, Federal Reserve Bank of Atlanta 2014 Financial Markets Conference, 15 April 2014).

Tang K and Xiong W "Index Investment and the Financialization of Commodities" (2012) 68(6) *Financial Analysts Journal* 54.

Tutt A "An FDA for Algorithms" (2017) 69 *Administrative Law Review* 83.

Twigg-Flesner C "Guiding Principles for Updating the Product Liability Directive for the Digital Age" (Briefing Note, European Law Institute, 2021).

Vitali S, Glattfelder JB and Battiston S "The Network of Global Corporate Control" (2011) 6(10) *PLoS ONE* e25995.

Vladeck DC "Machines Without Principals: Liability Rules and Artificial Intelligence" (2014) 89 *Washington Law Review* 117.

Wagner G "Liability for Artificial Intelligence: A Proposal of the European Parliament" (2021) 28 *Maastricht Journal of European and Comparative Law* 793.

Wallach W "Implementing Moral Decision-Making Faculties in Computers and Robots" (2007) 22(4) *AI & Society* 463.

Watson S "Investor Protection in New Zealand: A Tale of Two Statutes" (2014) 20 *New Zealand Business Law Quarterly* 245.

Wendehorst C "Strict Liability for AI and Other Emerging Technologies" (2020) 11 *Journal of European Tort Law* 150.

Witting C "Duty of Care: An Analytical Approach" (2005) 25 *Oxford Journal of Legal Studies* 33.

Wood D "AI Is the Next Cyber: Why Insurers May Carve Out Standalone Cover" *Insurance Business Australia* (online, 24 February 2026).

Yadav Y "How Algorithmic Trading Undermines Efficiency in Capital Markets" (2015) 68 *Vanderbilt Law Review* 1607.

Yadav Y "Insider Trading and Market Structure" (2016) 63 *UCLA Law Review* 968.

Yadav Y "The Failure of Liability in Modern Markets" (2016) 102 *Cornell Law Review* 1031.

Yeung K "Algorithmic Regulation: A Critical Interrogation" (2018) 12 *Regulation & Governance* 505.

G. Online and Institutional Resources

Choi BH "Negligence Liability for AI Developers" *Lawfare* (26 September 2024) <www.lawfaremedia.org/article/negligence-liability-for-ai-developers>.

Commodity Futures Trading Commission "CFTC Orders Panther Energy Trading LLC and Its Principal Michael J Coscia to Pay $2.8 Million…" (Release No 6649-13, 22 July 2013) <www.cftc.gov/PressRoom/PressReleases/6649-13>.

Brown D "Securities Law: New Zealand" in *Encyclopaedia of New Zealand Law* (online ed, LexisNexis).

Commerce Commission *The Fair Trading Act and New Technologies* (Discussion Paper, March 2020).

Financial Markets Authority "FMA Files Market Manipulation Case against Kok Ding Cheng for Trading of Rua Shares" (Media Release MR No 2022-47).

Financial Markets Authority "Former Promisia Chairman Admits Insider Trading Breaches" (Media Release MR No 2019-44).

London Stock Exchange "About the LSE: History" London Stock Exchange <https://www.londonstockexchange.com>.

Morgan Lewis *Algorithmic Trading: Hong Kong Monetary Authority Discusses Supervisory Expectations and Evolving Market Manipulation Risk* (March 2020).

NZX "Organisation Structure" NZX <https://www.nzx.com/about-nzx/organisation-structure>.

NZX "Market Participants" NZX <https://www.nzx.com/services/market-participants>.

NZX RegCo "About NZ RegCo" NZX RegCo <https://www.nzx.com/regulation/nzregco/about-nzregco>.

Pinsent Masons "Revised EU Product Liability Regime Expands Liability for AI and Software Providers" (2024).

United States Securities and Exchange Commission "SEC Charges New York-Based High-Frequency Trading Firm with Fraudulent Trading to Manipulate Closing Prices" (Press Release 2014-229, 16 October 2014).

United States Securities and Exchange Commission “SEC Charges Former Day Trader with Market Manipulation Scheme” (Litigation Release No 24989, December 2020).

United States Securities and Exchange Commission “SEC Charges Webcast Host for Role in Market Manipulation Scheme” (Press Release 2021-206, 2021).