

Stochastic Control and Systemic Stability: Governing AI in Healthcare and Algorithmic Financial Markets

Sana Muthalif

DEPARTMENT OF 5 YEAR INTERGRATED LAW, BBA LLB

CHRIST ACADEMY INSTITUTE OF LAW, BENGALURU

Abstract

The swift implementation of Artificial Intelligence (AI) and Machine Learning (ML) systems in the healthcare industry and trading infrastructure has always resulted in the development of very high-stakes decision-making which are a complex of adaptive systems, having properties of opacity, feedback, and nonlinear risk propagation. The existing regulatory structures are mainly static and and sectoral fragmented. Therefore, they are not suitable for the regulation of adaptive systems, which are time-varying. The current research work formulates the regulation of AI in the healthcare industry and financial markets as a problem of dynamic risk minimization. The central research problem concerns the design of regulatory architectures that minimize expected social loss:¹

$$f(x) = a_0 + \sum_{n=1}^{\infty} \left(a_n \cos \frac{n\pi x}{L} + b_n \sin \frac{n\pi x}{L} \right)$$

$$RminE[L(Ht, Mt | R)],$$

where (Ht) denotes population health outcomes, (Mt) represents market stability outcomes, and ($\{R\}$) is the regulatory regime, subject to constraints on innovation efficiency, information asymmetry, and institutional feasibility. In healthcare, AI-enabled diagnostic and decision-support systems—regulated under adaptive Software as a Medical Device framework like those which are articulated by the U.S. Food and Drug Administration—operating as a probabilistic classifier which are vulnerable to the distributional shift ($(P_t \neq P_{t+1})$) and model drift. In financial markets, algorithmic trading systems governed under frameworks including the Markets in Financial Instruments Directive II (MiFID II) generate endogenous volatility through latency-sensitive strategic interaction and high-frequency feedback dynamics.

The current study advances majorly three objectives: (i) construction of a unified stochastic representation of various algorithmic risk across healthcare and financial markets; (ii) evaluating the contemporary regulatory instruments, including the risk-tiered architecture of the Artificial Intelligence Act, through the control-theoretic and principal-agent lenses; and (iii) proposing an adaptive governance model with the integration of ex ante certification with continuous ex post supervisory feedbacks. Methodologically, the research also synthesizes the preliminary stochastic process modelling, network contagion analysis, and dynamic game theory. Healthcare AI systems are modelled as a sequential decision process with cumulative regret:

$$RT = t = 1 \sum T(Lt(\pi t) - Lt(\pi *)),$$

¹ 1 OECD, *OECD Artificial Intelligence Papers* (OECD Publishing 2023).

The main argument is that both intelligent care systems and high-frequency trading architectures can produce systemic externalities due to correlated error structures and feedback amplification. In the healthcare domain, correlated misclassification errors can always lead to population-level harm, whereas in the financial domain, correlated trading strategies can lead to various flash crashes.² These phenomena have a similar structure to the stability thresholds in complex networks, where systemic collapse usually happens if the spectral radius of the interaction matrix is satisfied by the condition ($\rho(A) > 1$). Thus, regulatory design must incorporate dynamic monitoring, stress testing, auditability constraints, and risk measures which are similar to Conditional Value-at-Risk (CVaR) but adjusted for the clinical harm.

By framing AI governance as a problem of stochastic control and systemic stability, the study advances a mathematically grounded framework for safeguarding both clinical integrity and market stability in algorithmically mediated environments.

Keywords:

Artificial Intelligence Governance /Algorithmic Trading /Systemic Risk/ Stochastic Control/ Market Integrity/ Adaptive Regulation

1. Introduction

Artificial Intelligence (AI) and Machine Learning (ML) are increasingly shaping various decision-making processes in the sectors where errors may produce severe social consequences. Two of the most critical domains are **healthcare** and **financial markets**, both of which now rely on algorithmic systems for prediction, classification, optimisation, and automated execution. In healthcare, AI is used in diagnostics, triage, treatment recommendation, and clinical decision support systems. Whereas In financial markets, the algorithmic and high-frequency trading systems influences liquidity, pricing, and market behaviour³.

While these technologies promise us efficiency, speed, and precision, they also generate numerous **systemic risks**. AI systems in healthcare may produce biased and inaccurate outcomes due to data drifting. Opacity training, or even misclassifications. Likewise, in algorithmic ⁴trading systems, one can witness amplify volatility, trigger liquidity shocks, and contribution to market instability through rapid feedback loops. These harms are not always isolated, rather, they may spread across institutions, users, and infrastructures across the world.

The present study also argues that conventional legal and regulatory frameworks are and will insufficient in future because they are largely **static**, **sector-specific**, and **reactive**. On the other hand, AI systems are **dynamic**, **adaptive**, and **probabilistic**. Therefore, this research proposes a more robust legal and governance framework thereby conceptualising AI regulations as a problem of **stochastic control**, **risk minimisation**, and **systemic stability**.

2. Statement of the Problem

Existing regulatory models that are used for governing AI in healthcare and algorithmic finance are fragmented and often fail to account for the **adaptive**, **nonlinear**, and **interdependent** nature of algorithmic systems, which are of

² World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021); International Organization of Securities Commissions, *Artificial Intelligence and Machine Learning in Capital Markets: Use Cases and Challenges* (IOSCO 2021)

³ International Organization of Securities Commissions, *Artificial Intelligence and Machine Learning in Capital Markets: Use Cases and Challenges* (IOSCO 2021)

⁴ World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021); International Organization of Securities Commissions, *Artificial Intelligence and Machine Learning in Capital Markets: Use Cases and Challenges* (IOSCO 2021)

utmost importance. This creates a major legal and governance gap: regulators are attempting to control systems that evolve faster than the rules designed to govern them.

In healthcare, this may lead to various unsafe diagnostic outputs, clinical harm, and accountability deficits that can cause irreparable harm. In financial markets, it may also result in flash crashes, manipulation risks, and broader threats. The central problem, therefore, can be concluded as:

How can legal and regulatory frameworks be designed to govern AI systems in a manner that minimises systemic harm while preserving innovation and institutional feasibility?

3. Research Questions

This study seeks to answer the following questions:

1. **How can risks generated by AI systems in healthcare and algorithmic financial markets be understood through a common analytical framework?**
2. **To what extent do existing regulatory models adequately address the dynamic and systemic risks posed by such systems?**
3. **Can principles of stochastic control and systemic stability inform a more adaptive model of AI governance in coming era?**
4. **What kind of regulatory architecture can balance innovation, accountability, safety, and public trust?**

4. Objectives of the Study

Primary Objective

To develop a **unified regulatory framework** for the governing AI systems in healthcare and algorithmic financial markets through the lens of **stochastic control and systemic stability**.

Secondary Objectives

1. To examine the operation of AI and ML systems in various healthcare and algorithmic trading.
2. To identify the nature of systemic and correlated risks emerging from these systems.
3. To analyse the existing legal and regulatory frameworks governing AI in both sectors.
4. To evaluate if the current models are adequate to address dynamic algorithmic harm.
5. To propose an **adaptive governance model** based on continuous monitoring, auditability, and risk-sensitive regulation.

5. Hypothesis

This study is based on the following hypothesis: The Static and fragmented regulatory frameworks are inadequate to govern adaptive AI systems in healthcare and financial markets; an adaptive and mathematically informed governance model based on stochastic control and systemic stability would provide a more effective and efficient mechanical framework for minimising systemic harm.

6. Scope of the Study

The scope of this research is interdisciplinary, situated at the intersection of:

- Law and Technology
- Artificial Intelligence Governance

- Healthcare Regulation
- Financial Market Regulation
- Risk Theory and Systems Analysis

The study is confined at two high-stakes sectors:

1. **Healthcare AI**, including various AI-enabled diagnostics, decision-support systems, and Software as a Medical Device (SaMD).
2. **Algorithmic Financial Markets**, especially dealing with the automated trading and high-frequency trading systems.

7. Significance of the Study

This study is significant for several reasons:

- It contributes to the growing discourse on **AI regulation in any of the high-risk sectors**.
- It introduces a **comparative and unified framework** for analysing risks in both healthcare and financial markets.
- It may also assist the policymakers, regulators, legal scholars, and institutional decision-makers in understanding why **adaptive governance** is necessary in algorithmically mediated environments.

The research is especially relevant in the context of emerging global AI regulation, including risk-based regulatory models and post-deployment monitoring obligations.

9. Research Methodology

This study adopts a doctrinal and analytical research methodology with an interdisciplinary orientation.

Nature of Research

- **Doctrinal**
- **Analytical and comparative**
- **Interdisciplinary**

Method Used

The study will rely on:

- Analysis of **primary legal materials** such as statutes, regulations, policy documents, and institutional frameworks;
- Examination of **secondary sources** including books, journal articles, reports, and academic commentary;
- Comparative evaluation of governance models in **healthcare AI** and **algorithmic finance**;
- Conceptual application of **stochastic process theory**, **systemic risk analysis**, and **control-based regulatory reasoning**.

10. Chapterisation

CHAPTER 1

INTRODUCTION AND CONCEPTUAL FRAMEWORK

1.1 Introduction

Artificial Intelligence (AI) and Machine Learning (ML) have become integral to contemporary governance, commerce, and public welfare. Their deployment in **healthcare** and **financial markets** has fundamentally altered the architecture of various decision-making arenas by shifting key judgments from human cognition to probabilistic computational systems. These sectors are especially significant because they involve **high-stakes outcomes**, that is the healthcare AI affects diagnosis, treatment, and patient safety, while financial AI affects liquidity, market pricing, investment behaviour, and systemic stability⁵.

The growing use of AI in these domains has generated considerable optimism regarding efficiency, speed, and predictive accuracy. Yet, these systems also produce **novel legal and regulatory challenges**. Unlike the traditional tools, AI systems are adaptive, they are data-dependent, and are capable of altering the behaviour over time in response to environmental inputs. Consequently, harms generated by such systems are not always isolated or linear. Instead, they often manifest through **feedback loops**, **distributional shifts**, **correlated decision errors**, and **network contagion**.

1.2 Background of the Study

Historically, regulation has assumed that the object of governance remains sufficiently stable to permit ex ante rulemaking and ex post enforcement. This assumption is increasingly strained in the context of current machine-learning systems. A clinical diagnostic model trained on one patient population may become unreliable when disease prevalence, imaging quality, or patient demographics shift. Likewise, a trading algorithm calibrated to one market environment may destabilise prices when interacting with similarly structured automated strategies under conditions of stress.

Thus, both healthcare AI and algorithmic trading may be understood as a **complex adaptive system**. A complex adaptive system is one in which multiple interacting agents produce aggregate outcomes that cannot be easily predicted from the behaviour of individual components alone.⁶ In legal terms, this means that regulatory responsibility cannot be and will not be confined to the isolated acts of negligence or rule violation; it must also address **emergent and systemic harms**.

The study therefore proceeds from the premise that regulation must be designed not only to control **individual algorithmic errors**, but also to contain **aggregate instability**⁷.

1.3 Research Problem

The central problem of this study lies in the mismatch between **dynamic algorithmic systems** and **static regulatory frameworks**. Existing regulatory structures in healthcare and financial markets are largely fragmented, sector-specific, and reactive. They are not fully equipped to address:

- model drift and data shift,
- probabilistic uncertainty,
- opaque decision pathways,
- correlated algorithmic failures,

⁵ World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021); International Organization of Securities Commissions, *Artificial Intelligence and Machine Learning in Capital Markets: Use Cases and Challenges* (IOSCO 2021)

⁶ European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final.

⁷ European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final.

- cascading systemic effects.

This raises a foundational legal question:

How should law regulate AI systems whose risks are probabilistic, evolving, and capable of producing systemic externalities?

The present research addresses this question by conceptualising AI governance as a problem of dynamic optimisation under uncertainty.

1.4 AI Governance as a Problem of Risk Minimisation

A central proposition of this research is that AI governance can be mathematically framed as a **social loss minimisation problem**. Let:

- (H_t) = healthcare outcomes at time (t)
- (M_t) = market stability outcomes at time (t)
- $(\mathcal{R})$ = regulatory regime
- $(L(\cdot))$ = aggregate social loss function

Then the objective of regulation may be represented as:

$$\left[\min_{\mathcal{R}} ; \mathbb{E} \left[L(H_t, M_t \mid \mathcal{R}) \right] \right]$$

This formulation means that the regulator seeks to choose a legal and institutional framework $(\mathcal{R})$ that minimises expected social harm arising from AI deployment across both sectors.

However, this optimisation is not unconstrained. It is subject to several real-world limitations:

$$\left[\begin{array}{l} \text{subject to } \quad \\ \begin{cases} I(\mathcal{R}) \geq \underline{I} \text{ \& \textit{innovation efficiency}} \\ A(\mathcal{R}) \leq \overline{A} \text{ \& \textit{information asymmetry}} \\ F(\mathcal{R}) \leq \overline{F} \text{ \& \textit{institutional feasibility}} \end{cases} \\ \end{array} \right]$$

This reflects an important legal insight: regulation must not only reduce risk, but must do so **without completely stifling innovation**, while remaining **administratively workable** and **epistemically realistic**.

1.5 Stochastic Control as a Regulatory Lens

Stochastic control is a branch of mathematics concerned with the complexive **decision-making in uncertain and evolving environments**. It is particularly useful in settings where future states depend partly on present actions and partly on random shocks⁸.

⁸ Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015).

This framework is highly relevant to AI governance because both healthcare AI and trading algorithms operate in uncertain environments where:

- the system evolves over time,
- outcomes are probabilistic rather than deterministic,
- interventions must often be made before all information is known.

A basic stochastic state equation may be represented as:

$$[X_{t+1} = f(X_t, U_t, \epsilon_t)]$$

Where:

- (X_t) = state of the system at time (t)
- (U_t) = control or regulatory intervention at time (t)
- (ϵ_t) = random disturbance or uncertainty term

In the present context:

- in healthcare, (X_t) may represent clinical model performance or patient-risk environment;
- in finance, (X_t) may represent liquidity conditions, volatility, or algorithmic interaction patterns;
- (U_t) may represent regulatory interventions such as audits, circuit breakers, certification, or post-deployment monitoring.

This allows us to think of regulation not as a one-time legal command, but as a **continuous supervisory control process**.

1.6 Systemic Stability and Network Effects

One of the most important insights of this study is that harm caused by AI is often **systemic**, not merely individual. In healthcare, correlated errors across hospitals or patient populations may produce widespread clinical harm. In financial markets, synchronised strategies may create flash crashes or self-reinforcing instability⁹.

A useful mathematical representation of systemic interaction is the network matrix (A) , where each entry (a_{ij}) reflects the influence of actor (j) on actor (i) . The stability of the system depends on the **spectral radius** of this interaction matrix:

$$[\rho(A) < 1 \quad \Rightarrow \quad \text{stable system}]$$

$$[\rho(A) > 1 \quad \Rightarrow \quad \text{risk of systemic amplification}]$$

⁹ Frank Pasquale, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015).

This formulation is highly relevant to legal regulation because it demonstrates that risk is not solely a function of the quality of an individual model, but also of **how multiple models interact within an institutional environment**. Thus, law must regulate not only algorithmic accuracy but also **structural interdependence**.

1.8 Significance of the Study

This study is significant because it contributes to the legal discourse on AI by moving beyond narrow debates around transparency and fairness and toward a more structural understanding of **systemic algorithmic governance**. It also bridges two usually separate regulatory domains—healthcare and finance—by demonstrating that both involve common problems of **uncertainty, adaptation, and cascading risk**.¹⁰

Its broader contribution lies in suggesting that future AI regulation must become **adaptive, feedback-sensitive, and resilience-oriented**, rather than merely static and compliance-based.

CHAPTER 2

AI IN HEALTHCARE: RISK, REGULATION, AND CLINICAL GOVERNANCE

2.1 Introduction

The healthcare sector has witnessed a rapid expansion in the use of AI and ML tools for diagnostics, triage, image interpretation, disease prediction, treatment planning, and administrative optimisation. These systems promise enhanced efficiency, early detection, and data-driven clinical decision support. However, because healthcare directly implicates **human life, bodily integrity, and patient autonomy**, errors in AI systems may produce consequences of exceptional gravity.

2.2 AI as a Probabilistic Decision System in Healthcare

Most AI systems in healthcare function as **probabilistic classifiers** or **predictive models**. Rather than delivering certain truths, they assign likelihoods to possible outcomes¹¹. For example, an AI model may estimate the probability that a patient has a tumour, is at risk of sepsis, or requires urgent intervention.

Formally, let:

- (X) = patient data (e.g., symptoms, imaging, vitals, history)
- $(Y \in \{0,1\})$ = true health condition
- $(\hat{Y} = f_{\theta}(X))$ = AI-generated prediction

Then the system learns a function:

$$[f_{\theta} : X \rightarrow \hat{Y}]$$

where (θ) represents the model parameters estimated from training data.

The model attempts to minimise prediction error through a loss function such as:

$$[$$

¹⁰ Martin Peterson, *An Introduction to Decision Theory* (2nd edn, Cambridge University Press 2017); Dimitri P Bertsekas, *Dynamic Programming and Optimal Control* (4th edn, Athena Scientific 2017).

¹¹ World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021).

$$\min_{\theta} \mathbb{E}_{(X,Y)} [\ell(f_{\theta}(X), Y)]$$

In legal and regulatory terms, this is crucial because the AI system is not a neutral tool; it is a **risk-producing inference mechanism** whose outputs influence clinical action.

2.3 Sources of Risk in Healthcare AI

Healthcare AI systems generate risk from multiple sources, including:

- **data bias**
- **misclassification**
- **lack of interpretability**
- **distributional shift**
- **automation bias**
- **institutional over-reliance**

A key problem is that models often perform well under training conditions but degrade in real-world deployment. This occurs when the statistical distribution of input data changes over time. If:

$$P_t(X, Y) \neq P_{t+1}(X, Y)$$

then the model is exposed to **distributional shift**. In simple terms, the population on which the model is used is no longer the same as the population on which it was trained.

This may arise due to:

- demographic changes,
- new disease patterns,
- altered imaging practices,
- local variations in hospital workflows.

From a legal perspective, this creates a serious accountability problem. A model may have been “safe” at the point of approval but become unsafe in operation¹². Static certification alone therefore becomes inadequate.

2.4 Clinical Error and Misclassification Harm

AI in healthcare often operates within binary or multi-class classification settings. Consider a simplified diagnostic model. Its performance may be evaluated through:

- **False Positives (FP)** – diagnosing disease when none exists
- **False Negatives (FN)** – failing to detect disease when it exists

¹² Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (Basic Books 2019).

The expected clinical harm can be represented as:

$$R = c_{\{FP\}} \cdot P(FP) + c_{\{FN\}} \cdot P(FN)$$

Where:

- $(c_{\{FP\}})$ = cost of false positive
- $(c_{\{FN\}})$ = cost of false negative

This formula is important because not all errors are equally harmful. In medical law and regulation, the severity of harm matters. For example, a false negative in cancer detection may be far more damaging than a false positive requiring additional tests.

Thus, healthcare regulation must move beyond generic accuracy metrics and toward **harm-sensitive performance evaluation**.

2.5 Sequential Learning and Cumulative Regret

Healthcare AI systems are often deployed in **sequential decision environments**, where the consequences of each decision may accumulate over time. This can be modelled through the concept of **cumulative regret**, which measures the performance loss of an algorithm relative to the best possible policy.

Let:

- (a_t) = action chosen by the AI at time (t)
- (a_t^*) = optimal action
- $(r(a_t))$ = reward or utility from chosen action

Then cumulative regret over (T) periods is:

$$\text{Regret}(T) = \sum_{t=1}^T [r(a_t^*) - r(a_t)]$$

In healthcare, “regret” may be interpreted as **avoidable clinical harm** caused by repeated suboptimal decisions. This¹³ is particularly relevant in triage systems, treatment recommendation engines, or predictive prioritisation systems.

A legally significant implication follows: a system that appears acceptable in isolated cases may still produce **unacceptable aggregate harm** when deployed at scale over time.

2.6 Correlated Errors and Population-Level Harm

One of the most underappreciated dangers of healthcare AI is the possibility of **correlated error**. If a flawed model is deployed across multiple hospitals or health systems, its errors are not random or isolated; they may become structurally replicated across thousands of patients.

¹³ Nuffield Council on Bioethics, *Artificial Intelligence (AI) in Healthcare and Research* (Nuffield Council on Bioethics 2023).

Let:

- (e_i) = error of AI system in institution (i)
- $(\text{Cov}(e_i, e_j))$ = covariance of errors across institutions

If:

$$[\text{Cov}(e_i, e_j) > 0]$$

then errors are positively correlated. This means that one systemic flaw may affect multiple institutions simultaneously.

From a governance perspective, this transforms AI risk from an issue of **individual malpractice** into one of **public regulatory concern**. The legal system must therefore ask not only whether a tool is safe for one patient, but whether its widespread deployment may create **population-scale harm**.

14

2.7 Regulatory Challenges in Healthcare AI

Healthcare law traditionally relies on frameworks such as:

- professional negligence,
- product liability,
- medical device regulation,
- informed consent,
- institutional accountability.

However, AI complicates each of these.

First, causation becomes difficult:

If an AI recommendation contributes to harm, who is responsible?

- the developer?
- the hospital?
- the physician?
- the data supplier?
- the regulator?

Second, explainability becomes contested:

If the system's internal reasoning is opaque, how can clinicians meaningfully supervise it?

¹⁴ Kevin P Murphy, *Probabilistic Machine Learning: An Introduction* (MIT Press 2022).

Third, safety becomes dynamic:

A model may continue to evolve post-deployment, rendering one-time approval inadequate.

This creates a need for **adaptive clinical governance**, where oversight continues throughout the operational lifecycle of the system.

2.8 Toward a Dynamic Regulatory Model for Healthcare AI

A more suitable governance approach would treat healthcare AI as a **continuously supervised risk system**. Instead of relying solely on pre-market approval, regulation should incorporate:

- ongoing post-deployment monitoring,
- performance drift detection,
- independent audits,
- real-world validation,
- mandatory reporting of algorithmic adverse events.

This can be conceptualised through a dynamic safety threshold. Let (q_t) represent model quality at time (t), and let ($q_{\min}$) represent the minimum acceptable safety threshold.

Then regulatory intervention is required whenever:

$$\left[\begin{array}{l} q_t < q_{\min} \end{array} \right]$$

Similarly, if clinical harm exceeds acceptable tolerance (h^*), then:

$$\left[\begin{array}{l} H_t > h^* \end{array} \right]$$

should trigger corrective legal or institutional action.

This approach is consistent with the broader thesis of the present study: that AI governance must function as a **feedback-sensitive control architecture**, not merely as static rule enforcement.

CHAPTER 3

ALGORITHMIC TRADING AND FINANCIAL MARKET STABILITY

3.1 Introduction

Financial markets have undergone a profound transformation with the rise of **algorithmic trading**, **high-frequency trading (HFT)**, and AI-driven execution strategies. Trading decisions that were once mediated by human judgment are now frequently made by automated systems capable of analysing data, placing orders, and responding to market signals

within milliseconds¹⁵. These developments have improved speed and market efficiency, but they have also introduced new forms of **systemic fragility**, **strategic opacity**, and **feedback-driven volatility**.

This chapter examines algorithmic trading as a legally significant form of automated decision-making and argues that, like healthcare AI, it must be understood as a **dynamic risk system** rather than merely a tool of operational efficiency.

3.2 Nature of Algorithmic Trading

Algorithmic trading refers to the use of computer programs to execute orders according to pre-specified instructions relating to variables such as:

- price,
- timing,
- volume,
- market depth,
- arbitrage opportunities,
- predictive signals.

More advanced systems incorporate machine learning to adapt strategies in response to market conditions.

A simplified trading decision can be represented as:

$$\begin{bmatrix} a_t = \pi(s_t) \end{bmatrix}$$

Where:

- (s_t) = market state at time (t)
- (a_t) = trading action at time (t)
- (π) = trading policy or strategy

The market state (s_t) may include:

- bid-ask spread,
- order book depth,
- historical volatility,
- news signals,
- correlated asset movements.

This means that algorithmic trading systems are essentially **real-time decision engines** operating in a highly interactive environment¹⁶.

¹⁵ U.S. Food and Drug Administration, *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan* (FDA 2021).

¹⁶ Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 *Georgetown Law Journal* 1147

3.3 Financial Markets as Strategic Multi-Agent Systems

Unlike many healthcare AI applications, financial algorithms do not operate in isolation. They exist in markets populated by **multiple strategic agents**, each responding to the actions of others. This makes financial markets especially vulnerable to **feedback effects** and **collective instability**.¹⁷

A useful framework is the **partially observed Markov game**, where each algorithm chooses actions based on imperfect information about the state of the market and the strategies of others.

Let:

- (N) = number of trading agents
- (s_t) = market state
- (a_t^i) = action of trader (i) at time (t)

Then the market evolves according to:

$$\begin{bmatrix} s_{t+1} = g(s_t, a_t^1, a_t^2, \dots, a_t^N, \epsilon_t) \end{bmatrix}$$

This is legally and economically significant because it implies that harmful outcomes may arise not from a single defective algorithm, but from **collective interaction among individually rational systems**.

3.4 Volatility, Feedback Loops, and Endogenous Risk

A major concern in algorithmic finance is the creation of **endogenous risk**—that is, risk generated internally by the market system itself rather than by external shocks.

Suppose the price process of an asset is represented as:

$$\begin{bmatrix} P_{t+1} = P_t + \alpha D_t + \epsilon_t \end{bmatrix}$$

Where:

- (P_t) = asset price at time (t)
- (D_t) = aggregate net demand from trading algorithms
- (α) = market impact coefficient
- (ϵ_t) = exogenous shock

If many algorithms respond to the same signal in the same direction, then (D_t) may become highly concentrated, producing sudden and self-reinforcing price movements.

This is especially dangerous where algorithmic behaviour creates **positive feedback loops**, such as:

- buying because prices are rising,

¹⁷ 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

- selling because volatility is increasing,
- withdrawing liquidity when stress is detected.

Such dynamics may amplify ordinary fluctuations into systemic instability.

3.5 Flash Crashes and Liquidity Spirals

The phenomenon of the **flash crash** illustrates the systemic danger of algorithmic trading. A flash crash occurs when automated selling, liquidity withdrawal, and strategic feedback interact to produce a sudden collapse in prices, often followed by rapid recovery.

A stylised liquidity function may be expressed as:

$$[L_t = L_0 - \beta V_t]$$

Where:

- (L_t) = market liquidity at time (t)
- (L_0) = baseline liquidity
- (V_t) = volatility
- $(\beta > 0)$ = sensitivity of liquidity to volatility

This shows that as volatility rises, liquidity may fall. If algorithms are programmed to withdraw or reduce exposure during periods of uncertainty, then a destabilising loop emerges:

$(x + a)^n = \sum_{k=0}^n \binom{n}{k} x^k a^{n-k}$ This feedback structure is legally important because it demonstrates that market failure may occur even in the absence of fraud or intentional misconduct¹⁸. Existing legal frameworks focused only on manipulation or insider abuse may therefore fail to capture structurally destabilising behaviour.

3.6 Correlated Strategies and Systemic Contagion

One of the most important parallels with healthcare AI is the issue of **correlated error or correlated strategy**. In financial markets, many algorithms are trained on similar datasets, optimise for similar signals, and react to similar conditions. This creates a high probability of **synchronised behaviour**.

Let:

- (x_t^i) = exposure or position of agent (i)
- (C) = correlation matrix of trading strategies

If:

¹⁸ European Securities and Markets Authority, *Guidelines on Systems and Controls in an Automated Trading Environment for Trading Platforms, Investment Firms and Competent Authorities* (ESMA 2012).

$$\left[\begin{array}{l} \rho(x_t^i, x_t^j) \approx 1 \\ \end{array} \right]$$

for many (i, j) , then market participants are effectively moving together. This increases the likelihood of contagion.

Systemic stability can again be modelled through the interaction matrix (A) . If the dominant eigenvalue or spectral radius exceeds unity:

$$\left[\begin{array}{l} \rho(A) > 1 \\ \end{array} \right]$$

then shocks are amplified rather than absorbed.

This has major regulatory implications. The problem is not merely whether each algorithm individually complies with market rules, but whether the **market ecology as a whole** remains stable under stress.

3.7 Regulatory Challenges in Algorithmic Finance

Traditional financial regulation has focused on:

- disclosure,
- anti-manipulation rules,
- best execution,
- market abuse prevention,
- prudential supervision.¹⁹

These remain important, but they are increasingly insufficient for AI-driven trading environments.

Key regulatory difficulties include:

(a) Opacity

Complex strategies may be difficult for regulators, exchanges, or even deploying firms to interpret in real time.

(b) Speed

Harm may occur in milliseconds, far faster than traditional enforcement mechanisms can respond.

(c) Collective Interaction

Even lawful individual strategies may produce unlawful or harmful aggregate outcomes.

(d) Accountability Gaps

Responsibility may be dispersed across developers, firms, infrastructure providers, and compliance systems.

¹⁹ Rama Cont, 'The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instability' in Jean-Philippe Bouchaud and others (eds), *Market Microstructure and Nonlinear Dynamics* (Springer 2018)

Thus, algorithmic finance requires a shift from static compliance toward **dynamic market supervision**.

3.8 Risk Measures and Legal Supervision

One way to conceptualise regulatory supervision in financial AI is through formal risk thresholds. A common risk measure in finance is **Value at Risk (VaR)** or its more robust extension, **Conditional Value at Risk (CVaR)**.

For a loss variable (X), CVaR at confidence level (α) is:

$$\text{CVaR}_{\alpha}(X) = \mathbb{E}[X \mid X \geq \text{VaR}_{\alpha}(X)]$$

This measures the expected loss in the worst tail of the distribution.

The present study argues that financial AI regulation should not only assess ordinary performance but also **tail-risk behaviour** under extreme market conditions. This suggests the need for:

- stress testing,
- kill-switches,
- circuit breakers,
- algorithmic audits,
- real-time surveillance.

These mechanisms function as forms of **regulatory control input** designed to keep the market within acceptable stability bounds.

3.9 Toward Adaptive Regulation of Financial AI

An adaptive governance model for algorithmic trading would require the regulator to monitor not just firm-level compliance, but **system-wide behavioural patterns**. This includes:

- concentration of strategies,
- market reaction speed,
- liquidity resilience,
- cross-platform contagion,
- abnormal order behaviour²⁰.

If market instability at time (t) is represented by (S_t), and the regulator defines a permissible threshold (S^*), then intervention is required whenever:

$$S_t > S^*$$

²⁰ U.S. Securities and Exchange Commission and Commodity Futures Trading Commission, *Findings Regarding the Market Events of May 6, 2010* (2010).

Possible interventions may include:

- throttling order flow,
- halting execution,
- triggering supervisory review,
- imposing mandatory recalibration,
- suspending high-risk strategies.

This reflects the broader logic of stochastic control: the law must not merely punish after collapse, but **continuously stabilise before collapse occurs**.

CHAPTER 4

TOWARDS AN ADAPTIVE GOVERNANCE MODEL FOR AI SYSTEMS

4.1 Introduction

AI systems in healthcare and financial markets generate dynamic and systemic risks that are not adequately addressed by static legal regulation.

The central claim is that AI regulation must function less like a one-time certification exercise and more like a **continuous control architecture** capable of observing, correcting, and stabilising evolving systems over time.

4.2 Why Static Regulation Fails

Traditional regulatory systems often operate through a familiar sequence:

1. define the object of regulation,
2. establish ex ante standards,
3. certify compliance,
4. impose sanctions after violation.²¹

This model works reasonably well when the regulated object is relatively stable. It becomes much less effective where the regulated system:

- learns from new data,
- changes behaviour over time,
- interacts strategically with its environment,
- generates harms through network effects.

In such cases, compliance at time (t_0) does not guarantee safety at time (t_1).

Let (q_t) represent the quality or safety of an AI system at time (t).

A static regulator assumes:

$$[q_{t_0} \approx q_{t_1} \approx q_{t_2}]$$

²¹ World Economic Forum, *A Framework for the Governance of Artificial Intelligence at the Edge* (WEF 2020).

]

But for adaptive systems, this assumption may fail:

$$\begin{bmatrix} q_{t+1} = q_t + \Delta_t \end{bmatrix}$$

Where (Δ_t) may be positive or negative depending on drift, retraining, market adaptation, or deployment conditions.

Thus, the legal architecture must become **time-sensitive**.

5.3 The Logic of Adaptive Governance

Adaptive governance is a model of regulation in which oversight is:

- **continuous** rather than episodic,
- **feedback-driven** rather than static,
- **risk-sensitive** rather than formally uniform,
- **system-aware** rather than purely entity-specific.

Its underlying objective is not merely to prevent every possible error, which is impossible, but to ensure that the system remains within acceptable **stability and safety boundaries**.

This can be expressed through a control objective:

$$\begin{bmatrix} \min_{U_t} ; \mathbb{E} \left[\sum_{t=0}^T L(X_t, U_t) \right] \end{bmatrix}$$

Where:

- (X_t) = system state at time (t),
- (U_t) = regulatory intervention,
- ($L(X_t, U_t)$) = social loss associated with the system and intervention.

The regulator's task is therefore not only to prohibit misconduct, but to choose interventions that minimise expected long-run social harm.

5.4 Core Components of an Adaptive Governance Model

A legally robust adaptive governance model for AI should include the following five components:

5.4.1 Ex Ante Risk Classification

Not all AI systems pose the same degree of risk. A foundational step is to classify systems according to the gravity of harm they may cause.

Let:

$$[R_i = p_i \cdot c_i]$$

Where:

- (R_i) = expected risk of system (i) ,
- (p_i) = probability of failure,
- (c_i) = severity of consequences.

A risk-tiered regulatory system would subject higher-risk systems to more intensive obligations. This is especially relevant for:

- diagnostic and treatment AI in healthcare,
- autonomous execution and market-impact systems in finance.

Such a model prevents regulatory overreach for low-risk tools while ensuring stronger control over high-risk infrastructures.

5.4.2 Continuous Post-Deployment Monitoring

Ex ante approval is insufficient unless accompanied by real-world monitoring. Every high-risk AI system should be subject to **continuous performance surveillance** after deployment.

Let:

- (m_t) = measured system performance at time (t) ,
- (m^*) = minimum acceptable threshold.

Then intervention is required whenever:

$$[m_t < m^*]$$

This simple threshold logic is crucial in both sectors.

In healthcare

Monitoring should track:

- diagnostic sensitivity and specificity,
- subgroup performance,

- adverse clinical events,
- model drift.

In finance

Monitoring should track:

- abnormal order behaviour,
- liquidity impact,
- concentration of strategies,
- stress sensitivity.

Legally, this implies a move toward **operational accountability**, not merely approval-stage accountability.

5.4.3 Mandatory Auditability and Explainability

A major governance challenge is that many AI systems are difficult to interpret. While complete transparency may not always be technically feasible, **auditability** must be legally required.

Auditability means that regulators and accountable institutions should be able to reconstruct:

- what the system did,
- why it did it,
- under what conditions it failed,
- who authorised deployment.²²

This can be formalised through an observability condition. If the true internal state of the system is (X_t) , and observable outputs are (Y_t) , then governance requires that:

[
 $P(X_t \mid Y_t)$ be sufficiently informative
]

In simple terms, enough information must exist to permit meaningful oversight.

This is particularly important because legal responsibility cannot operate where decision systems are functionally inscrutable.

5.4.4 Stress Testing and Scenario Simulation

High-risk AI systems should be evaluated not only under ordinary conditions, but also under **adverse or tail-risk conditions**.

Let (Ω) denote the set of plausible operational states, and $(\Omega^{\text{stress}}) \subset \Omega$ denote extreme but possible stress conditions.

²² Julia Black, 'Decentring Regulation: Understanding the Role of Regulation and Self-Regulation in a "Post-Regulatory" World' (2001) 54 *Current Legal Problems* 103.

Then robust governance requires evaluation over:

$$\left[\begin{array}{l} \forall \omega \in \Omega^{\text{stress}}, \quad \text{Assess } f_{\theta}(\omega) \end{array} \right]$$

In healthcare, stress testing may include:

- rare disease prevalence changes,
- imaging noise,
- demographic imbalance,
- hospital overload conditions.

In finance, stress testing may include:

- liquidity withdrawal,
- volatility spikes,
- order-book thinning,
- correlated strategic response.²³

This is a crucial regulatory principle: systems should be judged not only by how they perform when everything is normal, but by how they behave when the environment becomes unstable.

5.4.5 Dynamic Intervention Mechanisms

Adaptive regulation requires the capacity to intervene when systems begin to deviate from acceptable performance boundaries.

Possible intervention tools include:

- model suspension,
- human override requirements,
- recalibration mandates,
- kill-switches,
- execution throttles,
- temporary deployment bans,
- targeted audit triggers.²⁴

Let (X_t) denote system state and (X^*) denote the acceptable region of operation.

Then the regulatory objective is:

$$\left[\begin{array}{l} X_t \in X^* \quad \forall t \end{array} \right]$$

²³ U.S. Food and Drug Administration, *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan* (FDA 2021).

²⁴ Orly Lobel, 'The Renew Deal: The Fall of Regulation and the Rise of Governance in Contemporary Legal Thought' (2004) 89 *Minnesota Law Review* 342.

Whenever:

[
 $X_t \notin X^*$
]

the regulator or supervised institution must apply a corrective control (U_t).

This is the practical heart of stochastic control as applied to governance.

5.5 Governance Architecture Across the AI Lifecycle

A coherent adaptive regulatory system must cover the entire **lifecycle** of AI deployment.

(a) Design Stage

- data quality obligations
- documentation of intended use
- bias and robustness assessment²⁵

(b) Pre-Deployment Stage

- validation and certification
- scenario testing
- legal and ethical impact assessment

(c) Deployment Stage

- operational monitoring
- human oversight obligations
- logging and reporting

(d) Post-Deployment Stage

- audit trails
- incident reporting
- recalibration and recall mechanisms

This lifecycle approach is superior to fragmented governance because it recognises that algorithmic harm may emerge at **any stage**, not merely at the moment of design.

5.6 Sector-Specific Application of Adaptive Governance

Though the underlying framework is unified, implementation must remain context-sensitive.

²⁵ World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021).

Healthcare AI

An adaptive governance model should require:

- clinical validation,
- subgroup fairness review,
- patient safety reporting,
- physician override rights,
- continuous post-market surveillance.

Algorithmic Finance

An adaptive governance model should require:

- strategy registration,
- market-impact testing,
- real-time supervision,
- stress simulations,
- kill-switch and circuit-breaker protocols.²⁶

Thus, the framework is **general in principle** but **sector-specific in implementation**.

5.7 Public Law and Institutional Responsibility

A key normative implication of this study is that AI governance should not be treated solely as a matter of private ordering or corporate self-regulation. Because these systems create **public risks**, governance must be anchored in **public-law values**, including:

- accountability,
- proportionality,
- procedural fairness,
- transparency,
- institutional responsibility,
- democratic legitimacy.

This means that regulators should not merely certify systems; they must also shape the institutional conditions under which algorithmic systems operate.

In this sense, the role of law is not only corrective but also **architectural**.

5.8 Proposed Unified Governance Function

The adaptive model advanced in this study may be expressed through the following unified governance function:

$$\begin{aligned} &[\\ G_t &= \phi(R_t, M_t, A_t, S_t, O_t) \\ &] \end{aligned}$$

²⁶ Philippe Jorion, *Value at Risk: The New Benchmark for Managing Financial Risk* (3rd edn, McGraw-Hill 2007).

Where:

- (G_t) = governance intensity at time (t) ,
- (R_t) = risk level,
- (M_t) = monitoring output,
- (A_t) = audit findings,
- (S_t) = systemic interconnectedness,
- (O_t) = observed drift or instability.

This means that regulation should become more intensive when:

- risk is high,
- performance deteriorates,
- audit failures appear,
- systemic concentration increases,
- environmental drift becomes significant.

This is far more rational than uniform, static compliance requirements²⁷.

CONCLUSION AND RECOMMENDATIONS

Introduction

This study set out to examine whether existing legal and regulatory frameworks are adequate to govern AI systems deployed in **healthcare** and **algorithmic financial markets**. It argued that these sectors, though institutionally distinct, share a common structural challenge: both are increasingly governed by adaptive algorithmic systems that produce uncertain, nonlinear, and potentially systemic harms.

By bringing together legal theory, systems analysis, and mathematical modelling, the research has attempted to develop a more rigorous framework for understanding and governing AI in high-risk environments.

Summary of the Core Argument

The central thesis of this study is that AI governance must be understood as a problem of **stochastic control and systemic stability**.

Traditional regulatory systems assume that the object of governance is relatively stable, knowable, and separable from its environment. AI systems challenge all three assumptions. They are:

- **adaptive** rather than static,
- **probabilistic** rather than deterministic,
- **networked** rather than isolated²⁸.

²⁷ Benedict Kingsbury, 'The Concept of "Law" in Global Administrative Law' (2009) 20 *European Journal of International Law* 23.

²⁸ Cary Coglianese and David Lehr, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 *Georgetown Law Journal* 1147.

This means that legal governance cannot be confined to ex ante certification or ex post liability alone. Instead, regulation must function as a **continuous, feedback-sensitive architecture** capable of monitoring and correcting evolving systems over time.

Major Findings of the Study

The research yields five principal findings:

First, AI in healthcare and algorithmic finance can be analysed through a common structural framework.

Both sectors involve automated decision systems operating under uncertainty and generating outcomes with high social consequence.

Second, the most serious harms generated by these systems are often systemic rather than individual.

This includes correlated misdiagnosis in healthcare and synchronised instability in financial markets.

Third, conventional legal doctrines such as negligence, foreseeability, and fault remain relevant but are insufficient on their own.

They struggle to address diffuse causation, probabilistic harm, and multi-actor accountability.

Fourth, static regulatory models are normatively and practically inadequate.

A system may satisfy formal compliance at one point in time yet become unsafe as its environment changes.

Fifth, an adaptive governance model grounded in continuous monitoring, auditability, stress testing, and dynamic intervention offers a more suitable regulatory response.

These findings support the broader claim that future AI law must be **structural, anticipatory, and resilience-oriented**.

Normative Implications

The study has important implications for legal theory and policy.

(a) AI Regulation Must Become Dynamic

Law must shift from a one-time approval logic to an ongoing supervisory logic.

(b) Systemic Risk Must Become a Legal Category

Regulators should recognise correlated and cascading algorithmic harm as a distinct governance concern.

(c) Public Institutions Must Retain Oversight Capacity

AI governance cannot be outsourced entirely to private firms or internal compliance teams.

(d) Accountability Must Be Distributed but Traceable

Where harms emerge from multi-actor systems, law must ensure that responsibility remains reconstructable and enforceable.²⁹

(e) Governance Must Be Sector-Sensitive but Conceptually Unified

While healthcare and finance require different implementation strategies, both can be governed under a shared logic of adaptive risk control.

Recommendations

Based on the findings of this study, the following recommendations are proposed:

1. Mandatory Risk-Based Classification

All high-impact AI systems should be categorised according to the probability and severity of harm they may cause.

$$[R = p \cdot c]$$

This should determine the level of regulatory scrutiny.

2. Continuous Post-Deployment Monitoring

Regulation should require ongoing monitoring of model performance in real-world conditions rather than relying solely on initial certification.

$$[m_t < m^* \rightarrow \text{mandatory intervention}]$$

3. Compulsory Audit Trails

Developers and deployers of high-risk AI systems should be legally required to maintain auditable logs of system design, updates, outputs, and human interventions.

4. Stress Testing for High-Risk Systems

AI systems should be tested under adverse and edge-case conditions before and after deployment.

²⁹ European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final.

[
forall $\omega \in \Omega^{\text{stress}}$, quad $\text{evaluate performance}$
]

5. Dynamic Regulatory Intervention Tools

Regulators and supervised institutions should possess authority to:

- suspend unsafe systems,
- impose recalibration,
- require human override,
- trigger emergency controls,
- halt destabilising deployment.

6. Sector-Specific Supervisory Frameworks

Separate operational guidance should be developed for:

- clinical AI systems,
- algorithmic trading and financial execution systems.

This avoids over-generalisation while preserving a coherent conceptual framework.

7. Recognition of Systemic Algorithmic Risk in Law

Legislators and regulators should explicitly recognise that AI harms may arise from:

- concentration,
- common model architecture,
- correlated behaviour,
- infrastructure dependence,
- scale of deployment.

This is necessary to move beyond purely individualised legal reasoning.

6.6 Contribution of the Study

This research contributes to the literature in three major ways:

First, it offers a comparative legal analysis of healthcare AI and algorithmic financial systems.

Second, it introduces mathematical and systems-theoretic concepts—including stochastic control, network stability, and dynamic optimisation—into legal analysis in a meaningful and non-superficial way.

Third, it proposes a unified adaptive governance framework that may be used by scholars, regulators, and institutions to rethink AI regulation in high-risk sectors.

In doing so, the study attempts to move the legal conversation beyond abstract ethical principles and toward **institutional design under uncertainty**.

6.7 Limitations of the Study

Like all theoretical research, this study has certain limitations.

1. It is primarily **doctrinal and analytical**, and does not include empirical testing of specific AI models.
2. The mathematical formulations are **conceptual and explanatory**, rather than computational simulations.
3. The legal analysis is broad and comparative, and therefore does not exhaust the doctrinal complexity of every national regulatory framework.
4. Because AI regulation is rapidly evolving, some policy frameworks may continue to develop beyond the temporal scope of the present research.

These limitations, however, do not weaken the central theoretical contribution of the study.

6.8 Scope for Further Research

This study opens several directions for future work:

- empirical testing of systemic-risk indicators in deployed AI systems;
- comparative study of national AI regulatory models;
- doctrinal analysis of liability allocation in multi-actor AI environments;
- judicial approaches to algorithmic accountability;
- design of sector-specific stress-testing protocols for healthcare and finance.

Future scholarship may also extend this framework to other high-risk sectors such as:

- criminal justice,
- autonomous mobility,
- public administration,
- insurance,
- critical infrastructure.

6.9 Final Conclusion

The rise of AI in healthcare and financial markets marks not simply a technological shift, but a transformation in the nature of institutional decision-making itself. Systems once governed by human discretion are increasingly mediated by adaptive, probabilistic, and interconnected machine architectures. The law cannot govern such systems effectively through static compliance models alone.

This study has argued that the proper legal response lies in reconceiving AI governance as a problem of **stochastic control and systemic stability**. That shift matters because it recognises that the central task of law is no longer only to assign blame after harm occurs, but to **design resilient institutions capable of anticipating, monitoring, and containing algorithmic risk before systemic failure takes hold**.

In that sense, the future of AI regulation will depend not merely on whether law can keep pace with technology, but on whether law itself can become **adaptive, intelligent, and structurally responsive**.

11. Expected Findings

The study is expected to show that:

1. AI systems in healthcare and algorithmic finance create **comparable forms of systemic risk** despite operating in different institutional settings.
2. Current legal frameworks remain **reactive and fragmented**, making them inadequate for adaptive algorithmic systems.³⁰
3. A regulatory model based on **continuous supervision, dynamic risk assessment, auditability, and feedback-sensitive governance** is more suitable than traditional compliance structures.
4. Law must move from a purely static compliance approach to a **resilience-based governance model** for high-risk AI systems.

BIBLIOGRAPHY

A. Books

Bertsekas DP, *Dynamic Programming and Optimal Control* (4th edn, Athena Scientific 2017).

Coglianese C and Lehr D, *Regulating by Robot* (where separately cited as article, do not repeat as book if not used in book form).

Jorion P, *Value at Risk: The New Benchmark for Managing Financial Risk* (3rd edn, McGraw-Hill 2007).

Lewis M, *Flash Boys: A Wall Street Revolt* (Allen Lane 2014).

Lo AW, *Adaptive Markets: Financial Evolution at the Speed of Thought* (Princeton University Press 2017).

Murphy KP, *Probabilistic Machine Learning: An Introduction* (MIT Press 2022).

Pasquale F, *The Black Box Society: The Secret Algorithms That Control Money and Information* (Harvard University Press 2015).

Peterson M, *An Introduction to Decision Theory* (2nd edn, Cambridge University Press 2017).

Quionero-Candela J and others (eds), *Dataset Shift in Machine Learning* (MIT Press 2009).

Simon HA, *The Sciences of the Artificial* (3rd edn, MIT Press 1996).

³⁰ European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final.

Topol E, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* (Basic Books 2019).

Watts DJ, *Six Degrees: The Science of a Connected Age* (Vintage 2004).

Hildebrandt M, *Law for Computer Scientists and Other Folk* (Oxford University Press 2020).

Sunstein CR, *Laws of Fear: Beyond the Precautionary Principle* (Cambridge University Press 2005).

B. Journal Articles / Chapters

Black J, 'Decentring Regulation: Understanding the Role of Regulation and Self-Regulation in a "Post-Regulatory" World' (2001) 54 *Current Legal Problems* 103.

Coglianese C and Lehr D, 'Regulating by Robot: Administrative Decision Making in the Machine-Learning Era' (2017) 105 *Georgetown Law Journal* 1147.

Cont R, 'The Endogenous Dynamics of Markets: Price Impact, Feedback Loops and Instability' in Jean-Philippe Bouchaud and others (eds), *Market Microstructure and Nonlinear Dynamics* (Springer 2018).

Doshi-Velez F and Kim B, 'Towards A Rigorous Science of Interpretable Machine Learning' (2017) arXiv:1702.08608.

Kirilenko A and Lo AW, 'Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents' (2013) 27 *Journal of Economic Perspectives* 51.

Kingsbury B, 'The Concept of "Law" in Global Administrative Law' (2009) 20 *European Journal of International Law* 23.

Lobel O, 'The Renew Deal: The Fall of Regulation and the Rise of Governance in Contemporary Legal Thought' (2004) 89 *Minnesota Law Review* 342.

Sendak MP and others, 'A Path for Translation of Machine Learning Products into Healthcare Delivery' (2020) 3 *EMJ Innovations* 65.

Stiglitz JE, 'Risk and Global Economic Architecture' (2010) 104 *Journal of Risk and Financial Management* 1.

C. Reports / Institutional Publications

European Commission, *Proposal for a Regulation Laying Down Harmonised Rules on Artificial Intelligence* COM (2021) 206 final.

European Securities and Markets Authority, *Guidelines on Systems and Controls in an Automated Trading Environment for Trading Platforms, Investment Firms and Competent Authorities* (ESMA 2012).

Food and Drug Administration, *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan* (FDA 2021).

International Organization of Securities Commissions, *Artificial Intelligence and Machine Learning in Capital Markets: Use Cases and Challenges* (IOSCO 2021).

Nuffield Council on Bioethics, *Artificial Intelligence (AI) in Healthcare and Research* (Nuffield Council on Bioethics 2023).

OECD, *OECD Artificial Intelligence Papers* (OECD Publishing 2023).

Securities and Exchange Commission and Commodity Futures Trading Commission, *Findings Regarding the Market Events of May 6, 2010* (2010).

World Economic Forum, *A Framework for the Governance of Artificial Intelligence at the Edge* (WEF 2020).

World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: WHO Guidance* (WHO 2021).

D. Legislation / Legal Materials

Directive 2014/65/EU of the European Parliament and of the Council on markets in financial instruments [2014] OJ L173/349.