

Sun, Longguang (2026) *Essays on financial markets: aquaculture finance, energy transition, and intangible governance.* PhD thesis.

https://theses.gla.ac.uk/86362/

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses
https://theses.gla.ac.uk/
research-enlighten@glasgow.ac.uk



Essays on Financial Markets:

Aquaculture Finance, Energy Transition, and Intangible Governance

BY

Longguang Sun

SUBMITTED IN FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

ADAM SMITH BUSINESS SCHOOL

COLLEGE OF SOCIAL SCIENCES

UNIVERSITY OF GLASGOW

MARCH 2026

Abstract

This doctoral dissertation investigates the pricing, management, and mitigation of risk in modern financial markets, traversing the domains of natural resource economics, energy transition finance, and corporate governance. As the global economy enters an era of polycrisis defined by accelerating climate change, structural policy shifts, and complex agency conflicts, traditional valuation frameworks, which are often predicated on assumptions of linearity, stationarity, and limited dimensions, are increasingly insufficient. This research proposes and validates a unified computational and empirical framework to address these challenges, comprising four independent but thematically integrated essays.

The first essay addresses the curse of dimensionality in natural resource valuation by applying Deep Q-Learning to the optimal harvesting problem in aquaculture. Treating the harvesting decision as a real option, the study evaluates Deep Reinforcement Learning as a scalable, nonparametric alternative to regression-based methods such as Least-Squares Monte Carlo. In the baseline setting, the DQN policy produces harvesting times that are broadly comparable with the LSMC benchmark, although the estimated option values remain lower. The main advantage of the framework is therefore its capacity to accommodate richer state spaces without requiring a rapidly expanding set of manually specified basis functions.

The second essay bridges biological modeling and financial engineering to value a novel hedging instrument: the biomass futures contract. By integrating a stochastic biological growth model, which accounts for thermal stress, mortality shocks, and growth rates contingent on temperature, with a Schwartz dual factor pricing model, the study quantifies the hedging demand against production risks driven by climate change. The findings reveal that the economic value of biomass futures is highly dependent on the underlying state. These contracts offer substantial downside protection and fundamentally alter optimal harvesting strategies, primarily under scenarios of extreme environmental volatility.

The third essay develops a physics-informed real options framework for evaluating nuclear power investment over extended horizons under climate-policy uncertainty. A key innovation is the incorporation of a climate tipping point that introduces a structural break into the natural carbon-decay mechanism and, in turn, the carbon-pricing function. Comparing alternative continuation-value approximators shows that tree-based methods, particularly Random Forest and XGBoost, are well suited to representing non-smooth exercise bound-

aries and can produce economically meaningful differences at the margin relative to global polynomial smoothing. The model interprets the additional value associated with the anticipated crisis regime as a “fear premium”, reflecting ambiguity about a potentially irreversible environmental regime shift before the physical threshold is actually crossed.

The fourth essay extends the analysis to corporate finance, examining the role of Organizational Capital as an intangible governance mechanism. Using a comprehensive panel of U.S. firms, the study documents a significant negative association between Organizational Capital and future stock price crash risk. This relation is conditional on the surrounding governance environment and is stronger in several settings characterized by effective monitoring, managerial expertise, and stronger compensation incentives. Together, these essays advance the understanding of how financial agents price and manage complex risks, demonstrating the necessity of integrating advanced computational techniques with contextual physical and organizational insights.

Contents

Abstract iii

List of Figures ix

List of Tables x

Acknowledgements xii

1 Introduction 1

1.1 Background, Scope, and Motivation 1

1.2 Thesis Structure and Contributions 3

1.3 Methods and Techniques: Preliminaries 4

1.4 Theoretical Foundations 4

1.4.1 Stochastic Optimal Control 4

1.4.2 Optimal Stopping Theory 6

1.5 Numerical Methods: Function Approximation Techniques 7

1.5.1 Least-Squares Monte Carlo (LSM) 7

1.5.2 Tree-Based Ensemble Methods 8

2 Applying Machine Learning to Solve the Optimal Stopping Problem in Aquaculture 10

2.1 Abstract 10

2.2 Introduction 10
2.3 Methodology 12
2.3.1 Data Source and Processing 12
2.3.2 Biological and Economic Parameters 12
2.3.3 Stochastic Pricing Framework and Biological Dynamics 13
2.3.4 The Longstaff-Schwartz (LSM) Approach 15
2.4 Deep Q-Learning Framework for Optimal Harvesting 17
2.4.1 Optimal Stopping and the Snell Envelope 17
2.4.2 Markov Decision Process (MDP) Formulation 18
2.4.3 Deep Q-Network (DQN) Algorithm 19
2.4.4 Validation via Dual Upper Bound 21
2.5 Results and Discussion 21
2.5.1 Comparison of Models 21
2.5.2 Scalability in High-Dimensional Environments 22
2.5.3 Parameter Sensitivity Analysis 23
2.6 Conclusion 24

3 Biomass Futures Contracts for the Aquaculture Sector: An Exploration 25

3.1 Introduction 25
3.2 Model 29
3.2.1 Fish Growth 29
3.2.2 Biomass Futures Contracts 34
3.2.3 Salmon Futures 36
3.2.4 Harvesting 37

3.2.5 Estimation of the Continuation Value Using the Longstaff-Schwartz Approach . . . 38
3.2.6 Parameters . . . 40
3.3 Results . . . 42
3.4 Conclusion . . . 48

4 Valuation of Nuclear Power Investment with Real Options Theory 49

4.1 Introduction . . . 49
4.2 Model Description and Parameter Settings . . . 52
4.2.1 Modelling Uncertainty Factors in Nuclear Power Investment . . . 52
4.3 Model Solution . . . 62
4.3.1 Modified Longstaff–Schwartz Method (LSM) . . . 65
4.3.2 Machine Learning-Enhanced Regression Schemes . . . 67
4.3.3 Implementation and Comparative Framework . . . 69
4.4 Model Calibration . . . 70
4.5 Model Results and Discussion . . . 72
4.6 Conclusion . . . 76

5 Organizational Capital and Stock Price Crash Risk 78

5.1 Introduction . . . 78
5.2 Literature Review . . . 82
5.2.1 Organizational Capital . . . 82
5.2.2 Stock Price Crash Risk . . . 82
5.2.3 Theoretical Underpinning and Hypotheses Development . . . 82
5.3 Research methodology . . . 86
5.3.1 Sample and data . . . 86

5.3.2 Organizational Capital 87
5.3.3 Stock Price Crash Risk 88
5.3.4 Control Variable 89
5.3.5 Empirical models 90
5.4 Empirical analyses 90
5.4.1 Descriptive Statistics 90
5.4.2 Regression Analysis 91
5.4.3 Organizational Capital and Stock Price Crash Risk: Agency Theory Perspectives 98
5.4.4 Exclusion of computers, software, and electronic equipment firms (business equipment industry) 104
5.4.5 Alternative Measures of Organizational Capital 104
5.4.6 An IV approach with heteroskedasticity-based instruments 107
5.4.7 Heckman selection IV model 108
5.4.8 Propensity Score Matching (PSM) 110
5.4.9 Entropy Balancing 112
5.4.10 Evidence from the machine learning approach 113
5.5 Conclusion 114
5.6 Appendix 115
5.6.1 Appendix A. Variable definitions 115
5.6.2 Appendix B: Machine Learning XGBoost 117

6 Conclusion **119**

Bibliography **121**

List of Figures

3.1 Simulated water temperature conditions 42

3.2 Optimal Harvesting Points under Ideal Growth Conditions 44

3.3 Optimal Harvesting Points under Thermal Stress Conditions 45

3.4 Optimal Harvesting Points under Extreme Temperature Conditions 45

3.5 Harvesting Surface under Ideal Growth Conditions 46

3.6 Harvesting Surface under Thermal Stress Conditions 46

3.7 Harvesting Surface under Extreme Temperature Conditions 47

4.1 Sensitivity of Nuclear Option Value to Volatility and Discount Rate 75

5.1 Relative Variable Importance Rankings 118

List of Tables

2.1 Relevant Parameters for Fish Farming 13

2.2 Average Harvesting Time and Real Option Value Comparison 22

2.3 Sensitivity Analysis: Mortality Rate Parameters 23

3.1 Schwartz (1997) Two-Factor Model Parameters (Aligned with Simulation) . 40

3.2 Biological Model and Futures Contract Parameters 41

3.3 Real Option Value and Hedging Effectiveness Across Temperature Scenarios (Calibrated) 43

3.4 Mean Optimal Stopping Time and Classification Accuracy (Calibrated) . . 47

4.1 Calibration of Model Parameters for Different Nuclear Plant Scales 71

4.2 Comparison of Real Option Valuation Results (Extended Horizon) 73

5.1 **Summary Statistics** 91

5.2 **Baseline Regression: Stock Price and Organizational Capital** 92

5.3 **Mechanism of CEO Experience** 95

5.4 **Mechanism of CEO Compensation** 97

5.5 **Stock Price and Organizational Capital: Institutional Investors' Orientation** 99

5.6 **Stock Price and Organizational Capital: Analyst Monitoring** 99

5.7 **Stock Price and Organizational Capital: Financial Reporting Transparency** 101

5.8 **Stock Price and Organizational Capital: Managerial Risk Perception** . 103

5.9 **Exclusion of computers, software, and electronic equipment firms** . . . 104

5.10 **Alternative Measure: Organizational Capital** . . . 106

5.11 **Endogeneity Test: Lewbel (2012)** . . . 108

5.12 **Heckman Two-Stage Results** . . . 110

5.13 **Robustness Check: Propensity Score Matching (PSM)** . . . 111

5.14 **Robustness Check: Entropy Balancing** . . . 113

Acknowledgements

The journey of pursuing a PhD is rarely a solitary endeavor, and this dissertation would not have been possible without the support, guidance, and encouragement I have received along the way.

First and foremost, I would like to express my deepest gratitude to my supervisory team: Professor Christian Ewald, Professor Bart Taub, and Professor Georgios Sermpinis. Your patience, intellectual generosity, and insightful feedback have profoundly shaped my understanding of financial markets, quantitative modeling, and corporate governance. I am especially grateful for the freedom you gave me to explore diverse research ideas and for your guidance whenever I encountered theoretical or computational challenges.

I am deeply thankful to my fellow PhD classmates and friends at the Adam Smith Business School. The discussions, shared challenges, and moments of relief we experienced together made this demanding academic journey much more enjoyable.

A special word of thanks goes to my roommate and close friend, Li Chen, for your patience, companionship, and support during some of the most demanding stages of my PhD. I would also like to thank my dear friends Shuwen Tan, Tongtong Wang, Xihao Song, Xiaoqin Wei, and Leihan He for all the laughter, good times, and memorable moments we shared. Your friendship brought much happiness and balance to my life outside academia.

Lastly, and most importantly, I am deeply grateful to my parents and elders for their unconditional love, encouragement, and quiet sacrifices. Their support has been the foundation of my life and education and has enabled me to pursue this journey. This thesis is dedicated to them.

Author's Declaration

I declare that, except where explicit reference is made to the contribution of others, this dissertation is the result of my own work and has not been submitted for any other degree at the University of Glasgow or any other institution.

Printed Name: Longguang Sun

Signature: _____

Chapter 1

Introduction

1.1 Background, Scope, and Motivation

In today's financial and economic environment, uncertainty has become the norm. These uncertainties come not only from physical and environmental changes, but also from policy and institutional shifts, and corporate governance structures. In response to this complexity, the goal of this research is to develop a unified and scalable computational framework that can be applied across different fields to address dynamic valuation and risk management problems. The central idea of the framework is to integrate Stochastic Optimal Control Theory with modern Computational Economics techniques, in order to provide more robust and flexible decision support under uncertainty. Although the four essays in this thesis are independent, they are closely connected. Together, they show how advanced machine learning methods, such as Deep Reinforcement Learning and ensemble tree models, can improve traditional real options valuation and tail risk management, especially when dealing with high-dimensional state spaces, non-smooth structures, and intangible assets that are difficult to observe.

The first essay (Chapter 2) studies the optimal stopping problem in resource management, using the harvesting decision in aquaculture as a case study. At its core, the harvesting decision is a real option in which operators must determine the best timing under stochastic fluctuations in commodity prices and convenience yields (Ewald et al., 2017). Traditional numerical methods, such as the Least-Squares Monte Carlo (LSM) approach with polynomial basis functions, often encounter the curse of dimensionality and over-smoothing when facing complex and high-dimensional state variables or strongly nonlinear continuation value functions (Becker et al., 2019). To address these limitations, this research introduces Deep Q-Learning (DQL), a model-free Deep Reinforcement Learning technique for solving high-dimensional real option problems. DQL uses deep neural networks as universal function approximators to learn a scalable, nonparametric representation of the continuation value function. Numerical comparisons show that DQL can match the accuracy of the standard

Longstaff–Schwartz two-factor method while avoiding the dimensionality and smoothing issues that commonly affect polynomial regression.

The second essay (Chapter 3) extends the dynamic valuation framework by incorporating explicit environmental and biological uncertainties to address the unique risks faced by the aquaculture industry (Ewald et al., 2017). In this setting, the model links stochastic water temperature, modeled by an Ornstein–Uhlenbeck process, with fish growth and mortality, modeled by a CIR-type process, to describe the evolution of biomass $B(t)$ in a realistic and data-driven way. By modeling the physical production process directly, this chapter provides a rigorous approach to quantifying climate-induced production risks and evaluating the intrinsic value of a new financial instrument: the Biomass Futures Contract. In this chapter, the biomass futures contract is interpreted as a real option designed to hedge production uncertainty triggered by environmental shocks (such as thermal stress), rather than simply price volatility. The results show that the economic benefit of hedging increases substantially under adverse environmental conditions. Using the Longstaff–Schwartz approach to estimate continuation values (Longstaff and Schwartz, 2001), the analysis builds a direct connection between biological uncertainty and market-based hedging strategies, offering new financial engineering insights for risk management in agriculture and aquaculture. Under extreme temperature conditions, the results indicate a clear behavioral response: agents tend to harvest earlier and lock in profits even when biomass levels are low, in order to avoid further biological deterioration. This behavior highlights the interaction between the financial option value and biological constraints, and demonstrates how climate risk fundamentally alters optimal production and hedging decisions.

The third essay (Chapter 4) shifts the focus to macroeconomic and policy risks by developing a real options model for long-term nuclear power investment. A key innovation is the introduction of a climate tipping-point mechanism. The natural decay rate of atmospheric pollution, denoted by $\delta(M_t)$, changes once the pollutant concentration M_t exceeds a critical threshold M^*. This creates a non-smooth carbon-pricing environment that is difficult to represent with global polynomial approximations. The chapter therefore compares polynomial LSM with tree-based and threshold-based continuation-value estimators. Random Forest and XGBoost are naturally suited to partitioning the state space around a hard threshold and produce distinct valuation and timing implications near the crisis regime. Economically, the model interprets the additional option value associated with anticipation of the crisis regime as a "fear premium", reflecting the value of waiting when the timing and consequences of a potentially irreversible environmental shift are uncertain.

The fourth essay (Chapter 5) extends the analysis to corporate finance by examining the association between organizational capital and corporate tail risk, specifically future stock price crash risk (Kim et al., 2011b; Chang et al., 2017; Callen and Fang, 2015). Grounded in information asymmetry theory and the Resource-Based View, the chapter develops the argument that stronger organizational routines, internal processes, and information systems may

be associated with less managerial bad-news hoarding (Jin and Myers, 2006). Empirically, a one-standard-deviation increase in organizational capital is associated with approximately 0.040 lower NCSKEW, equivalent to about 19.5% of the sample mean. A key insight is that this negative association is not uniform across firms. It is stronger in several settings characterized by stronger monitoring, managerial expertise, and compensation incentives (Andreou et al., 2016; Kim et al., 2016). These cross-sectional patterns are consistent with the proposed governance mechanisms but are interpreted as observational rather than causal evidence.

Taken together, the fourth essay shows that the relationship between organizational capital and extreme downside risk is highly context-dependent. Organizational capital is associated with lower crash risk primarily in governance environments that provide stronger monitoring and incentive alignment. The chapter therefore treats OC as a potentially important governance-related intangible asset while remaining cautious about causal interpretation.

1.2 Thesis Structure and Contributions

This thesis contains an introductory chapter, four self-contained essay chapters, and a concluding chapter. The first essay (Chapter 2) examines the optimal fish-harvesting decision for aquaculturists as an optimal stopping problem under stochastic commodity prices and convenience yields. It evaluates Deep Q-Learning (DQL) as a nonparametric alternative to regression-based continuation-value approximation and emphasizes the scalability of neural-network methods when the state space becomes richer. The second essay (Chapter 3) incorporates stochastic fish growth, mortality, and temperature dynamics to quantify hedging demand against climate-induced production risk. Its principal contribution is the valuation of biomass futures as a financial instrument for managing biological risk, showing that the relative value of hedging improves as environmental conditions deteriorate and becomes positive in the extreme-stress calibration. The third essay (Chapter 4) develops a long-horizon real options model for nuclear power investment under climate-policy uncertainty. A central methodological contribution is a physics-informed state space with a climate tipping point that produces a non-smooth valuation problem; tree-based methods such as Random Forest and XGBoost are used alongside polynomial and threshold-regression benchmarks to evaluate the implications of that non-smoothness. The fourth essay (Chapter 5) investigates whether organizational capital mitigates stock price crash risk. The evidence indicates a negative association between organizational capital and future crash risk, with the relation varying across governance, monitoring, transparency, and managerial-incentive environments.

1.3 Methods and Techniques: Preliminaries

This section details the advanced computational and mathematical frameworks employed to address the complex dynamic stochastic control problems examined in this thesis. The discussion begins with the theoretical formulation of stochastic optimization and proceeds to the specific numerical approximation techniques, ranging from regression-based methods to advanced machine learning algorithms.

1.4 Theoretical Foundations

This study relies on the fundamental principles of dynamic optimization under uncertainty. The following subsections define the rigorous mathematical structures for stochastic control and optimal stopping problems, providing the necessary theoretical basis for the numerical methods employed in subsequent chapters.

1.4.1 Stochastic Optimal Control

We consider a finite time horizon $T < \infty$. Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a complete probability space equipped with a filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}$ satisfying the usual conditions (i.e., right-continuous and containing all $\mathbb{P}$-null sets).

State Dynamics and Admissible Controls

The dynamics of the system state, denoted by $X(t) \in \mathbb{R}^n$, are governed by the following Stochastic Differential Equation (SDE) (Øksendal, 2003):

$$dX(t) = \mu(t, X(t), u(t))dt + \sigma(t, X(t), u(t))dW(t), \quad X(0) = x_0, \tag{1.1}$$

where $W(t)$ is a d-dimensional standard Brownian motion adapted to $\mathbb{F}_t$. The functions $\mu : [0,T] \times \mathbb{R}^n \times \mathcal{U} \to \mathbb{R}^n$ and $\sigma : [0,T] \times \mathbb{R}^n \times \mathcal{U} \to \mathbb{R}^{n\times d}$ represent the drift and diffusion coefficients, respectively. We assume μ and σ satisfy the Lipschitz and linear growth conditions to ensure the existence and uniqueness of a strong solution to the SDE.

The process $u(t)$ represents the control strategy chosen by the decision-maker. It takes values in a compact metric space $\mathcal{U} \subset \mathbb{R}^k$.

Admissible Control. An $\mathcal{F}_t$-adapted process $u(\cdot)$ is said to be an *admissible control*, denoted

as $u \in \mathcal{A}$, if it satisfies the integrability condition:

$$\mathbb{E}\left[\int_0^T |u(t)|^2 dt\right] < \infty. \tag{1.2}$$

Objective Function and Value Function

The objective of the decision-maker is to maximize a performance criterion $J(t, x; u)$, defined as the expected cumulative reward:

$$J(t, x; u) = \mathbb{E}_{t,x}\left[\int_t^T f(s, X(s), u(s))ds + g(X(T))\right], \tag{1.3}$$

where $f(\cdot)$ is the continuous instantaneous reward rate and $g(\cdot)$ is the terminal payoff function. The value function $V(t, x)$ is defined as the supremum of the expected payoff over the set of all admissible controls $\mathcal{A}$:

$$V(t, x) = \sup_{u \in \mathcal{A}} J(t, x; u). \tag{1.4}$$

The Hamilton-Jacobi-Bellman (HJB) Equation

To solve for the value function, we employ the Dynamic Programming Principle (DPP). Associated with the controlled diffusion process is the second-order differential operator $\mathcal{L}^u$, known as the infinitesimal generator, acting on a function $\phi \in C^{1,2}([0, T] \times \mathbb{R}^n)$:

$$\mathcal{L}^u \phi(t, x) = \frac{\partial \phi}{\partial t} + \sum_{i=1}^n \mu_i(t, x, u)\frac{\partial \phi}{\partial x_i} + \frac{1}{2}\sum_{i,j=1}^n (\sigma\sigma^\top)_{ij}(t, x, u)\frac{\partial^2 \phi}{\partial x_i \partial x_j}. \tag{1.5}$$

According to the DPP, the value function $V(t, x)$ satisfies the Hamilton-Jacobi-Bellman (HJB) equation, a nonlinear Partial Differential Equation (PDE):

$$\sup_{u \in \mathcal{U}} \{\mathcal{L}^u V(t, x) + f(t, x, u)\} = 0, \quad \text{for } (t, x) \in [0, T) \times \mathbb{R}^n, \tag{1.6}$$

subject to the terminal condition $V(T, x) = g(x)$.

Verification Theorem. Under suitable smoothness conditions, if there exists a function $\phi \in C^{1,2}$ that satisfies the HJB equation and the boundary conditions, then $\phi(t, x) = V(t, x)$, and the control u^* that achieves the supremum in the HJB equation is the optimal control (Bertsekas, 2005).

1.4.2 Optimal Stopping Theory

Optimal stopping problems constitute a specific class of stochastic control where the control variable is the timing of a single, irreversible decision (Peskir and Shiryaev, 2006).

Problem Formulation

Let $X(t)$ be the underlying stochastic process defined on $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$. The problem is to identify an optimal stopping time τ^* that maximizes the expected discounted payoff:

$$V(x) = \sup_{\tau \in \mathcal{T}} \mathbb{E}_x \left[e^{-r\tau} g(X(\tau)) \right], \tag{1.7}$$

where $\mathcal{T}$ is the set of all $\mathbb{F}$-stopping times (i.e., $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$), $r > 0$ is the discount rate, and $g(\cdot)$ is the gain function (payoff upon exercise).

Variational Inequalities and Free Boundary Problems

Unlike standard control problems where the HJB equation is an equality, optimal stopping problems lead to a *Variational Inequality* (VI). The value function $V(x)$ must satisfy the following conditions involving the infinitesimal generator $\mathcal{L}$ of the process X:

$$\max \{(\mathcal{L} - r)V(x),\ g(x) - V(x)\} = 0. \tag{1.8}$$

This formulation implies that at any state x, one of two conditions must hold:

1. **Continuation Region ($\mathcal{C}$):** $V(x) > g(x)$. In this region, it is optimal to wait. The value function satisfies the differential equation $(\mathcal{L} - r)V(x) = 0$.
2. **Stopping Region ($\mathcal{S}$):** $V(x) = g(x)$. In this region, it is optimal to exercise the option immediately.

The boundary separating $\mathcal{C}$ and $\mathcal{S}$ is the *optimal exercise boundary* or *free boundary*. The optimal stopping time is characterized as the first entry time of the process $X(t)$ into the stopping region:

$$\tau^* = \inf \{t \geq 0 : X(t) \in \mathcal{S}\}. \tag{1.9}$$

Mathematically, the solution must also satisfy the *smooth pasting condition* ($\nabla V(x) = \nabla g(x)$) at the free boundary to ensure the value function is continuously differentiable.

1.5 Numerical Methods: Function Approximation Techniques

Analytical solutions to the high-dimensional stochastic control problems formulated above are generally intractable, particularly when the underlying dynamics exhibit regime-switching or non-smooth boundaries. Consequently, this thesis employs advanced numerical approximation methods to estimate the value function (or continuation value).

Mathematically, the problem is reduced to finding an approximation function $\widehat{V}$ from a hypothesis space $\mathcal{H}$ that minimizes the expected error against the true value function V:

$$\widehat{V} = \arg\min_{f \in \mathcal{H}} \mathbb{E}\left[\mathcal{L}\left(V(X), f(X)\right)\right], \tag{1.10}$$

where $\mathcal{L}$ is a loss function (typically squared error). The choice of $\mathcal{H}$ and the optimization algorithm distinguishes the methods discussed below.

1.5.1 Least-Squares Monte Carlo (LSM)

The Least-Squares Monte Carlo (LSM) method, proposed by Longstaff and Schwartz (2001), combines Monte Carlo simulation with dynamic programming. It approximates the conditional expectation of future payoffs (the continuation value) by projecting them onto a finite set of basis functions.

Continuation Value. Let t_n be a discrete time step. The continuation value $C(t_n, X_{t_n})$ is defined as the risk-neutral expectation of the optimal future payoff:

$$C(t_n, X_{t_n}) = \mathbb{E}^{\mathbb{Q}}\left[e^{-r\Delta t} V(t_{n+1}, X_{t_{n+1}}) \mid \mathcal{F}_{t_n}\right]. \tag{1.11}$$

Orthogonal Projection. The LSM method assumes that the continuation value lies within the span of a set of K basis functions $\{\phi_k\}_{k=1}^{K}$. The approximation $\widehat{C}(t_n, X)$ is given by the linear combination:

$$\widehat{C}(t_n, X) = \sum_{k=1}^{K} \beta_k \phi_k(X), \tag{1.12}$$

where the vector of coefficients $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_K)^\top$ is estimated via Ordinary Least Squares (OLS):

$$\widehat{\boldsymbol{\beta}} = \arg\min_{\boldsymbol{\beta} \in \mathbb{R}^K} \sum_{m=1}^{M} \left(Y^{(m)} - \sum_{k=1}^{K} \beta_k \phi_k(X_{t_n}^{(m)}) \right)^2. \tag{1.13}$$

Here, M is the number of simulation paths, and $Y^{(m)}$ is the realized discounted cash flow for the m-th path. While computationally efficient, LSM is limited by the choice of basis functions (typically polynomials), which may suffer from the Runge phenomenon or fail to capture sharp discontinuities (e.g., tipping points) in the value surface.

1.5.2 Tree-Based Ensemble Methods

To address the limitations of global polynomial approximation, this study employs nonparametric tree-based ensemble methods. These algorithms partition the feature space $\mathbb{R}^d$ into disjoint rectangular regions and model the response as a constant within each region.

Regression Tree. A regression tree $T(x;\Theta)$ partitions the state space into J terminal regions (leaves) $\{R_j\}_{j=1}^J$ with associated constants $\{w_j\}_{j=1}^J$:

$$T(x;\Theta) = \sum_{j=1}^{J} w_j \mathbb{I}(x \in R_j), \tag{1.14}$$

where $\Theta = \{R_j, w_j\}_{j=1}^J$ represents the tree parameters, and $\mathbb{I}(\cdot)$ is the indicator function.

The parameters are estimated by minimizing the empirical risk, typically defined as the sum of squared errors over the training data (x_i, y_i) for $i = 1, \dots, N$:

$$\min_{\Theta} \sum_{i=1}^{N} \left(y_i - \sum_{j=1}^{J} w_j \mathbb{I}(x_i \in R_j) \right)^2. \tag{1.15}$$

For a fixed partition structure $\{R_j\}_{j=1}^J$, the optimal constant w_j is simply the sample mean of the response variable within that region:

$$\hat{w}_j = \text{ave}(y_i \mid x_i \in R_j) = \frac{1}{N_j} \sum_{x_i \in R_j} y_i, \tag{1.16}$$

where N_j is the number of sample points falling into region R_j. The regions themselves are determined using a greedy recursive binary splitting algorithm that iteratively selects the splitting variable and split point to achieve the maximum reduction in the error criterion.

Random Forest

Random Forest, introduced by Breiman (2001), acts as a variance-reduction technique for decision trees. It employs *bagging* (Bootstrap Aggregating) combined with random feature selection to construct an ensemble of uncorrelated trees.

Let $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^N$ be the training set at time step t_n. The Random Forest estimator is defined as:

$$\widehat{V}_{RF}(x) = \frac{1}{B} \sum_{b=1}^{B} T_b(x;\Theta_b), \tag{1.17}$$

where each tree T_b is trained on a bootstrap sample of $\mathcal{D}_n$. The parameter vector Θ_b captures the randomness in both sample selection and the split-candidate set (a random subset of features is considered at each node split). By averaging B decorrelated trees, the method reduces

the variance of the estimator without increasing the bias, making it robust against overfitting in high-dimensional spaces (Hastie et al., 2009).

eXtreme Gradient Boosting (XGBoost)

XGBoost (Chen and Guestrin, 2016) is a scalable implementation of the Gradient Boosting Decision Tree (GBDT) framework. Unlike Random Forest, which builds trees in parallel, XGBoost builds an additive model sequentially to minimize a regularized objective function.

At iteration k, the model seeks a new function $f_k \in \mathcal{H}$ (a regression tree) to add to the existing ensemble $\hat{y}_i^{(k-1)}$:

$$\hat{y}_i^{(k)} = \hat{y}_i^{(k-1)} + f_k(x_i). \tag{1.18}$$

The objective function at step k is given by:

$$\text{Obj}^{(k)} = \sum_{i=1}^{N} l(y_i, \hat{y}_i^{(k-1)} + f_k(x_i)) + \Omega(f_k), \tag{1.19}$$

where l is a differentiable convex loss function, and $\Omega(f_k)$ is a regularization term penalizing model complexity:

$$\Omega(f) = \gamma T + \frac{1}{2}\lambda ||w||^2, \tag{1.20}$$

with T denoting the number of leaves and w the vector of leaf scores.

Second-Order Approximation. A key innovation of XGBoost is the use of the second-order Taylor expansion to approximate the loss function. The objective is simplified to:

$$\text{Obj}^{(k)} \approx \sum_{j=1}^{T} \left[\left(\sum_{i \in I_j} g_i \right) w_j + \frac{1}{2} \left(\sum_{i \in I_j} h_i + \lambda \right) w_j^2 \right] + \gamma T, \tag{1.21}$$

where $g_i = \partial_{\hat{y}^{(k-1)}} l(y_i, \hat{y}^{(k-1)})$ and $h_i = \partial^2_{\hat{y}^{(k-1)}} l(y_i, \hat{y}^{(k-1)})$ are the first and second-order gradients. The optimal weight w_j^* for leaf j has a closed-form solution:

$$w_j^* = -\frac{\sum_{i \in I_j} g_i}{\sum_{i \in I_j} h_i + \lambda}. \tag{1.22}$$

This formulation allows XGBoost to efficiently handle complex nonlinear boundaries by directly optimizing the structure of the tree based on the gradient and Hessian of the loss.

Chapter 2

Applying Machine Learning to Solve the Optimal Stopping Problem in Aquaculture

2.1 Abstract

This study presents an optimal fish-harvesting model for aquaculture based on the stochastic pricing framework of Ewald et al. (2017). We employ Deep Reinforcement Learning (DRL) to approximate the harvesting policy, using a neural-network framework inspired by Becker et al. (2020) and Ery and Michel (2021). Unlike the traditional Least-Squares Monte Carlo (LSMC) method, which depends on pre-specified basis functions, the model-free Deep Q-Network (DQN) architecture (Mnih et al., 2015) approximates state-action values directly. In the baseline calibration, DQN produces average harvesting times that are broadly comparable with the LSMC benchmark, although its option-value estimates are lower. The principal motivation for the DQN framework is therefore scalability: additional biological or environmental state variables can be incorporated without manually constructing an increasingly large polynomial basis.

2.2 Introduction

Fish Pool ASA, established in 2006 in Bergen, Norway, operates as a regulated marketplace for trading financial derivatives in salmon contracts, including a significant volume of fresh farmed salmon futures and options. The existing literature on salmon farming and aquaculture economics, such as Arnason (1992), Yu and Leung (2006), and Guttormsen (2008), provides extensive research into optimal feeding strategies and harvesting plans. While these studies lay important groundwork, they largely overlook the inherent price volatility in the salmon market. Although some attention has been given to price fluctuation, for example, Forsberg and Guttormsen (2006) propose a stochastic model for timber prices subject to random fluctuations, these models often fail to fully reflect the realities of dynamic market conditions,

particularly for perishable commodities like salmon.

While stochastic pricing models have been applied in forestry, such as the work by Chen et al. (2012) using lumber futures, these models often employ simplified "long-term" approximations that are suitable for long rotation periods. However, recognizing that salmon farming cycles are significantly shorter than forestry rotations, Ewald et al. (2017) argued that such approximations are insufficient. Instead, they implemented the full Schwartz (1997) two-factor model to capture the complex mean-reversion dynamics between spot prices and convenience yields.

Ewald et al. (2017) further integrated this pricing framework into an optimal harvesting problem, utilizing the Least-Squares Monte Carlo (LSMC) approach to value the "real option" of the fish farm. Their findings demonstrated that neglecting the futures market or managerial flexibility could lead to substantial undervaluation of the farm, particularly for risk-averse farmers. However, a critical limitation of the LSMC approach lies in its dependence on the regression of basis functions. As the realism of the biological model increases, the state space expands and the number of required basis functions can grow rapidly, particularly when stochastic mortality rates, variable feed costs, or environmental factors such as water temperature are introduced. This "curse of dimensionality" often forces researchers to simplify biological realities to fit computational constraints. This limitation motivates our application of Deep Reinforcement Learning (DRL), which bypasses the need for basis function selection and offers a scalable solution for approximating continuation values in high-dimensional environments.

Optimal stopping models are widely used across various domains. Ajdari et al. (2019), for instance, applied an optimal stopping model to determine when to stop radiotherapy treatments. Liyanage et al. (2019) leveraged optimal stopping in civic engagement platforms, while Dai et al. (2019) used Bayesian optimization coupled with optimal stopping for training neural networks. In financial engineering, Bertsekas and Tsitsiklis (1996) introduced Approximate Dynamic Programming (ADP) as a powerful tool for solving such stochastic control problems. The LSMC method developed by Longstaff and Schwartz (2001) became a benchmark for solving American option pricing problems.

When transition dynamics are unknown or complex, model-free learning methods such as Reinforcement Learning (RL) offer a promising alternative. Sutton and Barto (2018) introduced RL as a method to maximize cumulative reward through interaction with an environment. Goodfellow et al. (2017) emphasized the power of deep learning in extracting complex features from data. The combination of these fields, Deep Reinforcement Learning (DRL), has been explored in financial contexts by Li et al. (2009), Kohler et al. (2010), and more recently by Fathan and Delage (2021) and Chen et al. (2019). Becker et al. (2020) and Hu (2019) have further demonstrated the efficacy of neural networks in solving high-dimensional optimal stopping problems.

In this study, we aim to bridge the gap between advanced biological realism and computational feasibility in aquaculture economics. By leveraging the approximation power of deep neural networks within an RL framework, we develop a decision-support model that is not only robust in current low-dimensional settings but also capable of scaling to complex future scenarios. Specifically, we contribute to the literature by: (1) formulating the salmon harvesting problem as a Markov Decision Process (MDP) solvable by a Deep Q-Network (DQN) agent; (2) demonstrating that the DRL approach achieves valuation accuracy comparable to the LSMC benchmark without manual feature engineering; and (3) analyzing the learned policy structure to validate the agent's ability to capture economic rationale in a model-free manner.

2.3 Methodology

2.3.1 Data Source and Processing

Our dataset consists of 1,496 daily observations of salmon futures prices obtained from Fish Pool ASA in Bergen, Norway. The sample spans June 12, 2006, to March 22, 2012, matching the observation window used by Ewald et al. (2017). Using the same period allows the baseline pricing parameters, including those estimated through the Kalman filter, to remain comparable with that benchmark. The contracts are monthly salmon futures settled against the Fish Pool Index (FPI), a weekly spot-price index for fresh Atlantic salmon.

Following Ewald et al. (2017), we organize the observations into four maturity panels:

- **Panel A (Short-term):** maturities of 1–3 months;
- **Panel B (Medium-term):** maturities of 4–6 months;
- **Panel C (Long-term):** maturities of 7–9 months; and
- **Panel D (Mixed-term):** contracts with a range of maturities used to assess robustness across the term structure.

2.3.2 Biological and Economic Parameters

The parameters governing the fish farming process, summarized in Table 2.1, are derived from the empirical findings of Asche and Bjorndal (2011). Their work provides a comprehensive analysis of cost structures and production efficiency in the Norwegian salmon industry, and the selected values reflect a representative efficient firm within the sector.

Biologically, the model assumes an annual mortality rate of 0.1, reflecting natural population losses due to disease and environmental factors consistent with industry averages during the sample period. The production cycle is defined by a maximum time horizon of 3 years, with an initial recruitment of 10,000 fish and an asymptotic weight of 6 kg, which sets the upper biological bound for the growth model. In terms of production efficiency, the model utilizes a Feed Conversion Ratio (FCR) of 1.1. This ratio indicates that 1.1 kilograms of feed are required to produce 1 kilogram of salmon biomass, a figure characteristic of highly efficient modern aquaculture operations.

Economically, the cost parameters are calibrated to match the operational expenditure levels reported in Asche and Bjorndal (2011). The variable harvesting cost is set at 3 NOK per kg. The variable feeding cost, which constitutes the largest component of marginal production costs driven by global commodity prices for fishmeal and soy, is estimated at 7 NOK per kg per year.

Table 2.1: Relevant Parameters for Fish Farming

Parameters	**Value**
Mortality rate	0.1
Conversion rate	1.1
Number of recruits	10,000
Time horizon (years)	3
Asymptotic Weight (kg)	6
Variable harvesting cost per kg (NOK)	3
Variable feeding cost per kg per year (NOK)	7

2.3.3 Stochastic Pricing Framework and Biological Dynamics

To model the stochastic nature of salmon spot prices, we employ the two-factor framework originally proposed by Schwartz (1997) and adapted for aquaculture by Ewald et al. (2017). This framework is particularly suitable for commodity markets as it accounts for the mean-reverting property of prices and the stochastic behavior of the convenience yield.

Let $P(t)$ denote the spot price of salmon at time t, and $\delta(t)$ represent the instantaneous convenience yield. Under the risk-neutral measure $\mathbb{Q}$, the joint dynamics of the spot price and convenience yield are governed by the following system of stochastic differential equations (SDEs):

$$dP(t) = (r - \delta(t))P(t)dt + \sigma_1 P(t)d\tilde{Z}_1(t), \tag{2.1}$$

$$d\delta(t) = [\kappa(\alpha - \delta(t)) - \lambda]dt + \sigma_2 d\tilde{Z}_2(t). \tag{2.2}$$

Here, r represents the risk-free interest rate. The convenience yield $\delta(t)$ follows an Ornstein-Uhlenbeck process, where κ denotes the speed of mean reversion, α is the long-run mean of the convenience yield, and λ represents the market price of convenience yield risk. The terms σ_1 and σ_2 capture the volatility of the spot price and convenience yield, respectively. The stochastic innovations are driven by two standard Brownian motions, $\tilde{Z}_1(t)$ and $\tilde{Z}_2(t)$, which are correlated with a coefficient ρ:

$$d\tilde{Z}_1(t)d\tilde{Z}_2(t) = \rho dt. \tag{2.3}$$

A key advantage of the Schwartz two-factor model is its analytical tractability for pricing derivative contracts. Drawing on the derivations by Jamshidian and Fein (1990) and Bjerksund (1991), the theoretical price of a futures contract $F(P, \delta, T)$ with maturity T can be expressed in closed form as:

$$F(P, \delta, T) = \mathbb{E}_{\mathbb{Q}}[P(T) \mid \mathcal{F}_t] = P(t) \cdot \exp\left(-\delta(t)\frac{1 - e^{-\kappa T}}{\kappa} + A(T)\right), \tag{2.4}$$

where the deterministic term $A(T)$ accounts for the convexity adjustment and risk premia, defined as:

$$\begin{aligned} A(T) = &\left(r - \alpha + \frac{\lambda}{\kappa} + \frac{1}{2}\frac{\sigma_2^2}{\kappa^2} - \frac{\sigma_1\sigma_2\rho}{\kappa}\right)T + \frac{\sigma_2^2}{4}\left(\frac{1 - e^{-2\kappa T}}{\kappa^3}\right) \\ &+ \left(\alpha\kappa - \lambda + \sigma_1\sigma_2\rho - \frac{\sigma_2^2}{\kappa}\right)\left(\frac{1 - e^{-\kappa T}}{\kappa^2}\right). \end{aligned} \tag{2.5}$$

The model parameters are typically calibrated to historical futures data using the Kalman filter technique, allowing for the estimation of latent state variables.

To fully capture the value of the aquaculture operation, the financial model must be integrated with biological growth dynamics. We assume the total biomass $X(t)$ is the product of the number of fish $n(t)$ and the average weight per fish $w(t)$. The growth of an individual fish is modeled using the von Bertalanffy growth function, which creates a concave growth trajectory approaching an asymptotic limit:

$$w(t) = w_\infty \left(a - be^{-ct}\right)^3, \tag{2.6}$$

where w_∞ is the asymptotic maximum weight, and a, b, and c are biological constants determining the growth rate and initial conditions.

Simultaneously, the fish population is subject to natural mortality. The population dynamics follow a deterministic decay process:

$$dn(t) = -m(t) \cdot n(t)dt, \tag{2.7}$$

where $m(t)$ represents the instantaneous mortality rate. Consequently, the total harvestable biomass at any time t is given by $X(t) = n(t)w(t)$.

The farmer's objective is to maximize the expected present value of the profit. The instantaneous profit function upon harvesting at time t is defined as the revenue from biomass sales minus the variable harvesting costs:

$$H(t, P(t), X(t)) = P(t)X(t) - CH(t), \tag{2.8}$$

where $CH(t)$ denotes the total cost associated with the harvesting process. Furthermore, the operation incurs continuous feeding and maintenance costs, denoted by $CF(t)$, which accumulate over the farming cycle.

The decision to harvest can be formulated as an optimal stopping problem. The fish farmer seeks an optimal stopping time τ that maximizes the expected discounted cumulative profit under the risk-neutral measure $\mathbb{Q}$:

$$V_0 = \max_{\tau} \mathbb{E}_{\mathbb{Q}} \left[e^{-r\tau} \big(P(\tau)X(\tau) - CH(\tau) \big) - \int_0^{\tau} e^{-rt} CF(t) dt \right]. \tag{2.9}$$

This formulation captures the trade-off between the potential for biological growth and price appreciation versus the risks of mortality, price decline, and accumulating feeding costs.

2.3.4 The Longstaff-Schwartz (LSM) Approach

To serve as a benchmark for our Deep Learning approach, we employ the Least Squares Monte Carlo (LSM) method proposed by Longstaff and Schwartz (2001). This technique is the standard numerical procedure for solving optimal stopping problems, often referred to as American-style options, when the dimensionality of the state space renders grid-based methods computationally expensive or infeasible. In the context of the salmon harvesting problem described in Ewald et al. (2017), the decision to harvest is analogous to exercising an American option. At any discrete time step t_n (where n ranges from 0 to N), the fish farmer faces a binary choice: harvest immediately to realize profit $H(t_n)$, or continue farming to benefit from potential biological growth and price appreciation.

Following the Longstaff-Schwartz Least Squares Monte Carlo (LSM) approach, we approximate the value of the operation by simulating M paths of the state variables over the time horizon T. We employ a time discretization $\Delta t = \frac{T}{N-1}$ and simulate the two-factor model dynamics using the Euler-Maruyama scheme.

The valuation procedure is analogous to determining the optimal exercise policy for an American option. At each discrete time point t_n, the fish farmer must choose between harvesting immediately or waiting. This decision is governed by comparing the immediate Har-

vesting Value (VH) with the expected Continuation Value (VC). The optimal value function $V(t_n, x_{t_n})$ at time t_n and state x_{t_n} satisfies the Bellman equation:

$$V(t_n, x_{t_n}) = \max\left\{VH(t_n, x_{t_n}), \quad VC(t_n, x_{t_n})\right\} \tag{2.10}$$

The Harvesting Value VH is defined as the revenue from selling the biomass X_{t_n} at the current market price P_{t_n}, net of harvesting costs CH_{t_n}:

$$VH(t_n, x_{t_n}) = P_{t_n} X_{t_n} - CH_{t_n} \tag{2.11}$$

Here, the total harvesting cost CH_{t_n} is calculated as the product of the unit harvesting cost c_h and the total biomass X_{t_n} (i.e., $CH_{t_n} = c_h X_{t_n}$), covering variable expenses such as slaughtering and processing.

The Continuation Value VC represents the expected benefit of deferring the harvest. It accounts for the immediate feeding costs CF_{t_n} incurred over the interval Δt and the discounted expected value of the farm at the next time step:

$$VC(t_n, x_{t_n}) = -CF_{t_n}\Delta t + e^{-r\Delta t}\mathbb{E}_{\mathbb{Q}}[V(t_{n+1}, x_{t_{n+1}})|\mathcal{F}_{t_n}] \tag{2.12}$$

where $\mathcal{F}_{t_n}$ denotes the filtration representing the information set available to the fish farmer at time t_n.

The optimal stopping time τ is determined via backward induction starting from $t_N = T$. Harvesting occurs at the first time instance where the immediate value exceeds the continuation value:

$$\tau = \inf\{t_n : VH(t_n, x_{t_n}) > VC(t_n, x_{t_n})\} \tag{2.13}$$

Since the conditional expectation in the continuation value is unknown, LSM approximates it by projecting future cash flows onto a set of basis functions. The LSM algorithm addresses this by approximating the continuation value using a linear regression of realized discounted future payoffs against a set of basis functions of the current state variables. The procedure is implemented via backward induction. First, we simulate a large number of paths for the state variables using the discretized Schwartz two-factor model. At the terminal time T, the value is simply the harvest value, assuming a forced harvest or zero value if not harvested. The algorithm then proceeds recursively backwards from T. At each time step, we identify paths where immediate harvesting is profitable, known as "in-the-money" paths. For these paths, we regress the discounted future cash flows onto basis functions of the current state variables to estimate the expected continuation value. Following Ewald et al. (2017), we utilize a set of Laguerre polynomials as basis functions to ensure numerical stability, estimating the coefficients via Ordinary Least Squares (OLS). Specifically, we assume $VC(t_n, x_{t_n})$ can be represented as a linear combination of basis functions $L_j(x)$. In this study, we employ a

quadratic basis:

$$\widehat{VC}_{t_n} = \hat{a}_1 x_1^2 + \hat{a}_2 x_2^2 + \hat{a}_3 x_1 + \hat{a}_4 x_2 + \hat{a}_5 x_1 x_2 + \hat{a}_6 \tag{2.14}$$

The coefficients $\hat{a}_j$ are estimated by regressing the discounted future realized cash flows C_{t_k} (from time t_{n+1} onwards) onto the basis functions of the current state. The realized cash flow along a path is given by:

$$C_{t_k} = \begin{cases} -CF_{t_k}\Delta t + VH_{t_k} & \text{if } \tau = t_k \\ -CF_{t_k}\Delta t & \text{otherwise} \end{cases} \tag{2.15}$$

Finally, the estimated continuation value is compared to the immediate harvest value. If the immediate harvest value exceeds the estimated continuation value, the optimal decision is to harvest, and the stopping time for that path is updated. By iterating this process back to the initial time, we obtain the optimal exercise policy and the estimated present value of the fish farm. This value serves as the baseline ground truth for evaluating the accuracy of our DQN agent.

2.4 Deep Q-Learning Framework for Optimal Harvesting

To address the limitations of traditional regression-based methods in high-dimensional settings, we adopt a model-free Deep Reinforcement Learning (DRL) approach. This section details the transformation of the fish farmer's objective into an optimal stopping problem, its theoretical underpinnings via the Snell envelope, and the specific Deep Q-Network (DQN) architecture used for the solution. Furthermore, we introduce a dual martingale approach to validate the optimality of the learned policy.

2.4.1 Optimal Stopping and the Snell Envelope

Based on the valuation components defined previously, we formally pose the fish farmer's objective. Following the notation in Becker et al. (2020), let $(\Omega, \mathcal{F}, \mathbb{P}, \mathbb{F})$ be a filtered probability space supporting the stochastic processes P_t and δ_t.

Let $g(t, X_t)$ denote the intrinsic payoff function. In our context, this corresponds to the net profit from harvesting: $g(t, X_t) = P_t X_t - CH_t$. The objective is to find an admissible $\mathbb{F}$-stopping time $\tau \in \{t_0, t_1, \ldots, t_N\}$ that maximizes the expected discounted payoff starting from time zero:

$$V_0 = \sup_{\tau \in \mathcal{T}} \mathbb{E}\left[e^{-r\tau} g(\tau, X_\tau) - \sum_{k=0}^{\tau/\Delta t - 1} e^{-rt_k} CF_{t_k} \Delta t \right] \tag{2.16}$$

where the second term represents the accumulated discounted feeding costs up to the time of harvest.

Theoretically, the solution to this optimal stopping problem is characterized by the Snell envelope Z_n, which is the smallest supermartingale dominating the payoff process. Defined recursively, the value process V_n satisfies:

$$V_N = g(N, X_N), \quad V_n = \max(g(n, X_n), e^{-r\Delta t}\mathbb{E}[V_{n+1}|\mathcal{F}_n]). \tag{2.17}$$

The optimal stopping time τ^* is the first time the intrinsic payoff equals the continuation value implied by the Snell envelope:

$$\tau^* = \inf\{t_n : g(t_n, X_{t_n}) = V_n\} \tag{2.18}$$

In the context of Q-learning, we decompose this envelope into state-action values. The Q-function $Q(s, a)$ effectively approximates the conditional expectations inherent in the Snell envelope decomposition:

$$Q(s_n, \text{harvest}) = g(n, s_n), \tag{2.19}$$

$$Q(s_n, \text{continue}) = e^{-r\Delta t}\mathbb{E}[V_{n+1}(S_{n+1})|S_n = s_n] \approx e^{-r\Delta t}\mathbb{E}[\max_{a'} Q(S_{n+1}, a')|S_n = s_n]. \tag{2.20}$$

Thus, the Deep Q-Network acts as a function approximator for the continuation value component of the Snell envelope, allowing us to circumvent the computation of high-dimensional conditional expectations required by traditional backward induction.

2.4.2 Markov Decision Process (MDP) Formulation

To implement the sequential decision-making process described above, we explicitly reformulate it as a Markov Decision Process (MDP), defined by the tuple $(\mathcal{S}, \mathcal{A}, R, \mathcal{P}, \gamma)$.

State Space ($\mathcal{S}$): The state space is constructed to satisfy the Markov property. We define the state vector at time t_n as $s_n = [P_{t_n}, \delta_{t_n}, w_{t_n}, n_{t_n}, t_n]^\top$. It is important to note that the state space $\mathcal{S} \subseteq \mathbb{R}^5$ is **continuous** and high-dimensional, encompassing financial variables (spot price, convenience yield) and biological variables (average weight, population count).

Action Space ($\mathcal{A}$): The action space is binary, representing the farmer's operational choices. At each decision epoch t_n, the agent selects an action $a_n \in \{0, 1\}$, where:

- $a_n = 0$ (Continue): The farmer keeps the fish for another period to benefit from potential growth.
- $a_n = 1$ (Harvest): The farmer harvests the entire biomass and terminates the episode.

Reward Function (R): The reward function $R(s_n, a_n)$ serves as the feedback signal to the agent. It is designed to match the cash flows from the LSM framework:

$$R(s_n, a_n) = \begin{cases} -CF_{t_n} \Delta t & \text{if } a_n = 0 \text{ (Continue)} \\ P_{t_n} X_{t_n} - CH_{t_n} & \text{if } a_n = 1 \text{ (Harvest)} \end{cases} \tag{2.21}$$

Transition Dynamics ($\mathcal{P}$): The transitions are governed by the stochastic differential equations for P_t and δ_t (Schwartz model) and the deterministic biological growth functions for w_t and n_t, as defined in Section 3.1.

Discount Factor (γ): The discount factor corresponds to the risk-free rate, defined as $\gamma = e^{-r\Delta t}$.

2.4.3 Deep Q-Network (DQN) Algorithm

Network Architecture and Loss Function

To address the inherent limitations of tabular Q-learning in continuous state spaces, we employ the Deep Q-Network (DQN) algorithm (Mnih et al., 2015). We approximate the optimal action-value function $Q^*(s, a)$ using a deep neural network parameterized by weights w, denoted as $Q(s, a; w)$.

This approach leverages the *Universal Approximation Theorem* (Hornik et al., 1989), which states that a feed-forward neural network with sufficient hidden units can approximate any continuous function to arbitrary accuracy. This is a crucial advantage over LSM, which relies on projecting the continuation value onto a finite set of basis functions (e.g., polynomials) and may suffer from bias if the basis does not span the true value function space.

The network is trained to minimize the Mean Squared Bellman Error (MSBE). The loss function is given by:

$$L(w) = \mathbb{E}_{(s,a,r,s') \sim U(D)} \left[\left(\underbrace{R(s, a) + \gamma \max_{a'} \hat{Q}(s', a', w^-)}_{\text{TD Target}} - Q(s, a, w) \right)^2 \right] \tag{2.22}$$

where w^- denotes the parameters of the target network.

To effectively capture the nonlinear relationships, we construct a Multi-Layer Perceptron (MLP). The input layer consists of neurons corresponding to the normalized state vector. The network features two fully connected hidden layers (64 and 32 neurons respectively)

with Rectified Linear Unit (ReLU) activation functions. The final output layer consists of two linear neurons representing the Q-values for $\mathcal{A} = \{\text{Continue}, \text{Harvest}\}$.

Training Stability Mechanisms

Because general convergence guarantees for Q-learning do not extend to arbitrary nonlinear function approximation (Bertsekas and Tsitsiklis, 1996), we incorporate two key mechanisms to stabilize learning:

1. **Experience Replay:** We store transition tuples (s, a, r, s') in a replay buffer D. Randomly sampling mini-batches breaks the temporal correlation between consecutive samples, ensuring the data is approximately independent and identically distributed (i.i.d.).

2. **Target Network:** A separate network $\hat{Q}$ with weights w^- is used to generate the TD targets. w^- are updated periodically to match w, preventing the "chasing a moving target" instability.

Algorithm 1 details the complete training procedure.

Algorithm 1 Deep Q-Network for Optimal Salmon Harvesting

Initialize replay memory D to capacity N_{mem}
Initialize action-value network Q with random weights w
Initialize target network $\hat{Q}$ with weights $w^- = w$
for episode $= 1, M$ **do**
    Initialize state s_0 (at $t_0 = 0$)
    for step $n = 0, N - 1$ **do**
        With probability ϵ select random a_n, else $a_n = \arg\max_a Q(s_n, a; w)$
        Execute a_n, observe r_n and s_{n+1}
        if $a_n = 1$ (Harvest) **then**
            $r_n \leftarrow VH(t_n, x_{t_n})$, Terminate episode
        else
            $r_n \leftarrow -CF \cdot \Delta t$
        end if
        Store transition (s_n, a_n, r_n, s_{n+1}) in D
        Sample random mini-batch from D
        Set $y_j = r_j + \gamma \max_{a'} \hat{Q}(s_{j+1}, a'; w^-)$ (if non-terminal)
        Perform gradient descent on $(y_j - Q(s_j, a_j; w))^2$
        Every C steps, reset $\hat{Q} = Q$
        Update state $s_n \leftarrow s_{n+1}$
    end for
end for

2.4.4 Validation via Dual Upper Bound

The DQN policy yields a lower bound on the true value V_0. To assess the quality of this approximation, we outline a corresponding dual upper-bound procedure. Following the duality theory developed by Rogers (2002) and Andersen and Broadie (2004), and adapted for neural networks by Ery and Michel (2021), we construct a dual maximization problem based on the martingale representation theorem.

Let V_0 be the true optimal value. It is established that:

$$V_0 = \sup_{\tau \in \mathcal{T}} \mathbb{E}[g_\tau] = \inf_{M \in \mathcal{M}_0} \mathbb{E}\left[\max_{0 \le n \le N} (g(n, X_n) - M_n)\right], \tag{2.23}$$

where $\mathcal{M}_0$ denotes the space of martingales M with $M_0 = 0$.

To compute this empirically, we utilize the learned Q-values to approximate the optimal martingale. The martingale difference sequence ΔM_{n+1} is constructed using the Doob decomposition of the estimated value process:

$$\Delta M_{n+1} = \widehat{V}_{n+1}(X_{n+1}) - \mathbb{E}[\widehat{V}_{n+1}(X_{n+1}) | \mathcal{F}_n], \tag{2.24}$$

where $\widehat{V}_n(x) = \max_a Q(x, a; w^*)$ is the value function implied by the converged neural network. The conditional expectation can be estimated via nested Monte Carlo simulation. In an implementation of this diagnostic, the gap between the policy value (a lower bound) and the dual estimate (an upper bound) provides a useful measure of approximation quality. Because the numerical dual bound is not reported in the present results, we treat this construction as a validation framework rather than as direct empirical evidence of exact optimality.

2.5 Results and Discussion

2.5.1 Comparison of Models

We compare the performance of the DQN model with the benchmark Longstaff–Schwartz (LSMC) method. Both models were evaluated using 50,000 simulated paths with antithetic variates to reduce variance. Table 2.2 shows that the two methods generate broadly similar average harvesting times, while the DQN option-value estimates remain below the LSMC benchmark.

Although the optimal exercise boundary is not explicitly plotted, an analysis of the agent's decision-making patterns reveals a coherent "wait-and-see" structure typical of American-style options. The DQN agent autonomously established a dynamic threshold for the spot

Table 2.2: Average Harvesting Time and Real Option Value Comparison

	LSMC (Monte Carlo)		DQN (Deep Q-Learning)	
Contracts	**Harvest Time (yrs)**	**Option Value**	**Harvest Time (yrs)**	**Option Value**
Panel A	2.07	1,512,400	2.28	1,249,772
Panel B	2.23	1,646,700	2.36	1,152,817

price $P(t)$ relative to the biological mass $X(t)$. During early growth stages where biomass $X(t)$ is low, the agent consistently assigns higher Q-values to the "Continue" action, effectively creating a high reservation price for early harvesting. This reflects the option value of biological growth. Conversely, as the fish approach their asymptotic weight w_∞ and biological growth rates decline, the value of waiting diminishes, causing the agent to lower the price threshold required to trigger the "Harvest" action. Furthermore, the neural network successfully captures the mean-reversion dynamics of the convenience yield $\delta(t)$. In states of high $\delta(t)$—indicating market scarcity—the agent exhibits an increased propensity to harvest early to capture the convenience yield premium, a behavior consistent with the theoretical predictions of Schwartz (1997). The ability of the neural network to reconstruct these economically rational boundaries without explicit programming further validates the model-free learning capability of the proposed framework.

Table 2.2 also shows that the DQN option values are lower than the LSMC estimates. The gap is approximately 17.4% in Panel A and 30.0% in Panel B. These differences indicate that the two approximation schemes should not be treated as numerically equivalent in the present calibration. They may reflect differences in policy approximation, exploration, finite-sample training, and the treatment of high-value states. Accordingly, the DQN results are best interpreted as evidence that the agent learns a broadly similar harvesting policy, rather than as proof that it reproduces the LSMC valuation exactly.

2.5.2 Scalability in High-Dimensional Environments

While Table 2.2 shows broadly similar harvesting times under a baseline model driven by two stochastic financial factors (spot price and convenience yield), the complete DQN state vector also includes biological state variables and time. The primary potential advantage of the RL framework lies in its scalability when additional stochastic drivers are introduced. The traditional LSMC method relies on the regression of basis functions (e.g., Laguerre or Hermite polynomials). As the number of state variables increases, the number of required basis functions grows exponentially, leading to the well-known "curse of dimensionality". Consider a more complex, realistic aquaculture scenario incorporating environmental stochasticity, such as water temperature variability. Water temperature is a critical high-dimensional variable that nonlinearly impacts salmon growth rates, mortality (via extreme temperature events), and feed conversion ratios. In an LSMC framework, adding temperature, feed price volatil-

ity, and disease outbreak probabilities would require a prohibitively complex expansion of the basis function set to capture the cross-correlations between these factors. In contrast, the Deep Neural Network (DNN) architecture utilized in our DQN agent serves as a universal function approximator. It can inherently extract complex, nonlinear features from high-dimensional inputs without manual feature engineering. For example, water temperature, feed prices, or disease indicators can be appended to the existing state vector without requiring the researcher to enumerate a large set of polynomial interactions. By contrast, a conventional polynomial LSMC specification may require a rapidly expanding set of basis terms and cross-products as the number of stochastic drivers increases. Therefore, while our current simulation validates the DQN on a baseline model, the framework is specifically designed to outperform regression methods as the biological and economic realism of the model increases.

2.5.3 Parameter Sensitivity Analysis

To assess the robustness of the valuation model, we further examine the impact of stochastic mortality dynamics on the real option value. The mortality rate $m(t)$ is modeled by a Cox-Ingersoll-Ross (CIR) process, ensuring non-negativity:

$$dm(t) = \kappa_2(\theta - m(t))dt + \sigma_m\sqrt{m(t)}dW_t \tag{2.25}$$

where κ_2 represents the speed of mean reversion, θ is the long-run mean mortality rate, and σ_m denotes the volatility of the mortality process.

Table 2.3 illustrates the sensitivity of the option value and the average harvesting timing to variations in these parameters.

Table 2.3: Sensitivity Analysis: Mortality Rate Parameters

κ_2	θ	σ_m	Option Value (NOK)	Harvest Time (yrs)
2.0	0.15	0.5	774,815	2.53
2.5	0.15	0.5	789,240	2.52
4.0	0.15	0.5	846,049	2.48
4.0	0.15	0.8	730,010	2.56
1.5	0.15	0.3	1,073,774	2.35

The results indicate a distinct positive correlation between the speed of mean reversion (κ_2) and the real option value. Economically, a higher κ_2 implies that spikes in mortality rates (e.g., due to temporary disease outbreaks) dissipate more rapidly, pulling the rate back to the long-run mean θ. This stability preserves the biomass stock, thereby increasing the project's expected value. Conversely, higher volatility in mortality rates (σ_m) significantly reduces the option value. Unlike financial asset price volatility, which typically enhances option value, volatility in the underlying physical inventory (i.e., the risk of fish death) represents a pure downside volume risk. Higher σ_m increases the probability of severe biomass

depletion events, effectively eroding the asset base and delaying or precluding the optimal harvest, thus penalized by the valuation model.

2.6 Conclusion

This chapter presented a novel decision-support framework for optimal salmon harvesting by integrating the stochastic dynamics of the Schwartz (1997) two-factor pricing model with advanced Deep Reinforcement Learning (DRL) techniques. By formulating the harvesting decision as an optimal stopping problem within a Markov Decision Process (MDP), we successfully trained a Deep Q-Network (DQN) agent to approximate the optimal exercise policy for the real option of harvesting.

Our numerical results lead to three primary conclusions. First, the DQN agent learns economically coherent harvesting behavior and produces average harvesting times that are broadly comparable with the benchmark Least-Squares Monte Carlo (LSMC) method. The option-value estimates are lower than the LSMC benchmark, however, so the present evidence supports policy similarity more strongly than numerical equivalence. The exercise demonstrates that a neural-network approximator can be used without manually specifying a polynomial basis for the continuation value.

Second, while both methods perform comparably in lower dimensions, this study highlights the superior scalability of the DRL framework. Traditional regression-based methods like LSMC often suffer from the "curse of dimensionality" when additional stochastic factors are introduced. In contrast, the DQN framework utilizes deep neural networks as universal function approximators, making it uniquely positioned to handle high-dimensional complexities. This characteristic is particularly critical for future aquaculture models that must incorporate additional stochastic variables, such as regime-switching mortality rates, variable feeding costs, and environmental factors like water temperature volatility.

Finally, from a managerial perspective, this model provides aquaculturists with a robust tool for risk mitigation. By explicitly incorporating futures market data and stochastic biological factors, the model illustrates how managerial flexibility, specifically the timing of the harvest, adds significant value to the fish farm. The approach confirms that integrating financial derivatives into operational decision-making allows farmers to better navigate market uncertainties and maximize profitability. This study thus serves as a foundational step toward applying sophisticated machine learning algorithms to complex, high-dimensional resource management problems in aquaculture economics.

Chapter 3

Biomass Futures Contracts for the Aquaculture Sector: An Exploration

3.1 Introduction

Climate change has made fish farming more uncertain. Warmer oceans and extreme weather harm fish growth and survival. Salmon farming is especially sensitive to temperature changes, which can affect fish metabolism and lead to higher death rates. Fish farmers need better tools to manage these risks. According to the IPCC, sea surface temperatures have risen by 0.13°C per decade since 1901. Marine heatwaves are also more frequent, disrupting ecosystems. These trends call for solutions that blend biological insight with financial planning.

Managing risks in aquaculture is complex. Factors like water temperature, feed quality, fish density, and diseases directly influence fish growth. Water temperature is critical because it directly affects fish metabolism and growth efficiency. Brett (1979) shows that even small changes in water temperature can significantly affect how fish grow. Water temperature plays a crucial role in forecasting harvest outcomes. The effects of climate change make it increasingly difficult to anticipate water temperatures, posing challenges for fish farmers seeking to manage their operations effectively. Consequently, salmon farmers frequently need to adapt to environmental shifts, which can influence their feeding and harvesting strategies. These adjustments may negatively impact fish growth and ultimately diminish profitability.

Salmon aquaculture has evolved into a global industry, with farm-raised salmon now supplying a significant share of global fish consumption. The success of salmon farms depends a lot on water temperature. When temperatures are too high or too low, salmon grow slowly and may die. Farmers can face significant financial losses due to challenges in managing fish production. They might change how much they feed the fish, adjust the number of fish they raise, or shift their harvesting schedule to cope. However, these changes can be costly. Additionally, unpredictable environmental conditions make it harder for these strategies to work and complicate long-term planning for fish farmers. Researchers have extensively stud-

ied how to manage renewable natural resources under uncertain conditions, especially in the fisheries sector. Pizarro and Schwartz (2021) develop and implement a stochastic optimal control approach to value marine fisheries, demonstrating how uncertainty in environmental and economic conditions affects resource valuation and decision-making. Their framework highlights the importance of integrating stochastic models to optimise resource management under uncertainty. Similarly, salmon farming faces comparable challenges, where environmental variability introduces substantial uncertainty in production and profitability.

Thus, developing a financial instrument such as previously proposed biomass futures contracts could provide a more robust mechanism for managing the economic risks posed by environmental fluctuations. By allowing farmers to hedge against unfavourable biomass levels, such instruments would help stabilize income and introduce a level of predictability currently lacking in the face of climate variability.

The introduction of biomass futures explicitly addresses a "missing market" in aquaculture risk management. While existing salmon futures (such as those traded on Fish Pool ASA) effectively hedge against market price volatility, they do not offer protection against production or quantity risks. If a farmer loses a significant portion of their biomass due to a sudden temperature spike or disease outbreak, locking in a high selling price provides little financial relief for the lost volume. However, establishing a functional biomass futures market faces several practical obstacles. These include the difficulty of implementing standardized and verifiable biomass measurement across different farms, the basis risk arising from localized environmental conditions that diverge from the regional index, and the broader challenges of generating sufficient market liquidity and designing robust contract settlement mechanisms acceptable to both producers and financial speculators.

Marine biologists have created models predicting fish growth under various environmental conditions. These models allow farmers and policymakers to see how temperature and feeding rates affect fish growth. Research in aquaculture growth modelling has focused on species such as salmon (Bjorndal, 1988), trout (Logan and Johnston, 1992), shrimp (Leung et al., 1990), and tilapia (Fischer and Grant, 1994). These studies offer essential insights into how environmental factors affect aquaculture production. However, they often overlook the need for financial risk management strategies. As ecological conditions become increasingly unpredictable, it is essential to integrate biological growth factors with financial instruments. This strategy will enhance the management of the inherent risks associated with aquaculture production.

Integrating price (P) and quantity (Q) futures has been extensively researched as a robust risk management tactic in agricultural and financial markets. Aase (2004) demonstrated that combining area yield futures and futures options enables farmers to hedge income risks effectively by dynamically replicating revenue by trading futures. This approach leverages the relationship between price and quantity uncertainties to reduce income volatility. Similar

strategies are emerging in aquaculture, where environmental and market risks are closely intertwined. Previous studies have primarily focused on separate applications of price or quantity instruments, yet limited research explores their joint use in addressing dual uncertainties. Connecting biological elements with financial strategies can aid in stabilizing income and enhancing decision-making in unpredictable scenarios. This method presents an innovative solution by integrating biomass futures with price-oriented hedging techniques. It links environmental shifts with market approaches, offering valuable insights and practical uses.

In 1992, the Chicago Board of Trade (CBOT) launched Crop Yield Insurance (CYI) futures and options contracts. This was an essential step in agricultural finance. These financial products helped farmers, agribusinesses, and others protect themselves against the risks of changing crop yields. Historically, agricultural risk management has relied heavily on government-backed insurance and subsidy programs, which assist farmers in recovering losses from adverse weather conditions, pest infestations, and other uncontrollable factors. Government support helps by providing a safety net, but it has some problems. These include delays in payments, influence from politics, and market issues. Because of this, the agricultural sector needs more financial certainty and flexibility to manage risks. Although CYI futures contracts initially provided a private market solution, they were eventually phased out and replaced by other risk management instruments, such as weekly options for new crops currently traded on the Chicago Mercantile Exchange (CME); see (CME, 2025). To keep things relatively simple, though, in this work, we focus on contracts aligned with the original products issued by the CBOT.

CYI futures introduced a market-based alternative, allowing farmers to hedge directly against income losses caused by lower-than-expected crop yields. Market participants can utilize futures contracts to lock in anticipated yield levels and safeguard against potential losses. These contracts are grounded in crop yield estimates provided by the U.S. Department of Agriculture (USDA). The USDA regularly updates these estimates, serving as a standard reference for the settlement of contracts. For example, the Iowa Corn Yield Insurance futures contract was settled based on the USDA's monthly estimate of corn yield per acre in Iowa. By enabling farmers to hedge against yield shortfalls through futures contracts, CYI helped shift part of the risk management burden from government programs to the open market. This approach helped farmers become more flexible and improve their financial management. It also helped others in the farm supply chain, like storage facilities and transportation companies, manage their risks more effectively.

The introduction of CYI futures changed how people viewed and handled agricultural risks. Previously, farmers had few options to manage risks from changing yields. Traditional insurance can cover some losses, but it often pays out slowly and has limited coverage. This can lead to cash flow problems and poor decision-making. With CYI futures, farmers and others can manage risks more actively. Storage companies, transportation providers, and input suppliers can protect themselves against supply chain issues. These new tools for man-

aging risks provide more flexibility, better planning, and more efficient use of resources in agriculture. Trading of CYI futures also sparked innovation in financial markets, creating new financial products tailored to the unique risks of farming. The application of CYI futures contracts within a mean-variance-based framework has been explored in (Vukina et al., 1996), for example.

Although CYI futures were eventually phased out, their initial implementation offered a valuable model for using financial instruments to manage production risks. This approach has since influenced other sectors exposed to environmental uncertainty, such as aquaculture and particularly salmon farming. The temperature of the water is a critical factor in determining the growth rates of salmon, and fluctuations in temperature, which are caused by climate change, can have severe impacts on fish biomass and overall production. In the aquaculture industry, temperature changes that deviate from the optimal range can slow growth and increase mortality rates, leading to significant financial losses for fish farmers. Therefore, much like CYI futures offered a solution to mitigate risks associated with unpredictable crop yields, there is the potential to introduce a similar financial tool as biomass-based futures contracts for the aquaculture industry. These contracts would enable fish farmers to hedge against the risk of reduced biomass due to adverse environmental conditions, thereby protecting their revenues and promoting more sustainable farming practices. From a theoretical point of view, these contracts have been explored in (Aase, 2004). In our current work, we focus more on the practical issues and assess how beneficial such contracts could be to aquaculture businesses with regard to revenue creation and risk management.

Aquaculture-based biomass futures would function similarly to CYI futures and related products in the agricultural sector, allowing salmon farmers to hedge against the risk of yield shortfalls caused by climate-induced changes in water temperature or disease outbreaks, for example. In such a system, futures contracts would be tied to biomass estimates derived from environmental and production data regularly collected by regulatory bodies or independent monitoring agencies. As water temperature changes threaten the stability of the aquaculture industry, financial tools could help salmon farmers secure their income and make better production choices. Fish farmers can manage risks better by making key choices, such as when to harvest and how to feed their fish. They can use information about fish growth and survival along with market strategies. This helps them connect their farming decisions with financial planning, protecting them from environmental risks.

Biomass futures can help connect ecological risk management with financial planning. This tool offers a way to manage risks while planning financially. León-Santana and Hernández (2006) made significant progress in analysing how water temperature affects the optimal management of aquaculture farms. Their work helped to understand how environmental changes can affect the economy. However, we still need better models that include financial tools to protect against these risks. As climate change creates more uncertainty in production, predicting and managing environmental risks is becoming more critical. Aquaculture

can use biomass futures to stabilise revenues and make better long-term decisions. This can help salmon farmers operate more efficiently and promote sustainable practices by linking financial rewards to good environmental care.

The application of financial engineering techniques to the problem of environmental risk management in aquaculture draws on well-established methods in other industries. The optimal stopping problem, which is used extensively in financial markets to determine the best time to exercise American-style options, has been applied across various sectors, from gambling and medical decision-making to urban planning. Bertsekas and Tsitsiklis (1996) developed Approximate Dynamic Programming (ADP) to compute optimal decisions under uncertainty. In contrast, Longstaff and Schwartz (2001) introduced the Least-Squares Monte Carlo (LSMC) method, a benchmark model in option pricing. These models offer potent frameworks for managing financial risks, and their application to aquaculture could provide fish farmers with new ways to mitigate the economic impacts of climate variability. Farmers can develop more effective risk management strategies that account for biological and economic factors by linking environmental risks by engaging with these financial instruments.

In addition to financial modelling, the pricing of real options has become an essential area of research in industries characterised by high levels of uncertainty. In salmon farming, deciding when to harvest fish is akin to an option: the farmer must choose the optimal time to "exercise" the harvest, balancing the risks of environmental changes against the potential economic gains of delaying the harvest for better growth outcomes. Ewald et al. (2017) developed a two-factor model incorporating spot prices and stochastic convenience yields to analyze the optimal stopping problem in salmon farming. By adapting this model to include biomass futures, we can further refine the decision-making process, offering a more holistic approach to risk management that accounts for market and environmental factors. Harvesting fish is essentially a “real option,” meaning farmers choose when to harvest based on balancing growth potential against environmental risks. Ewald et al. (2017) developed a model linking fish futures prices, risk perception and optimal harvesting times. Our paper extends their approach by introducing biomass futures, combining biological growth models with financial market data to better manage biological and climate-related risks. We show how these financial instruments can improve economic outcomes using real-world data from Fish Pool ASA.

3.2 Model

3.2.1 Fish Growth

Since the value of biomass futures contracts is directly linked to the biomass of farmed fish, it is crucial to understand the underlying biological dynamics that drive biomass growth.

Therefore, we introduce a fish growth model to quantify these factors. Fish growth is a dynamic process that farmers seek to optimize to maximize profitability. Since variations in market supply partially drive salmon price fluctuations, understanding the biological factors influencing fish growth and survival is essential for predicting long-term price behavior. The model in this paper assumes that three factors influence growth: the fish's weight $w(t)$, the size of food ration $r(t)$ given to the fish, and the water temperature $T(t)$, where t denotes time. Following León-Santana and Hernández (2006), the effects of these three variables on fish growth are modelled multiplicatively. To simplify notation and enhance interpretability, we define a growth function $G(w, r, T)$ as:

$$G(w, r, T) = \Gamma f_1(w) f_2(r) f_3(T), \tag{3.1}$$

where Γ accounts for additional growth-influencing factors and is estimated via calibration. The function $f_1(w)$ captures the relationship between fish weight and growth rate. Growth follows an increasing but concave trajectory, i.e. $f_1' > 0$ and $f_1'' < 0$. The function $f_2(r)$ represents the effect of feeding levels on growth. The ration size $r(t)$ is expressed in normalized terms, where $r = 0$ represents no feeding, and $r = 1$ corresponds to the maximum feeding level that leads to satiation. It also follows an increasing, concave pattern with a minimum value at the maintenance ration $r = r_m$, the lowest ration required to sustain weight, and a maximum at $r = 1$. The function $f_3(T)$ captures the impact of temperature on growth. The stochastic differential equation (SDE) for fish growth can then be written as:

$$dw(t) = G(w(t), r(t), T(t))dt + \sigma_w w(t) dW_w(t). \tag{3.2}$$

This general model follows the principles of fish physiology (Brett, 1979). Initially developed for salmon growth models, it has been extended to other species (Muller-Feuga, 1990). The term $\sigma_w w(t) dW_w(t)$ represents the stochastic fluctuations in fish weight due to environmental and biological uncertainties. The coefficient σ_w quantifies the magnitude of these fluctuations. The presence of the multiplicative term $w(t)$ aligns with empirical observations that growth fluctuations are proportional to fish size (Gamito, 1998). This formulation is beneficial for capturing the effects of unpredictable environmental conditions, such as sudden changes in water quality, feed efficiency variations, and minor genetic differences among fish. The weight function $f_1(w)$ is assumed to be log-linear (Ricker, 1979), implying that while weight gain is positively correlated with fish size, the percentage growth rate decreases as the fish gets larger:

$$f_1(w) = C w^m. \tag{3.3}$$

Empirical studies suggest that this function better describes fish growth than alternative

models, particularly for species such as seabream (Hernández et al., 2003). The parameter m must be less than one ($m < 1$), which aligns with estimates for warm-blooded vertebrates, where $m \approx 0.7$ (Brett, 1979). The ration size functions $f_2(r)$ describe fish appetite and its influence on growth. The ration level varies from starvation ($r = 0$) to the maximum ration ($r = r_M$). The normalized growth rate $Z(r)$, which represents the growth rate as a percentage of its maximum possible value, is defined as follows:

$$f_2(r_n) = \frac{Z(r_n)}{Z(r_c)}. \tag{3.4}$$

The function $Z(r)$ is determined by the normalized conversion rate $Y(r)$, which accounts for the efficiency with which food is converted into body mass. The conversion rate function is given by:

$$Y(r) = r\left(1 + \frac{(r_m - r_o)^2(1-r)^2}{(r_u - 1)(3r_m r_o - r_m - 2r_o^2)(r - r_m)}\right) \tag{3.5}$$

This formulation ensures that the conversion rate remains a function of the feeding level, adjusting for diminishing returns and inefficiencies at extreme feeding levels. r_m represents the maintenance ration, the minimum food required to sustain weight.r_o is the optimal ration, where food conversion efficiency is maximized. r_u is the upper limit ration, beyond which additional food intake does not contribute to growth. The normalized growth rate follows:

$$Z(r) = \frac{r_n}{Y(r)}, \quad Z' > 0, \quad Z'' < 0. \tag{3.6}$$

This formulation ensures that fish growth remains biologically realistic, reflecting both the benefits of increased feeding and the inefficiencies associated with excessive food intake. By incorporating these constraints, the model provides a flexible yet accurate representation of the relationship between feeding and fish growth. The function $f_3(T)$ describes the impact of water temperature on fish growth. Temperature affects metabolic activity, with excessively low and high temperatures reducing growth efficiency. Following the widely used Brett (1979) growth-temperature model, we define $f_3(T)$ as:

$$f_3(T) = \exp\left(-\frac{(T - T_{\text{opt}})^2}{2\sigma_T^2}\right), \tag{3.7}$$

where T_{opt} is the optimal growth temperature (typically 15-17°C for Atlantic salmon). σ_T controls the sensitivity of growth to deviations from T_{opt}. This function ensures that growth is maximized at T_{opt} and declines symmetrically as temperature deviates from this optimum. The Gaussian-like structure is a standard way to model thermal growth responses in aquacul-

ture. The dynamics of water temperature $T(t)$ are simulated using a mean-reverting process, specifically employing an Ornstein-Uhlenbeck process. The equation for this process is given by

$$dT(t) = \kappa_T(\theta - T(t))dt + \sigma_T dW_T(t) \tag{3.8}$$

where $T(t)$ represents the water temperature at time t. Here κ_T denotes the rate of mean reversion, dictating the speed at which the water temperature reverts to its long-term average. θ is the long-term average temperature. The parameter σ_T represents the standard deviation of temperature fluctuations, representing the magnitude of environmental randomness. $W_T(t)$ is a Wiener process, reflecting the random fluctuations in temperature changes, assumed to be uncorrelated with the Brownian motions $\tilde{Z}_1(t)$ and $\tilde{Z}_2(t)$. This model generates a time series by simulating the water temperature in continuous small steps, each based on the current temperature and its deviation from the long-term average. This approach is particularly suited for simulating temperature changes in natural environments, as it captures the characteristic fluctuations over time while accounting for the impact of environmental randomness. The fish stock $n(t)$ evolves according to the differential equation:

$$dn(t) = -m(t)n(t)dt + \sigma_n n(t)dW_n(t). \tag{3.9}$$

The mortality rate $m(t)$ is modeled using a Cox-Ingersoll-Ross (CIR) process:

$$dm(t) = \kappa_m(\theta_m - m(t))dt + \sigma_m\sqrt{m(t)}dW_m(t). \tag{3.10}$$

where $m(t)$ is the mortality rate, κ_m is the speed of mean reversion, θ_m is the long-run mean mortality rate, σ_m is the volatility of the mortality process, and $dW_m(t)$ denotes an additional Brownian-motion increment assumed to be uncorrelated with the previously introduced Brownian motions. The total biomass $B(t)$ is then given by:

$$B(t) = n(t)w(t). \tag{3.11}$$

Since $w(t)$ and $n(t)$ are both subject to stochastic influences, biomass exhibits multiplicative stochasticity. The fish weight follows the growth dynamics described in Equation (2). This formulation ensures that biomass evolution captures the interacting effects of growth, mortality, feeding, and environmental temperature fluctuations. Applying Itô's Lemma, we compute the differential of $B(t)$, incorporating the stochastic dynamics of both weight and population size.

$$dB(t) = w(t)dn(t) + n(t)dw(t) + dn(t)dw(t). \tag{3.12}$$

Substituting the stochastic differential equations for $dw(t)$ and $dn(t)$:

$$dB(t) = [G(w(t), r(t), T(t))n(t) - m(t)B(t)]\,dt + \sigma_w n(t)w(t)dW_w(t) + \sigma_n n(t)w(t)dW_n(t). \tag{3.13}$$

Here, the drift term consists of two components: the first term, $G(w, r, T)n(t)$, represents the total growth contribution of individual fish, while the second term, $-m(t)B(t)$, captures the reduction in biomass due to mortality. The stochastic terms $\sigma_w n(t)w(t)dW_w(t)$ and $\sigma_n n(t)w(t)dW_n(t)$ reflect the random fluctuations in individual fish weight and the total number of fish, respectively. While it is in principle possible to work with the dynamic (13) directly, it would cause many computational issues, as an explicit formula for the biomass futures price would not be available. Instead, we continue with an approximation of the process described by (13) that maintains key features and can be calibrated to fit (13) as closely as possible, while still allowing for an explicit expression of the futures price.

A key assumption in this model is that biomass follows a mean-reverting process, which accounts for the natural tendency of biological systems to stabilize around an equilibrium level due to environmental and resource constraints. Thus, we seek to approximate the biomass dynamics by a mean-reverting stochastic differential equation of the form:

$$dB(t) = \kappa_B(\alpha_B - \ln B(t))B(t)dt + \sigma_B B(t)dW_B(t), \tag{3.14}$$

where κ_B represents the speed of mean reversion, α_B is the long-run equilibrium level in logarithmic form, and σ_B is the volatility of biomass fluctuations. In conventional mean-reverting processes, κ_B determines the rate at which biomass returns to its equilibrium level. $\kappa_B > 0$ is the rate of mean reversion for $B(t)$. The parameter α_B reflects the long-run equilibrium level of $B(t)$ in log form, allowing α_B to take on negative values. $\sigma_B > 0$ represents the standard deviation of the relative changes in $B(t)$. We have $B(t) > 0$ for all t, ensuring the non-negativity of the biomass process. By modeling $B(t)$ in this form, the log of $B(t)$, denoted $\ln B(t)$, follows an Ornstein-Uhlenbeck process, which has well-defined moments.

To ensure that the simplified mean-reverting biomass model (Equation 3.14) accurately captures the complex biological dynamics under different environmental scenarios, particularly the impact of high mortality on the long-term biomass level, we employ a moment matching calibration technique. Instead of relying on heuristic approximations, we calibrate the parameters κ_B, α_B, and σ_B to match the statistical properties of the full biological simulation defined by Equations (3.1) to (3.13).

We first generate a large set of Monte Carlo paths ($N = 10,000$) using the full biological model, which explicitly accounts for stochastic temperature $T(t)$, growth function $G(\cdot)$, and

the Cox-Ingersoll-Ross mortality process $m(t)$. Let $Y_T = \ln B(T)$ denote the logarithm of the total biomass at the maturity horizon T generated by this full simulation. We compute the first two empirical moments of the terminal log-biomass: the empirical mean $\hat{\mu} = \mathbb{E}[Y_T]$ and the empirical variance $\hat{v} = \mathrm{Var}[Y_T]$.

The theoretical moments for the simplified Ornstein-Uhlenbeck process for $\ln B(t)$ are known analytically. By equating the analytical moments to the empirically derived moments from the full simulation, we can solve for the effective parameters. Specifically, we fix the mean-reversion speed κ_B (which can be estimated from the autocorrelation of the simulated paths) and then determine the long-run equilibrium level α_B and the volatility σ_B as follows:

$$\alpha_B = \frac{\hat{\mu} - y_0 e^{-\kappa_B T}}{1 - e^{-\kappa_B T}}, \tag{3.15}$$

$$\sigma_B = \sqrt{\frac{2\kappa_B \hat{v}}{1 - e^{-2\kappa_B T}}}, \tag{3.16}$$

where $y_0 = \ln B(0)$ is the initial log-biomass.

Crucially, this calibration procedure ensures that the parameter α_B is directly linked to the mortality regime embedded in the full simulation. In scenarios with high mortality (e.g., Extreme Temperature Group), the empirical mean biomass $\hat{\mu}$ will be significantly lower. Consequently, the calibrated α_B derived from the equation above will decrease, reflecting a lower long-term equilibrium biomass. This adjustment corrects the economic anomaly where high-stress scenarios might otherwise appear unduly profitable; by accurately lowering the expected biomass level, the valuation model properly penalizes the project value for permanent biological losses, even in the presence of hedging instruments.

3.2.2 Biomass Futures Contracts

The biomass futures contract is structured around a reference level K. To make entry into the futures position costless at inception, we set $K = F(0, T)$, where $F(t, T) = \mathbb{E}_Q[B(T) \mid \mathcal{F}_t]$ denotes the expected biomass at maturity T under the pricing measure Q. The biomass index $B(t)$ is influenced by factors such as the average weight of the fish $w(t)$, the total number of fish $n(t)$, and various external determinants like water temperature, feeding conditions, and disease prevalence. The relationship between parameter values and moments is more complex for the geometric mean reversion process. Using (14), we consider the conditional expectation and variance of $\ln B(t)$ under the filtration $\mathcal{F}_t$, which represents all available information up to time t. Following Tvedt (2022), the conditional expectation and variance of $\ln B(\tau)$ under $\mathcal{F}_t$ are given by:

$$E\left[\ln B(\tau) \mid \mathcal{F}_t\right] = e^{-\kappa_B(\tau-t)} \ln B(t) + \left(\alpha_B - \frac{\sigma_B^2}{2\kappa_B}\right)\left(1 - e^{-\kappa_B(\tau-t)}\right), \tag{3.17}$$

$$\mathrm{Var}\left[\ln B(\tau) \mid \mathcal{F}_t\right] = \frac{\sigma_B^2}{2\kappa_B}\left(1 - e^{-2\kappa_B(\tau-t)}\right) \tag{3.18}$$

The conditional expectation of $\ln B(\tau)$ converges to $\alpha_B - \frac{\sigma_B^2}{2\kappa_B}$, and the variance converges to $\frac{\sigma_B^2}{2\kappa_B}$, as $\tau \to \infty$. This formulation ensures that $B(t)$ exhibits mean-reverting behaviour with a log-normal distribution. Based on the properties of a log-normally distributed variable, the biomass futures price is given by:

$$\begin{aligned} F(t,T) &= E\left[B(T) \mid \mathcal{F}_t\right] \\ &= \exp\left(e^{-\kappa_B(T-t)} \ln B(t) + \left(\alpha_B - \frac{\sigma_B^2}{2\kappa_B}\right)(1 - e^{-\kappa_B(T-t)}) \right. \\ &\quad \left. + \frac{\sigma_B^2}{4\kappa_B}(1 - e^{-2\kappa_B(T-t)})\right) \end{aligned} \tag{3.19}$$

Here we assume for simplicity that the dynamics of $B(t)$ under the pricing measure Q is the same as under the real-world measure. This can be generalized quite easily.

Equation (20) shows that the biomass futures price depends on the current log biomass level $\ln B(t)$, the speed of mean reversion κ_B, the long-run mean α_B, and the volatility σ_B. These parameters capture both biological uncertainty and market-based valuation. The long-run expected level of $B(t)$ is given by:

$$\lim_{\tau\to\infty} E\left[B(\tau) \mid \mathcal{F}_t\right] = \exp\left(\alpha_B - \frac{\sigma_B^2}{4\kappa_B}\right). \tag{3.20}$$

For a given α_B, the logarithm of the long-run average decreases with higher σ_B and with lower κ_B. Therefore, changes in σ_B or κ_B impact not only the variance of the biomass process but also the expected value. This dependency implies that variations in these parameters can have a complex influence on option values and optimal switching strategies. Salmon farmers will use the biomass futures as insurance against poor biomass growth. This works as follows. Salmon farmers go short biomass futures at a reference level K at time 0 when the farming cycle starts and then close out by engaging in the equivalent long position at the harvesting time τ. At a given nominal of M they then realize a profit of

$$\text{Payoff} = (K - F(\tau,T)) \times M. \tag{3.21}$$

The financial transaction adds to the profit from harvesting and selling the salmon. To understand how this acts as insurance, note that if the biomass is very low at harvest, then $F(\tau, T)$ will be low, and the payoff in (1) will be high. If the biomass is very high at the time of harvest, then the payoff will be very low (and negative). In this way, the total revenue's dependence on biomass at harvest is dampened, and the biomass risk is effectively hedged. Entering into futures contracts should be costless. Hence the reference level K should be set at the biomass futures price $F(0, T)$.

3.2.3 Salmon Futures

Assume $P(t)$ is the salmon spot price at time t. Following Schwartz (1997), the dynamics of the spot price and the convenience yield under the risk-neutral pricing measure are given by:

$$dP(t) = (r - \delta(t))P(t)\,dt + \sigma_1 P(t)\,d\tilde{Z}_1(t), \tag{3.22}$$

$$d\delta(t) = (\kappa(\alpha - \delta(t)) - \lambda)\,dt + \sigma_2\,d\tilde{Z}_2(t), \tag{3.23}$$

with correlated Brownian motions:

$$d\tilde{Z}_1(t)\,d\tilde{Z}_2(t) = \rho\,dt. \tag{3.24}$$

The spot price $P(t)$ reflects both market demand and underlying biological supply conditions in the salmon industry. In particular, supply is heavily influenced by biomass dynamics, including growth rates, mortality, and environmental conditions such as temperature. As fish reach harvest weight, their availability affects the market spot price through its impact on total biomass supplied. On the other hand, the convenience yield $\delta(t)$ captures short-term benefits from holding inventory, which in aquaculture may reflect production cycles, seasonal variation, or perishability risk. This framework is justified empirically in Ewald et al. (2017), who apply the two-factor model from Schwartz (1997) to salmon futures prices and estimate its parameters using the Kalman filter. Their findings confirm that the model fits well, with all key coefficients significant at the 1 percent level. Notably, the mean-reversion rate κ, long-run average convenience yield α, and the risk premium λ are all positive and economically reasonable. The spot price volatility σ_1 is relatively stable, while the volatility of the convenience yield σ_2 is more variable, reflecting dynamic inventory and biological conditions. The futures price $F(P, \delta, T)$ within that framework can be derived analytically:

$$F(P, \delta, T) = P \cdot \exp\left(-\delta \cdot \frac{1 - e^{-\kappa T}}{\kappa} + A(T)\right), \tag{3.25}$$

where the adjustment term $A(T)$ is given by:

$$A(T) = \left(r - \alpha + \frac{\lambda}{\kappa} + \frac{1}{2}\frac{\sigma_2^2}{\kappa^2} - \frac{\sigma_1\sigma_2\rho}{\kappa} \right) T + \frac{1}{4}\sigma_2^2 \left(\frac{1 - e^{-2\kappa T}}{\kappa^3} \right) + \left(\alpha\kappa - \lambda + \sigma_1\sigma_2\rho - \frac{\sigma_2^2}{\kappa} \right) \left(\frac{1 - e^{-\kappa T}}{\kappa^2} \right). \tag{3.26}$$

The use of this pricing model is supported by the empirical analysis in Ewald et al. (2017) and allows us to represent the dynamics of the salmon futures price in a tractable form. The spot price $P(t)$, driven by market expectations and the convenience yield, reflects both supply-side constraints and demand conditions in the salmon industry. While Ewald et al. (2017)'s model captures the price component, it does not explicitly account for the physical quantity of fish available in the market. In this study, we complement the price process with a stochastic model of biomass dynamics, representing the quantity $Q(t)$ available for harvesting. This allows us to study profit as the product $P(t) \cdot Q(t)$, linking financial market expectations with biological production realities. By incorporating environmental variability and growth uncertainty, the following section introduces a fish profit model that captures both price and supply fluctuations in a unified framework.

3.2.4 Harvesting

We consider the salmon harvesting problem to be of the same type as in Ewald et al. (2017). This framework links the fish growth model with the two-factor pricing framework of Schwartz (1997) for salmon futures. The contribution of the current paper is to add biomass futures to this framework. The model developed above provides a foundation for understanding the biomass dynamics under different environmental conditions, enabling improved predictions and risk management for the harvest profit function:

$$\mathrm{H}(t) = P(t)B(t) - CH(t),$$

where $P(t)$ is the biomass unit price, $B(t)$ the biomass volume and $CH(t)$ represents harvesting costs. Biological factors influence the harvest timing, including fish growth, health, water temperature, and feeding conditions. Farmers aim to harvest when biomass is near its peak, maximizing growth efficiency and minimizing costs. The maturity T of the biomass futures contract should extend beyond the expected harvesting time; the baseline calibration uses a three-year horizon. However, we assume that the farmer closes out the biomass futures contract at the time of harvest τ, which provides the payoff defined above. The feeding cost $CF(t)$ represents the variable cost farmers incur throughout the salmon farming cycle. Un-

like a fixed cost per biomass unit, feeding costs depend on the growth trajectory of the fish. To account for this, a conversion ratio measures the relationship between feeding quantity and fish weight gain. By incorporating this biologically realistic feeding cost structure, the model accounts for the dynamic nature of fish growth and resource consumption, providing a more accurate representation of production costs in aquaculture. The optimization problem for the farmer facilitating biomass futures can now be expressed as:

$$\max_{\tau} \mathbb{E}_{\mathbb{Q}} \left(e^{-r\tau} \left[P(\tau)B(\tau) + (K - F(\tau, T)) \times M - CH(\tau) \right] - \int_0^{\tau} e^{-rt} CF(t)dt \right), \quad (3.27)$$

where $P(\tau)$ is the spot price, $B(\tau)$ is the biomass, M is the monetary factor, $CH(\tau)$ represents the harvesting cost, and $CF(t)$ represents the feeding cost. The model Ewald et al. (2017) without biomass futures can be obtained by setting the monetary factor $M = 0$. This enables us to directly compare the two models and assess the impact of biomass futures contracts. In the next section, we will assess the impact of this arrangement on farm lease value.

3.2.5 Estimation of the Continuation Value Using the Longstaff-Schwartz Approach

The Longstaff-Schwartz approach provides a computationally efficient method for approximating the continuation value in optimal stopping problems. In our setting, this relates to the harvesting decision in salmon farming as described above. The method combines Monte Carlo simulation with least-squares regression to estimate the expected continuation value at each decision point. It captures the farmer's dynamic decision-making process, in which the decision to harvest or postpone is made repeatedly over time based on the trade-off between current and future expected payoffs, appropriately accounting for risk. Given that harvesting depends on both biological and financial conditions, the Longstaff-Schwartz method is particularly suitable for identifying optimal harvesting times. **Path Simulation.** We simulate M_{sim} Monte Carlo paths over a time horizon T, with a discretized time step $\Delta t = \frac{T}{N-1}$. The underlying state variables include the spot price $P(t)$, the convenience yield $\delta(t)$, and the biomass level $B(t)$. These follow the stochastic processes defined in Equations (23), (24), and (14), which model the dynamics of spot price, convenience yield, and biomass, respectively. The Euler–Maruyama scheme is used to numerically approximate each path. **Valuation Procedure.** At each decision time t_n, the farmer compares the immediate harvesting value (VH) with the expected continuation value (VC). The optimal value function satisfies the following Bellman equation:

$$V\left(t_n, x_{t_n}\right) = \max\Big\{ P_{t_n} B_{t_n} + (K - F(t_n, T)) \cdot M - CH_{t_n} \\ - CF_{t_n}\Delta t,\ e^{-r\Delta t}\mathbb{E}_{\mathbb{Q}}[V(t_{n+1}, x_{t_{n+1}}) \mid \mathcal{F}_{t_n}]\Big\}. \tag{3.28}$$

Here, M represents the monetary multiplier in the biomass futures contract. The immediate harvesting value is given by:

$$VH(t_n) = P_{t_n} B_{t_n} + (K - F(t_n, T)) \cdot M - CH_{t_n}. \tag{3.29}$$

The continuation value reflects the expected value of postponing the harvesting decision. It is computed as

$$VC(t_n) = -CF_{t_n}\Delta t + e^{-r\Delta t}\mathbb{E}_{\mathbb{Q}}[V(t_{n+1}) \mid \mathcal{F}_{t_n}]. \tag{3.30}$$

Under the assumption of no arbitrage and using the risk-neutral measure $\mathbb{Q}$, the continuation value can also be expressed as the expected present value of future cash flows until the stopping time τ:

$$VC(t_n) = \mathbb{E}_{\mathbb{Q}}\left[\sum_{k=n+1}^{\tau} e^{-r(t_k - t_n)} C_{t_k} \mid \mathcal{F}_{t_n}\right]. \tag{3.31}$$

The per-period cash flow C_{t_k} is defined as:

$$C_{t_k} = \begin{cases} VH_{t_k} - CF_{t_k}\Delta t, & \text{if } \tau = t_k \\ -CF_{t_k}\Delta t, & \text{otherwise.} \end{cases} \tag{3.32}$$

Regression-Based Approximation. The continuation value function $VC(t_n)$ does not have a known analytical form. To approximate it, we follow the approach of Longstaff and Schwartz (2001), which involves least-squares regression on simulated cash flows. We use a quadratic basis expansion involving the three state variables: spot price, convenience yield, and biomass. The approximation takes the following form:

$$\widehat{VC}_{t_n} = \hat{a}_0 + \sum_{i=1}^{3} \hat{a}_i x_i + \sum_{i=1}^{3} \hat{a}_{i+3} x_i^2 + \sum_{1 \le i < j \le 3} \hat{a}_{ij} x_i x_j \tag{3.33}$$

The state vector is $x = [P(t_n), \delta(t_n), B(t_n)]^\top$, and the coefficients $\hat{a}_j$ are estimated by regressing the realized discounted future cash flows onto these basis functions. To improve

the efficiency of estimation, we perform the regression only on those paths where immediate harvesting yields a positive net payoff. This focuses computational resources on paths where the option is in-the-money and harvesting is a relevant decision.

3.2.6 Parameters

To implement the model, we use parameter values sourced from the literature, fish farm data, or calibrated based on the simulation requirements to ensure economic consistency. For clarity, we categorise these parameters into two groups. Table 3.1 summarises the financial market parameters related to the salmon futures price dynamics following the Schwartz (1997) two-factor model. Table 3.2 provides the biological and operational parameters governing fish growth, mortality, and the futures contract specification.

The financial parameters in Table 3.1 are aligned with the estimation by Ewald et al. (2017) (Panel D, Mixed-term Panel), which provides a balanced view of price dynamics across short, medium, and long-term maturities. We set the risk-free rate r to 3.03%. Regarding operational costs, we fix the variable harvesting cost at 3 NOK/kg and the feeding cost at 7 NOK/kg per year, with a farming cycle horizon of $T = 3$ years.

Table 3.1: Schwartz (1997) Two-Factor Model Parameters (Aligned with Simulation)

Parameter	Description	Value
r	Risk-free interest rate	0.0303
P_0	Initial Spot Price (NOK/kg)	40.4
δ_0	Initial Convenience Yield	0
κ	Mean-reversion speed of convenience yield	1.664
α	Long-term mean level of convenience yield	0.460
σ_1	Volatility of spot price	0.214
σ_2	Volatility of convenience yield	0.448
ρ	Correlation between Brownian motions	0.806
λ	Market price of convenience yield risk	0.690

Table 3.2 summarizes the key biological and contract parameters. Unlike standard models that treat biomass parameters (α_B, σ_B) as fixed inputs, our physics-informed framework treats them as calibrated outputs derived from the underlying biological dynamics. Therefore, Table 3.2 lists the fundamental physical parameters that drive the simulation.

Crucially, the mortality parameter θ_m is not fixed but is regime-dependent, varying across the three temperature scenarios to reflect different levels of environmental stress. Similarly, the contract multiplier M is set to 10 in the baseline calibration. The multiplier is chosen conservatively to represent partial rather than full hedging and to prevent the derivative payoff from dominating the value of the physical production position.

The parameters in Table 3.2 are structured to separate deterministic biological traits from stochastic environmental factors. The growth function relies on the allometric exponent $m =$

Table 3.2: Biological Model and Futures Contract Parameters

Parameter	Description	Value	Source
Growth and Feeding			
Γ	Growth scaling factor	1	León-Santana and Hernández (2006)
C	Growth rate coefficient	1	Calibration
m	Allometric weight exponent	0.58	Brett (1979)
r_o	Optimal feeding ratio	0.50	Brett (1979)
r_m	Maintenance ration	0.12	Brett (1969)
r_c	Recommended feeding ratio	0.80	León-Santana and Hernández (2006)
c_r	Feed conversion rate	1.1	Industry Standard
Environmental and Mortality			
T_{opt}	Optimal growth temperature (°C)	16.5	Brett (1979)
$\sigma_{T,growth}$	Growth sensitivity to temp. dev.	2.5	Calibration
κ_T	Temperature mean reversion rate	0.5	Calibration
σ_T	Temperature volatility (°C)	3.5	Calibration
κ_m	Mortality mean reversion rate	0.3	Calibration
θ_m	Long-term mean mortality rate	*Scenario*	0.15 - 0.55
σ_m	Mortality rate volatility	0.05	Calibration
Stochastic Noise and Contract			
σ_w	Indiv. weight fluctuation	0.02	Estimation
M	Futures monetary multiplier	10	Calibrated

0.58, ensuring concavity in growth as fish mature. The temperature sensitivity parameter $\sigma_{T,growth} = 2.5$ controls how strictly fish growth is penalized when water temperatures deviate from the optimum $T_{opt} = 16.5$°C. A lower value here implies higher sensitivity to thermal stress.

A key distinction in our parameterization is the treatment of mortality θ_m. Instead of a single value, we define a range: 0.15 for Ideal conditions, 0.30 for Thermal Stress, and 0.55 for Extreme conditions. This variation is the primary driver of the physical asset loss in our simulations. Finally, the futures contract multiplier is set to $M = 10$, consistent with Table 3.2. This conservative calibration allows the futures payoff to offset part of the loss from adverse biomass outcomes without dominating the value of the underlying physical position.

3.3 Results

According to fish management practices in Scotland, the optimal temperature range for Atlantic salmon spans 6-20°C, with maximal growth occurring at 16-17°C. Thermal stress begins at approximately 17.5°C, and extreme stress occurs at around 18.5°C. To analyze the impact of biomass futures, we evaluate three groups under different water temperature conditions: Group 1 represents ideal growth conditions with $\kappa_T = 0.5$ per year, $\theta = 16.5$°C, $\sigma_T = 3.5$°C; Group 2 reflects thermal stress conditions with $\kappa_T = 0.5$ per year, $\theta = 17.5$°C, $\sigma_T = 3.5$°C; Group 3 simulates extreme temperature conditions with $\kappa_T = 0.2$ per year, $\theta = 18.5$°C, $\sigma_T = 3.5$°C.

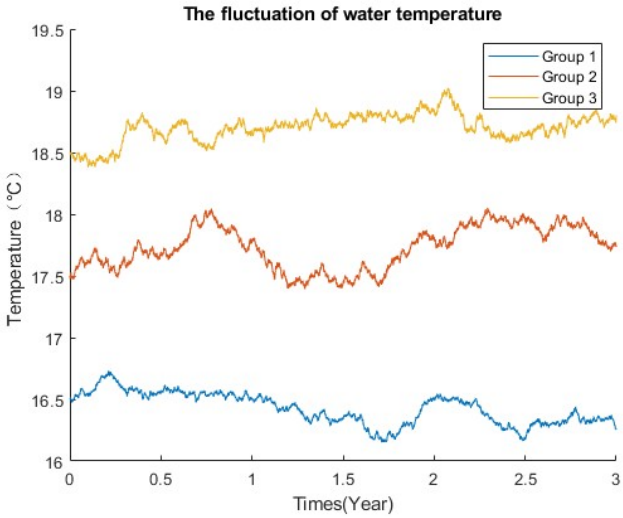


Figure 3.1: Simulated water temperature conditions

We apply the Longstaff-Schwartz method to solve the optimal stopping problem in fish farming, utilizing the two-factor model parameters estimated in Table 3.1. These parameters reflect futures contracts traded on the Fish Pool market, while additional biological and economic parameters relevant to fish farming are listed in Table 3.2.

To enhance computational efficiency while capturing the path-dependent nature of the option, we employ antithetic variates, simulating 5,000 sample paths over a three-year horizon with 1,000 time steps. The time discretization follows a fine-grained approach ($\Delta t = T/(N-1)$), enabling the farmer to continuously reassess harvesting strategies in response to stochastic environmental shifts.

Table 3.3: Real Option Value and Hedging Effectiveness Across Temperature Scenarios (Calibrated)

Group	**Unhedged Value (NOK)**	**Value with Hedged Biomass (NOK)**	**Increment Due to Hedge (NOK)**	**Increment Percentage**
1 (Ideal)	723,558	658,039	-65,519	-9.06%
2 (Stress)	693,260	653,729	-39,530	-5.70%
3 (Extreme)	615,171	641,331	+26,160	+4.25%

Table 3.3 summarizes the valuation results under the physics-informed moment matching framework, incorporating a conservatively calibrated contract multiplier to avoid overhedging. The results demonstrate a rational trade-off between risk mitigation and profit maximization.

In the Ideal Growth scenario (Group 1), the biological asset performs optimally, with the unhedged real option value peaking at NOK 723,558. However, the hedged portfolio value is lower at NOK 658,039, resulting in a negative increment of 9.06%. This represents the implicit "insurance premium" paid by the farmer. Under favorable conditions, biomass accumulates reliably, causing the futures price $F(\tau, T)$ to rise above the strike price K. Consequently, the short futures position incurs a loss, which partially offsets the high revenue from the physical sale of the fish. The optimal harvesting time is approximately 2.76 years, allowing the fish to reach their maximum growth potential.

In the Extreme Temperature scenario (Group 3), the unhedged value drops significantly to NOK 615,171. This decline reflects the physical reality of thermal stress: high mortality rates ($\theta_m \approx 0.55$) significantly reduce the harvestable biomass. Crucially, however, the hedging instrument provides a positive payoff, raising the total project value to NOK 641,331 (an increment of 4.25%). As the expected biomass trajectory collapses due to high mortality, the futures price $F(\tau, T)$ falls below the strike price K, generating a profit from the short position.

Most importantly, the simulation results now satisfy the condition of economic consistency. The total hedged value in the Extreme scenario (NOK 641,331) remains lower than the value in the Ideal scenario (NOK 723,558 or NOK 658,039). This pattern is consistent with the intended role of the biomass futures contract as a partial hedge: it mitigates downside risk without making the adverse high-mortality state more valuable than the favorable state. The contract softens the blow of environmental disaster, recovering approximately 4.25% of value relative to the unhedged baseline, but it does not fully replace the lost revenue from physical production.

The optimal stopping times also reflect a rational adaptation to risk. In the Extreme scenario, the average harvesting time shortens slightly to 2.67 years (compared to 2.76 years in the Ideal scenario). This behavior indicates a strategy of "cutting losses": given the high

mortality rate and the profitable short futures position, the optimal policy is to harvest slightly earlier to lock in the remaining physical value and the financial gains from the hedge, rather than risking further biological depletion.

To further elucidate these decision dynamics, Figures 3.2 to 3.4 visualize the optimal harvesting boundaries. The decision rule is nonlinear: under extreme stress, the harvesting threshold lowers, indicating a "rush to exit" to salvage value, whereas in ideal conditions, the boundary suggests patience to exploit the full growth potential.

The incremental benefit of hedging grows with environmental stress, highlighting the role of biomass futures as an effective risk management tool when biological uncertainty intensifies. This pattern can be explained by the relationship between the expected biomass path and the futures price $F(\tau, T)$. Under ideal conditions, biomass grows steadily, and the futures price increases over time. Since farmers short the futures at $K = F(0, T)$, they incur a loss when $F(\tau, T) > K$. However, in stress or extreme scenarios, the expected biomass path declines or becomes volatile, which leads to lower $F(\tau, T)$, resulting in a positive hedge payoff $K - F(\tau, T)$. Thus, the hedge becomes more valuable precisely when biological conditions are unfavorable, acting as a form of income insurance against poor production outcomes. Which of the three scenarios will materialize is uncertain. Nevertheless, although hedging may reduce value in favorable scenarios, it offers substantial protection under adverse conditions, confirming its value as a climate risk management instrument for the aquaculture industry. To better understand the harvesting decision dynamics, we visualize the optimal harvesting points based on salmon spot price P, convenience yield δ, and biomass B in Figures 3.2 to 3.4. The distribution of optimal harvesting points in Figures 2–4 reveals necessary threshold behaviour. Harvesting is more likely when biomass levels are high, even if prices or convenience yields are moderate. However, harvesting is triggered earlier and at lower biomass levels under extreme temperature conditions, especially when convenience yields are low. This suggests that farmers act conservatively under stress, preferring to lock in gains before further biomass deterioration occurs. The separation between harvesting and waiting decisions appears nonlinear, reinforcing the need for data-driven decision rules.

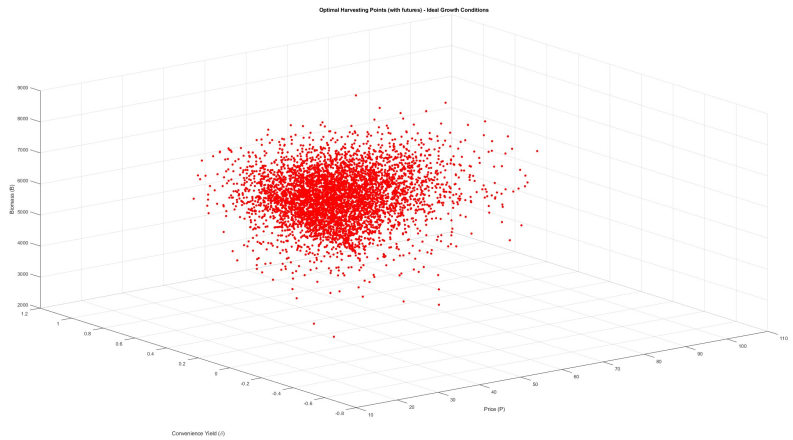


Figure 3.2: Optimal Harvesting Points under Ideal Growth Conditions

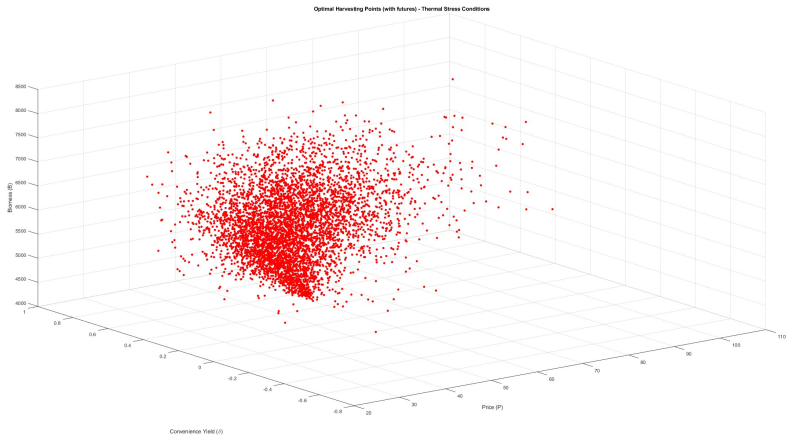


Figure 3.3: Optimal Harvesting Points under Thermal Stress Conditions

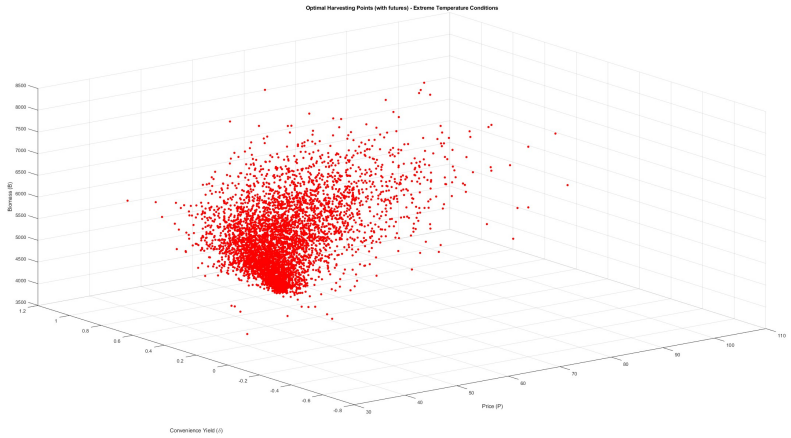


Figure 3.4: Optimal Harvesting Points under Extreme Temperature Conditions

Due to the complexity of the harvesting decision, which depends on three continuous state variables (P, δ, B), the optimal strategy forms a multidimensional surface within the state space. To approximate this boundary, we apply a supervised machine learning approach to model the relationship between the salmon spot price P, the convenience yield δ, and the biomass level B at the time of harvest.

Specifically, we employ a nonparametric bagged decision-tree ensemble to estimate the functional form of the harvesting surface. By aggregating predictions from multiple decision trees trained on bootstrap samples, this model effectively captures complex nonlinear interactions between price and convenience yield. Training the ensemble on simulated optimal harvesting points allows us to recover the underlying biomass threshold function, $\hat{B}(P, \delta)$, which implicitly defines the harvesting rule. This data-driven approach enables the visualization of how the critical biomass threshold varies across diverse market conditions.

Figures 3.5–3.7 present the fitted harvesting surfaces under different temperature scenarios. In these figures, the semi-transparent black dots represent the simulated harvesting points, while the gray surfaces depict the estimated harvesting thresholds. These surfaces delineate the state space into two regions: the exercise region, where immediate harvesting

is optimal, and the continuation region, where waiting is preferred. Unlike parametric models that assume a closed-form solution, the bagged decision tree ensemble flexibly separates these regions without imposing a rigid functional form. Consequently, this approach provides a rigorous method for visualizing and interpreting the optimal stopping boundary as a hypersurface, even in the absence of an analytical solution. Future extensions could explore alternative classifiers or attempt to approximate the boundary using interpretable functions or logistic regression for additional transparency.

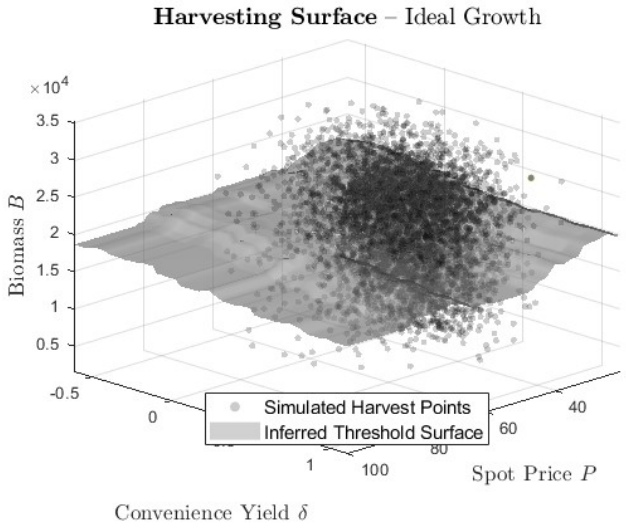


Figure 3.5: Harvesting Surface under Ideal Growth Conditions

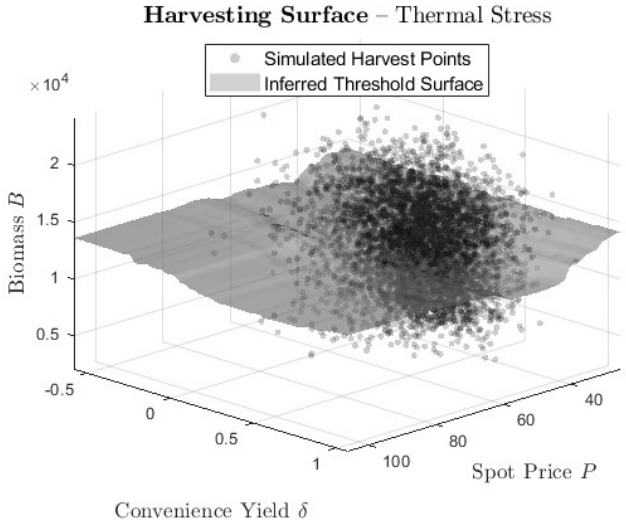


Figure 3.6: Harvesting Surface under Thermal Stress Conditions

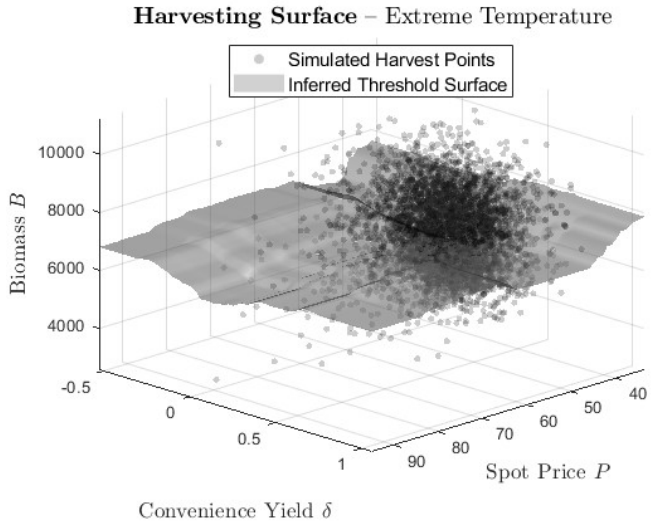


Figure 3.7: Harvesting Surface under Extreme Temperature Conditions

This regression-based machine learning framework offers a flexible and interpretable method for characterizing optimal harvesting strategies. It accurately captures the nonlinear interactions between price, convenience yield, and biomass, and provides practical insights into how these factors jointly determine the timing of optimal harvesting decisions.

Table 3.4: Mean Optimal Stopping Time and Classification Accuracy (Calibrated)

Group	**Mean Stopping Time (No Futures)**	**Mean Stopping Time (With Futures)**	**Classification Accuracy**
1 (Ideal)	2.85 years	2.76 years	91.08%
2 (Stress)	2.80 years	2.74 years	91.90%
3 (Extreme)	2.75 years	2.67 years	92.80%

Table 3.4 reports the average optimal stopping times and the classification accuracy of the bagged decision-tree models across the three temperature scenarios. The results illustrate how the introduction of a conservatively calibrated hedging instrument ($M = 10$) influences operational decisions.

A striking feature of the new results is the temporal stability of the optimal harvesting strategy. Unlike speculative hedging strategies that might encourage premature harvesting ("cashing in" the insurance), the partial hedging approach incentivizes farmers to maintain production cycles closer to the biological optimum.

In Ideal Growth Conditions, the mean stopping time with futures is 2.76 years, only slightly shorter than the typical unhedged maturity cycle (approximated at 2.85 years). This minor contraction reflects a rational lock-in effect: as biomass reaches marketable size and futures prices rise (resulting in a liability in the short position), the farmer harvests slightly earlier to cap the derivative loss while securing high physical revenues.

Under Extreme Temperature Conditions, the harvesting time adjusts to 2.67 years. Crucially, this is not a panic-driven "bailout" (which would be characterized by a collapse in stopping time to < 1 year). Instead, it represents a resilient strategy. Despite high mortality

rates ($\theta_m = 0.55$), the presence of the futures contract mitigates the financial risk of holding the stock. The farmer can afford to wait until the fish reach a larger weight (approx. 2.67 years) because the futures payout ($K - F(\tau, T)$) compensates for the fish that die during this waiting period. The decision to harvest roughly one month earlier than in the Ideal scenario (2.67 vs. 2.76) strikes an optimal balance between salvaging remaining physical biomass and realizing the hedging gain.

Finally, the classification accuracy of the decision boundary improves as environmental conditions worsen, rising from 91.08% in the Ideal scenario to 92.80% in the Extreme scenario. This suggests that the optimal decision rule becomes more unambiguous under stress. In favorable conditions, the trade-off between waiting for growth and harvesting for profit is nuanced (borderline cases are frequent). In contrast, under extreme mortality pressure, the signal to "harvest and hedge" becomes mathematically sharper, making the decision boundary easier for the logistic regression model to classify.

3.4 Conclusion

This paper extends the real option framework proposed by Ewald et al. (2017) by incorporating biological and environmental risk factors and introducing biomass futures contracts that help salmon farmers to effectively manage these risks. Our model captures fish growth, mortality, and temperature dynamics through stochastic dynamics and simulation and applies the Longstaff–Schwartz algorithm to solve the optimal harvesting problem in the presence of three stochastic factors: salmon spot price, convenience yield, and biomass volume. By linking biological processes with biomass futures contracts, the framework evaluates how these financial instruments influence harvesting decisions in the presence of risk. The results indicate that biomass futures do not always enhance profitability. Under ideal growth conditions, using futures can lead to lower expected profits. As environmental conditions deteriorate, the cost of hedging becomes smaller: the hedge remains value-reducing in the thermal-stress calibration but turns value-enhancing in the extreme-temperature scenario. This pattern indicates that the relative value of insurance rises with biological risk. This suggests that the value of hedging is highest when biological risks are elevated. The proposed framework offers a practical approach to managing both environmental and financial risk in aquaculture. By reducing income volatility, it supports more stable and informed production planning. Although the model is developed for Atlantic salmon, the approach can be adapted to other species and farming systems. Future research may explore its application to a wider range of aquaculture contexts and assess its robustness under different parameter settings and contract designs to promote long-term resilience and sustainability.

Chapter 4

Valuation of Nuclear Power Investment with Real Options Theory

4.1 Introduction

The global energy sector is currently undergoing a profound transformation driven by the urgent need to mitigate climate change. According to the Intergovernmental Panel on Climate Change (IPCC), the continuous accumulation of greenhouse gases (GHG) in the atmosphere is pushing the planetary climate system toward irreversible changes. To limit global warming in accordance with the Paris Agreement, nations must rapidly decarbonize their power grids. In this context, nuclear power has regained strategic importance. Unlike solar and wind energy, which are variable and weather-dependent, nuclear power provides carbon-free baseload electricity that is essential for grid stability.

However, nuclear projects are notoriously difficult to value using traditional financial methods. They are characterized by immense capital costs, long construction lead times, and significant exposure to regulatory and market uncertainties (OECD Nuclear Energy Agency, 2011). Standard Net Present Value (NPV) analysis often undervalues these projects because it fails to capture the managerial flexibility, such as the option to delay investment, that is crucial in such an uncertain environment (Dixit and Pindyck, 1994).

To address these valuation challenges, a growing body of literature has applied Real Options Analysis to the energy sector. This methodology treats investment opportunities as financial options, allowing researchers to quantify the value of waiting for new information. Rothwell (2006) was among the first to apply this framework to new nuclear builds. He highlighted how uncertainties regarding future net revenues and policy implementation create a substantial "risk premium," often justifying the delay of nuclear investments until regulatory conditions stabilize. Building on this, Zhu (2012) developed a more comprehensive simulation model that simultaneously considered multiple risk factors, including volatile construction costs, electricity prices, and the rare but impactful risk of nuclear accidents. His work

demonstrated that the uncertainty in construction costs is a dominant factor in determining the optimal investment timing. Comparative studies have further illustrated the unique position of nuclear power relative to other clean technologies. Thomas and Chrysanthou (2012) utilized real options to compare nuclear power with wind energy. They found that due to different cost structures and learning rates, the investment signals for these technologies diverge significantly under market uncertainty.

While previous studies have laid a solid foundation, they share a common methodological limitation regarding how they model climate policy risk. In most existing real options models, the price of carbon is treated as an exogenous stochastic variable. Typically, it is modeled as a Geometric Brownian Motion (GBM). This implies that carbon prices follow a random walk, mathematically similar to stock prices. However, this assumption disconnects the financial market from the physical reality of climate change. In the real world, climate policy is not random. Instead, it is reactive. As the environment deteriorates, regulations tend to tighten drastically.

To address this gap, this study develops a "physics-informed" valuation framework. This framework builds upon a rich lineage of environmental economics literature. Pindyck (2000) pioneered the analysis of policy timing under environmental irreversibility. He demonstrated that the inability to reverse ecological damage creates a significant value to reduce emissions early. Expanding on this, Saphores (2004) and Saphores and Carr (2019) explored optimal pollution control under ecological uncertainty. However, these seminal works typically assumed a constant rate of natural pollutant decay. This implies that nature's capacity to absorb pollution remains the same, regardless of the accumulated stock. This assumption fails to capture the risk of systemic environmental collapse.

In this paper, we advance this literature by introducing a state-dependent decay rate. This is the core innovation of our framework. Unlike previous models where the Earth's absorption capacity is constant, we posit that the natural decay rate degrades as the pollution stock accumulates. Specifically, as the atmospheric concentration approaches a critical physical threshold—a climate tipping point—the environment's ability to recover diminishes rapidly. Economically, this physical mechanism creates a "regime-switching" dynamic. When the atmospheric state is far from the tipping point, the system is stable. However, as pollution levels near the critical threshold, the risk of a catastrophic loss in natural absorption capacity rises. Rational actors anticipate this structural break. This leads to a surge in the shadow price of carbon before the tipping point is even reached. We term this phenomenon the "fear premium." By using this physics-based logic, our model captures a jump in carbon pricing that is triggered by physical accumulation, rather than random market noise.

To rigorously formalize the concept of the "fear premium" within an asset pricing framework, our model draws upon the intersection of the rare disaster risk literature and Knightian uncertainty. Traditional rational expectations models struggle to generate substantial risk

premiums without assuming implausibly high levels of risk aversion. However, as demonstrated by Barro (2006) and Rietz (1988), the potential for rare macroeconomic disasters can fundamentally alter asset pricing by shifting the stochastic discount factor. In the context of climate change, Weitzman (2009) illustrates that the economic consequences of catastrophic climate events generate profound tail risks. Crucially, because humanity lacks historical data to estimate the precise probability distribution of a climate tipping point, agents face severe ambiguity, or Knightian uncertainty. Following the robust control framework of Hansen and Sargent (2001), rational but ambiguity-averse investors will evaluate intertemporal plans under the worst-case probability measure. It is precisely this aversion to ambiguity regarding the tipping point's threshold that amplifies the social cost of carbon (Lemoine and Traeger, 2014). In our model, the "fear premium" mathematically emerges from this ambiguity aversion: markets price the possibility of an environmental regime shift based on robust, pessimistic beliefs, driving up the option value of zero-emission assets long before the threshold is physically breached.

However, integrating such detailed climate physics into a financial valuation model introduces significant computational and structural challenges. The resulting value function exhibits a sharp discontinuity or "kink" at the tipping point, creating a non-smooth surface that is difficult to approximate using traditional numerical methods. Standard Least Squares Monte Carlo (LSM) relies on polynomial basis functions, which tend to smooth over these critical discontinuities, potentially underestimating the risk premium.

To overcome this barrier and empirically validate the most effective valuation approach for physics-informed models, we conduct a comparative analysis of distinct algorithmic architectures. While recent literature has explored Neural Networks for high-dimensional option pricing (Han et al., 2018), their performance in the presence of hard physical thresholds remains under-explored. We hypothesize that tree-based ensemble methods (e.g., Random Forest and XGBoost), which naturally partition state space via binary splits, offer superior performance in capturing the regime-switching dynamics compared to continuous global approximators.

To precisely contextualize our contribution, it is imperative to delineate our approach from foundational models, particularly the comprehensive framework developed by Zhu (2012). While we adopt Zhu's robust specifications for modeling construction budget dynamics, operating costs, and compound Poisson accident risks, we introduce four critical methodological departures that fundamentally alter the valuation logic. First, we replace the standard exogenous random-walk carbon price with an endogenous, physics-informed carbon price derived from the accumulated pollutant stock and macroeconomic conditions. Second, we introduce a tipping-point driven regime shift that reflects the abrupt loss of environmental absorption capacity. Third, this physical mechanism generates a non-smooth exercise boundary in the value function, breaking the assumption of smooth continuations found in traditional models. Finally, rather than merely applying machine learning as a computational substitute, we

explicitly utilize the structural differences between tree-based algorithms and global smooth approximators to resolve the specific mathematical discontinuities induced by the climate physics.

Therefore, this paper contributes to the literature not merely as an algorithmic extension, but by fundamentally bridging physical climate science and financial engineering. First, it applies a physics-informed climate model to the valuation of nuclear energy, explicitly linking financial value to physical climate risks and identifying the "fear premium" associated with tipping points. Second, it provides a rigorous methodological validation, demonstrating that tree-based machine learning algorithms are uniquely suited, and significantly outperform traditional polynomial methods, in valuing long-term assets subject to these physics-induced discontinuous risks.

4.2 Model Description and Parameter Settings

This model is developed with reference to existing simulation-based real options approaches in the nuclear power literature. Assuming the total duration for evaluating the nuclear power project (including both construction and operation) is T years, we divide this period into N intervals, each of length $\Delta t = T/N$, and define $t_n = n\Delta t$ for $n = 0, 1, \ldots, N$.

4.2.1 Modelling Uncertainty Factors in Nuclear Power Investment

Nuclear Power Investment Cost

We denote by $K(t)$ the expected remaining capital requirement, or investment budget, at time t needed to complete the nuclear project. Let $I(t)$ represent the instantaneous rate of investment expenditure controlled by the firm. The project commences with an initial estimated budget $K(0) = K_0 > 0$. Construction is deemed complete at the stochastic stopping time τ, defined as the first moment the remaining budget is exhausted:

$$\tau \equiv \inf\{t \geq 0 \,:\, K(t) = 0\}.$$

The total realized outlay is thus given by $\int_0^\tau I(s)\,ds$.

Following the framework of Dixit and Pindyck (1994) and Zhu (2012), we model the evolution of the remaining budget $K(t)$ subject to technical uncertainty. Unlike market risk, technical uncertainty—such as engineering challenges, regulatory delays, or design modifications—resolves only as the construction progresses. The dynamics of $K(t)$ are governed

by the following controlled diffusion process:

$$dK_t = -I_t\,dt + \beta\sqrt{I_t K_t}\,dW_t^{(K)}, \tag{4.1}$$

where $\beta > 0$ is the volatility parameter scaling the magnitude of construction cost uncertainty, and $W^{(K)}$ is a standard Wiener process. Since construction shocks are idiosyncratic to the project and not correlated with the market portfolio, we treat the dynamics of $K(t)$ identically under both the physical measure $\mathbb{P}$ and the risk-neutral measure $\mathbb{Q}$.

Equation (4.1) captures three critical empirical features of large-scale nuclear builds. First, the deterministic drift term $-I_t\,dt$ reflects the expected reduction in the remaining budget as capital is deployed. Second, the diffusion term $\beta\sqrt{I_t K_t}$ implies that the variance of the budget change is proportional to both the investment rate I_t and the remaining scope K_t. This specification effectively models the concept of learning-by-doing, where uncertainty regarding cost overruns is revealed only through the act of investment. If investment stops ($I_t = 0$), the budget remains constant, and no new information is generated; conversely, accelerating construction increases the instantaneous risk of discovering unexpected costs. Third, as the project nears completion and K_t approaches zero, the absolute magnitude of uncertainty diminishes, reflecting the progressive resolution of technical unknowns.

For a discrete time step Δt, the moments of the budget change are approximated as:

$$\mathbb{E}\left[\Delta K_t \mid \mathcal{F}_t\right] = -\,I_t\,\Delta t, \quad \mathrm{Var}\left(\Delta K_t \mid \mathcal{F}_t\right) = \beta^2 I_t K_t\,\Delta t. \tag{4.2}$$

Hence the instantaneous variance rate is $\beta^2 I_t K_t$. As construction progresses and K_t declines, the variance shrinks, which reflects learning, scope reduction, and progressive de–risking documented by Dixit and Pindyck (1994). This property also implies that uncertainty concentrates early in the build. To numerically simulate the budget evolution and ensure economic realism, we discretize Equation (4.1) using an Euler-Maruyama scheme. Because the continuous-time SDE resembles a square-root diffusion but does not satisfy a strict Feller condition, a standard unconstrained discretization could erroneously produce negative values for the remaining budget. To prevent this, we introduce a max-operator in the discrete-time update:

$$K_{t+\Delta t} = \max\left\{0,\, K_t - I_t\Delta t + \beta\sqrt{I_t K_t \Delta t}\,\varepsilon^{(K)}\right\}, \tag{4.3}$$

where $\varepsilon^{(K)} \sim \mathcal{N}(0,1)$. This specification is justified both economically and mathematically. Economically, K_t represents the remaining required capital; negative values are economically meaningless. Mathematically, the max-operator effectively imposes an *absorbing boundary at zero*. The interpolation step acts as a weak approximation that is strictly consistent with our definition of the stopping time $\tau \equiv \inf\{t \geq 0 : K(t) = 0\}$, representing the exact moment of project completion.

In each period, the decision maker chooses the investment rate I_t from a compact inter-

val $[0, I_{\max}]$. As we will formally establish in Section 3, the absence of convex adjustment costs means that the optimal policy reduces to a binary spending rule: either $I_t = 0$ (pause) or $I_t = I_{\max}$ (build at capacity). This reproduces the result in Dixit and Pindyck (1994) and Zhu (2012) that, absent convex adjustment costs, interior controls are dominated by corner solutions. Under $I_t = I_{\max}$, the expected time to completion is approximately $K_t / I_{\max}$ when diffusion is small, while equation (4.1) preserves residual randomness in τ through the volatility term.

Nuclear Generating Cost

During the operation phase, the nuclear generating cost C_t is composed of two main components: (i) fixed and variable operation and maintenance (O&M) expenditures, and (ii) fuel costs, which include uranium procurement and enrichment. Note that nuclear waste disposal fees are treated separately in the cash flow model as a fixed statutory charge. For most large-scale nuclear plants, O&M costs are relatively stable over time, whereas fuel costs are subject to resource price volatility. Although fuel constitutes a smaller portion of the total levelized cost compared to fossil generation, uranium market prices can fluctuate with changes in global demand and mining output. We assume these fluctuations drive the stochastic behavior of the aggregate generating cost.

Following Zhu (2012), we model the dynamics of the unit generating cost, C_t, under the physical probability measure $\mathbb{P}$ as a Geometric Brownian Motion (GBM):

$$dC_t = \alpha_C \, C_t \, dt + \sigma_C \, C_t \, dW_t^{(C)}, \tag{4.4}$$

where α_C represents the expected long-run growth rate of generating costs, and $\sigma_C > 0$ denotes the volatility parameter. The term $W_t^{(C)}$ is a standard Wiener process representing stochastic cost shocks. We impose the condition $\operatorname{corr}(dW_t^{(K)}, dW_t^{(C)}) = 0$, reflecting the empirical observation that construction phase uncertainties are distinct from, and largely independent of, post-completion operational cost fluctuations.

For valuation purposes, we move to the risk-neutral measure $\mathbb{Q}$ by adjusting for the market price of cost risk. Let ν_C denote the cost risk premium. By applying Girsanov's theorem, the drift is adjusted, and the dynamics are governed by:

$$dC_t = (\alpha_C - \nu_C) \, C_t \, dt + \sigma_C \, C_t \, d\widetilde{W}_t^{(C)}, \tag{4.5}$$

where $\widetilde{W}_t^{(C)}$ is a standard Brownian motion defined under the risk-neutral measure $\mathbb{Q}$. Note that the risk-adjusted drift $(\alpha_C - \nu_C)$ accounts for the systematic risk associated with electricity generation costs.

To ensure numerical stability and positivity of the cost process, we employ the exact so-

lution of the GBM for time discretization rather than a standard Euler–Maruyama approximation. The evolution of C_t over a time step Δt is given by:

$$C_{t+\Delta t} = C_t \exp\left[\left(\alpha_C - \nu_C - \frac{1}{2}\sigma_C^2\right)\Delta t + \sigma_C \sqrt{\Delta t}\,\varepsilon^{(C)}\right], \tag{4.6}$$

where $\varepsilon^{(C)} \sim \mathcal{N}(0,1)$ are i.i.d. standard normal variates. This specification avoids discretization bias and ensures $C_t > 0$ for all t.

Nuclear Accident Risk

In addition to economic and market risks, nuclear power projects face operational safety risks arising from low-probability, high-impact events such as reactor accidents. Although modern third-generation reactor designs offer significant improvements in passive safety systems and containment measures, the potential for accidents cannot be entirely eliminated. Even rare accidents can cause severe financial consequences through plant shutdowns, repair costs, and regulatory penalties.

We model accident risk using a compound Poisson process q_t, which captures both the random occurrence of accidents and the magnitude of associated losses. Over a short time interval dt, the process evolves as:

$$dq_t = \begin{cases} 0, & \text{with probability } 1 - \lambda\, dt, \\ u, & \text{with probability } \lambda\, dt, \end{cases} \tag{4.7}$$

where λ is the average annual accident arrival rate and u denotes the financial loss incurred if an accident occurs.

The loss magnitude u is drawn from a discrete distribution reflecting three accident severity levels:

$$u = \begin{cases} u_S, & \text{with probability } \eta_S, \quad \text{(minor accident)}, \\ u_M, & \text{with probability } \eta_M, \quad \text{(moderate accident)}, \\ u_L, & \text{with probability } \eta_L, \quad \text{(severe accident)}, \end{cases} \qquad \eta_S + \eta_M + \eta_L = 1. \tag{4.8}$$

To integrate the accident model with the optimal control problem and maintain computational tractability within the backward induction framework, we operationalize these risks directly through the cash flow function. We define the instantaneous penalty U_t in the cash flow equation as exactly the realized jump increment, such that $U_t = dq_t$. The three severity levels are implemented as follows:

We categorize these operational risks into three severity levels. A minor accident (u_S) represents a small-scale technical fault or safety incident that does not interrupt plant operation, with the penalty reflecting limited maintenance expenditures. A moderate accident (u_M) involves a fault requiring significant maintenance and safety inspections. While in reality this may reduce generation capacity for an extended period, for mathematical tractability, the economic impact of this downtime and repair is translated into an equivalent lump-sum financial loss subtracted in the period the jump occurs. Finally, a severe accident (u_L) represents a major catastrophic incident involving massive costs for decommissioning, environmental remediation, and compensation. Rather than dynamically altering the state space to permanently terminate the simulation path's capacity, which would disrupt the regression mechanics of the LSM algorithm, we model a severe accident by imposing an overwhelmingly large instantaneous financial penalty u_L. This penalty is calibrated to be sufficiently large to effectively destroy the remaining net present value of the project, thereby mirroring the economic consequences of a permanent shutdown without requiring structural breaks in the state variable evolution.

The compound Poisson formulation thus allows the model to incorporate both the stochastic timing and the random magnitude of accident-related losses. By mapping dq_t directly to U_t, the occurrence of any accident immediately depresses the operating cash flow in that period, thereby properly penalizing the project valuation and influencing the optimal investment timing.

Electricity Price Dynamics

To capture the joint dynamics of spot electricity prices and their associated convenience yield, we adopt the two-factor framework of Schwartz (1997). We acknowledge the inherent limitations of applying this classical model directly to electricity markets. Unlike traditional commodities such as oil or metals, electricity is essentially non-storable, highly spiky, and frequently subject to regime-switching jumps and seasonality. Because the Schwartz (1997) framework fundamentally assumes a storable commodity, the traditional interpretation of the convenience yield as the benefit of holding physical inventory is weaker in this context.

However, to maintain analytical tractability while capturing market realities, we explicitly reinterpret the convenience yield $\delta(t)$ in the context of power markets as a proxy for "capacity value" or security of supply. Let $P(t)$ denote the spot electricity price at time t. The stochastic convenience yield $\delta(t)$ represents the implicit strategic benefit of holding dispatchable, grid-balancing generation capacity (like nuclear energy) rather than relying on volatile spot market purchases during periods of supply scarcity.

Under the physical probability measure $\mathbb{P}$, the model is specified as:

$$\frac{dP(t)}{P(t)} = \big(\mu - \delta(t)\big)dt + \sigma_1 dZ_1(t) \tag{4.9}$$

$$d\delta(t) = \kappa\big(\alpha - \delta(t)\big)dt + \sigma_2 dZ_2(t) \tag{4.10}$$

where: μ is the expected spot price drift; α is the long-term mean level of the convenience yield; $\kappa > 0$ is the speed of mean reversion for $\delta(t)$; σ_1 and σ_2 are the volatilities of $P(t)$ and $\delta(t)$ respectively; $Z_1(t)$ and $Z_2(t)$ are standard Brownian motions with correlation ρ, i.e., $dZ_1(t)\,dZ_2(t) = \rho\,dt$.

The mean-reverting process for $\delta(t)$ follows an Ornstein–Uhlenbeck specification, ensuring that convenience yields do not drift indefinitely but revert toward their long-term mean α. A positive correlation ρ implies that when the convenience yield is high, which indicates greater scarcity or value of dispatchable electricity, the spot price is also likely to be elevated.

For risk-neutral valuation under measure $\mathbb{Q}$, the drift terms are adjusted to incorporate market prices of risk. Denoting λ as the market price of convenience yield risk, the risk-neutral dynamics become:

$$\frac{dP(t)}{P(t)} = \big(r - \delta(t)\big)dt + \sigma_1 d\widetilde{Z}_1(t), \tag{4.11}$$

$$d\delta(t) = \big(\kappa(\alpha - \delta(t)) - \lambda\big)dt + \sigma_2 d\widetilde{Z}_2(t) \tag{4.12}$$

with $d\widetilde{Z}_1(t)\,d\widetilde{Z}_2(t) = \rho\,dt$ under $\mathbb{Q}$.

This specification allows the model to jointly capture short-term price shocks driven by demand fluctuations and market volatility (through $P(t)$), and longer-term structural dynamics in the marginal value of dispatchable electricity (through $\delta(t)$). While alternative approaches such as mean-reverting jump diffusions or structural supply-demand models exist for electricity markets, the reinterpreted two-factor approach is particularly suitable for our nuclear power investment analysis. It avoids importing a commodity model without justification by successfully proxying the economic premium of reliable baseload generation, all while maintaining the necessary computational tractability for the multidimensional Least Squares Monte Carlo (LSM) framework.

Carbon Pricing Dynamics

Unlike standard real option models that treat carbon prices as an exogenous geometric Brownian motion (GBM), we adopt a physics-informed framework. This approach explicitly links financial carbon risks to the physical accumulation of greenhouse gases (GHG) and the policy response to environmental tipping points.

The theoretical foundation for modeling environmental irreversibility and optimal policy timing has been well-established in the literature. Pindyck (2000) pioneered the analysis of policy timing under environmental irreversibility, demonstrating that the inability to reverse environmental damage creates a significant "option value" to reduce emissions early. Building on this, Saphores (2004) and Saphores and Carr (2019) further explored the optimal timing of emission limits under ecological uncertainty, highlighting how stochastic environmental processes influence the decision to implement regulations. However, these seminal works typically assumed a constant rate of natural pollutant decay, implying that the environment's ability to recover remains invariant regardless of the pollution stock.

We extend this lineage by introducing a critical innovation: the modeling of the natural decay rate, $\delta(M_t)$, not as a constant, but as a state-dependent function. This captures the concept of environmental irreversibility in a more physically rigorous manner: as the concentration M_t exceeds a critical tipping point M^*, the Earth's capacity to absorb carbon (e.g., via ocean sinks or forests) diminishes.

The core of their derivation rests on the dynamics of the aggregate atmospheric pollutant stock, M_t. Unlike standard asset price processes, M_t follows a controlled stochastic differential equation (SDE) that reflects the physical laws of accumulation and decay:

$$dM_t = \beta(E_t - \delta(M_t)M_t)dt + \sigma_M dW_t^M \tag{4.13}$$

Here, E_t represents the global emission rate, and β is a conversion factor. The state-dependent decay rate $\delta(M)$ is mathematically formalized as:

$$\delta(M) = \begin{cases} \delta_0, & \text{if } M \leq M^* \\ \delta_0 \exp(-\gamma(M - M^*)), & \text{if } M > M^* \end{cases} \tag{4.14}$$

where $\gamma > 0$ determines the speed of collapse. M^* is a pre-defined constant parameter representing the critical atmospheric concentration threshold (e.g., 450 ppm) beyond which the natural carbon sink capacity degrades irreversibly. While M^* represents the objective physical threshold, the exact atmospheric concentration that will trigger this collapse is not known with certainty by market participants. Instead of operating under a standard rational expectations framework with a known probability measure $\mathbb{P}$, agents face Knightian uncertainty regarding the true distribution of M^*. To incorporate this, we interpret the pricing dynamics through the lens of ambiguity aversion. Let $\mathbb{Q}$ denote the worst-case robust measure adopted by ambiguity-averse agents solving a max-min expected utility problem (Hansen and Sargent, 2001). The discrepancy between the objective expectation of climate damages and the expectation under the robust measure $\mathbb{Q}$ mathematically generates the "fear premium". As M_t approaches the estimated vicinity of M^*, the set of plausible worst-case distributions expands, causing the shadow price of carbon π_t to exhibit a highly convex jump (Lemoine and Traeger, 2014). Thus, the fear premium is not merely a reflection of expected cash flow losses,

but a formal asset pricing premium demanded by agents to compensate for the unquantifiable structural ambiguity of a permanent ecological regime shift.

To fully capture the uncertainty in climate policy decisions, the framework incorporates a second state variable, θ_t, which accounts for stochastic economic fluctuations and randomness in technological innovations. Following the mathematical framework, θ_t evolves according to a geometric Brownian motion (GBM):

$$d\theta_t = \eta\theta_t dt + \kappa\theta_t dW_t^{\theta} \tag{4.15}$$

where η represents the drift and κ denotes the volatility of these economic shocks. Crucially, the model allows for interdependence between the environmental and economic systems. The stochastic driver of the pollutant stock (W_t^M) and the driver of the economic variable (W_t^{θ}) are correlated with a coefficient ρ:

$$d\langle W_t^M, W_t^{\theta}\rangle = \rho dt \tag{4.16}$$

This correlation ρ captures the real-world phenomenon where economic activity often moves in tandem with environmental pressures. Consequently, the social benefit function $B(M_t, \theta_t)$ in the optimization problem depends on both the physical state of the climate and the prevailing economic/technological conditions.

The theoretical basis for the carbon price in this framework is the Social Cost of Carbon (SCC), π_t, which is derived from solving the central social planner's optimal control problem aimed at maximizing total expected discounted social welfare, $W(M_t, \theta_t)$. The objective functional is defined as maximizing the total expected discounted social welfare. While the general problem allows for continuous emission adjustments, our framework follows Lemma 2.2 of the source model, which restricts the control to a binary switch (e.g., switching from a high-emission to a low-emission regime). This transformation simplifies the general control problem into an optimal stopping problem:

$$V = \sup_{\tau} \mathbb{E}\left[\int_0^{\infty} e^{-rt} B(M_t^{\tau}, \theta_t)dt - \int_0^{\infty} k(E_t^{\tau})dt\right], \tag{4.17}$$

where τ represents the optimal timing to intervene. This formulation is mathematically equivalent to the American option valuation structure, justifying our use of Real Options Analysis.

The dynamics of this value function are governed by the Hamilton-Jacobi-Bellman (HJB) equation. In the continuation region (prior to the policy intervention τ), the welfare function $W(M_t, \theta_t)$ satisfies:

$$rW = \sup_{E}\left\{ B(M_t, \theta_t) - k(E) + \beta(E - \delta(M_t)M_t)\frac{\partial W}{\partial M} + \mathcal{L}_{\theta}W + \frac{1}{2}\sigma_M^2 \frac{\partial^2 W}{\partial M^2}\right\}, \tag{4.18}$$

where $\mathcal{L}_{\theta}$ represents the differential operator associated with the economic variable θ_t. This

HJB equation explicitly connects the instantaneous welfare flow to the expected evolution of the system, highlighting how the physical accumulation term $\beta(E - \delta(M_t)M_t)$ interacts with the marginal value of the pollutant stock, $\frac{\partial W}{\partial M}$.

While the theoretical objective function is defined over an infinite horizon, solving this numerically requires a dual-horizon approximation. Consistent with the Modified LSMC Algorithm proposed in the framework, we approximate the infinite integral by truncating the valuation horizon to T_2 (where $T_2 \gg \tau$), allowing for a computationally feasible solution while capturing the long-term inertia of the climate system.

In the continuous-time HJB framework, the SCC, π_t, is identified as the shadow price of the pollutant stock:

$$\pi_t(M_t, \theta_t) = -\frac{\partial W(M_t, \theta_t)}{\partial M_t}. \tag{4.19}$$

At this juncture, it is critical to distinguish between the global social planner's optimization, which endogenously determines the carbon price, and the local firm's investment problem. We assume the nuclear power investor acts as a competitive *price-taker*. Because a single nuclear plant's emissions or avoided emissions are negligible relative to the global atmospheric stock, the investor's actions do not dynamically alter the global trajectory of M_t. Consequently, the investor does not jointly solve the global climate HJB equation alongside its own investment decision. Instead, the firm treats the joint evolution of (M_t, θ_t) and the resulting carbon price $\pi_t(M_t, \theta_t)$ as an exogenous, state-dependent pricing process that feeds directly into its cash flow expectations.

This theoretical shadow price π_t is a continuous, nonlinear function that incorporates the economic penalty of environmental damage. Because the underlying physical decay $\delta(M_t)$ drops precipitously near the threshold M^*, the social value function $W(M_t)$ becomes highly convex. This convexity generates a severe "fear premium"—a sharp acceleration in the carbon price as society anticipates the catastrophic loss of environmental reversibility. Consequently, the function $\pi_t(M_t, \theta_t)$ exhibits a structural break, or a "kink", in its slope and curvature. This explicitly induces non-smoothness in the firm's continuation region, rendering traditional smooth polynomial approximations, such as those used in standard LSM algorithms, fundamentally inadequate.

In our numerical implementation, we approximate this theoretical shadow price using a parametric function shifting between regimes, calibrated to the distance from the tipping point Λ, which represents the critical threshold M^*. To ensure the economic state variable θ_t properly influences the project cash flows, thereby reflecting periods of high or low macroeconomic willingness to pay for climate mitigation, we incorporate θ_t as a multiplicative macroe-

conomic factor normalized around a baseline of 1. The realized carbon price is defined as:

$$\pi_t(M_t, \theta_t) = \theta_t \cdot \begin{cases} \max(0, \pi_{\text{base}} + \alpha_1(M_t - M_0)), & \text{if } M_t \leq \Lambda \quad \text{(Pre-tipping)} \\ \max(\pi_{\text{crisis}}, \pi_{\text{crisis}} + \alpha_2(M_t - \Lambda)), & \text{if } M_t > \Lambda \quad \text{(Post-tipping)} \end{cases} \tag{4.20}$$

By incorporating θ_t, the carbon price reflects both the fundamental physical regime, driven by M_t, and prevailing economic shocks, driven by θ_t. Here, the parameter constraint $\alpha_1 \ll \alpha_2$ captures the sudden, steep increase in the marginal social cost of carbon once the physical system transitions to the collapse regime.

Formally, the structural dependence on the threshold variable M_t, as defined in Equation (4.14), means that the system's behavior is classified as a Threshold Regime-Switching Model (Tan and Wu, 2025). According to the comprehensive classification by Tan and Wu (2025), this model category is characterized by parameters (in our case, the underlying physical decay rate) that vary deterministically based on the state of an observed variable (M_t) relative to a threshold (M^*).

However, our framework diverges fundamentally from standard econometric threshold models, where regimes are typically determined by lagged financial indicators or purely statistical artifacts (Tong, 1978; Tan and Wu, 2025). In stark contrast, the structural break in our model is physics-informed: the threshold variable is the physical pollutant stock M_t, driven by the stochastic laws of accumulation and decay described in Equation (4.13). Unlike reduced-form models that might overlay an exogenous Poisson jump process to simulate catastrophe, our framework effectively generates an endogenous regime switch through the state-dependent decay rate $\delta(M_t)$. When M_t exceeds the tipping point M^*, the decay rate decreases continuously according to a convex function, causing the value function and the resulting shadow price π_t to exhibit jump-like behavior without requiring an external jump process. This explicit and external physical constraint implies that the resulting nonlinear behavior within the Social Cost of Carbon is not a statistical anomaly but rather an economic manifestation of a fundamental physical reality, namely the catastrophic loss of environmental reversibility predicted by climate science.

Federal Nuclear Subsidy

In addition to the implicit social value of avoided carbon emissions, we incorporate explicit, statutory financial incentives available under current institutional regimes. Specifically, in line with recent U.S. federal energy policy, such as the Inflation Reduction Act (IRA), new-generation nuclear power plants frequently qualify for production-based tax credits (PTCs). These credits are designed to directly incentivize zero-carbon baseload generation by providing a fixed financial support rate per unit of electricity generated during a defined eligibility window following the plant's commissioning.

This statutory subsidy is incorporated into our model as:

$$\mathrm{PTC}(t) = \begin{cases} s_0\, Q_{\mathrm{elec}}(t), & \text{if } t \le T_{\mathrm{policy}}, \\ 0, & \text{otherwise.} \end{cases} \tag{4.21}$$

where s_0 is the per-MWh subsidy rate. In the baseline calibration used below, s_0 is set to \$33/MWh, and T_{policy} denotes the duration of the subsidy period in years. Mathematically, this subsidy formula relies on the indicator function, $\mathbf{1}_{\{t \le T_{\mathrm{policy}}\}}$, which equals one if the condition is met and zero otherwise.

It is important to distinguish this localized, direct tax credit from the global physics-informed carbon price (π_t) discussed previously. While π_t captures the theoretical shadow price of climate irreversibility from a social planner's perspective, the PTC represents a tangible, near-term cash flow realization governed by existing tax law. The inclusion of this subsidy effectively reduces the levelized cost of electricity for the project, improving its competitiveness against alternative carbon-intensive generation technologies. In real-options terms, it acts as a powerful catalyst that shifts the optimal exercise boundary for investment: by significantly augmenting early-period cash flows, the PTC mitigates downside risk during the plant's initial operating years, thereby potentially accelerating the optimal timing of capital deployment.

4.3 Model Solution

Once the investment is completed at the stopping time τ, the nuclear plant operates continuously until the end of the valuation horizon T_2. To align with the physics-informed framework which approximates the long-term impacts of climate change, we distinguish between two time horizons: the decision horizon T_1, during which the investment option can be exercised, and the extended valuation horizon T_2 ($T_2 > T_1$), which captures the long-term operational cash flows and the evolution of the climate state variables.

The electricity generation capacity of the plant is determined by a set of technical operational parameters. These parameters include the fixed installed capacity (M_{inst}, denoted in MW), the plant's utilization rate, characterized by the capacity factor ($\mathrm{cap}(t_n)$), and the availability rate ($\mathrm{avai}(t_n)$), alongside the annual operating hours (h).

Specifically, the **capacity factor** is defined as the ratio of the actual electrical energy output over a given period to the maximum possible output if the plant were operated at full nameplate capacity continuously. It reflects the efficiency of plant utilization. The **availability rate** measures the percentage of time the plant is operational and capable of producing electricity, accounting for necessary downtime such as refueling and maintenance outages.

The precise values for $\text{cap}(t_n)$ (set at 95% in the baseline) and $\text{avai}(t_n)$ (set at 90% in the baseline) are calibrated based on empirical performance data for advanced nuclear reactors (see Table 4.1, Appendix).

The electricity generation capacity in period t_n is determined by:

$$Q_{\text{elec}}(t_n) = M_{\text{inst}} \cdot \text{cap}(t_n) \cdot \text{avai}(t_n) \cdot h, \tag{4.22}$$

where M_{inst} is the installed capacity, $\text{cap}(t_n)$ is the capacity factor, $\text{avai}(t_n)$ is the availability rate, and h denotes the annual operating hours.

Assuming the Production Tax Credit (PTC) is included in the taxable income,[1] the instantaneous operating cash flow $CF(t)$ is given by:

$$CF(t) = \Big([P(t) - C(t) - c_{\text{waste}}] \cdot Q_{\text{elec}}(t) + \pi_t(M_t, \theta_t) \cdot \gamma \cdot Q_{\text{elec}}(t) + \text{PTC}(t) \Big)(1 - \text{Tax}) - U_t. \tag{4.23}$$

In this formulation, $P(t)$ and $C(t)$ denote the stochastic electricity price and generating cost, respectively, while c_{waste} represents the unit nuclear waste disposal cost in USD/kWh. The term $\pi_t(M_t, \theta_t)$ constitutes the physics-informed carbon price. We assume the nuclear investor acts as a price-taker whose individual project scale is too small to alter the global trajectory of the pollutant stock M_t. Integrating this term into the firm's cash flow requires careful institutional interpretation to avoid double-counting with baseline electricity prices, which may already partially internalize regulatory costs. Here, we evaluate the project from a social planner's perspective, where the product $\pi_t \cdot \gamma \cdot Q_{\text{elec}}$ represents the monetary value of avoided carbon emissions (i.e., saved carbon taxes) compared to a marginal carbon-intensive generation technology, with γ being the emission reduction factor. Finally, U_t represents the financial penalty jump associated with an operational accident. To maintain computational tractability within the simulation framework, we model even severe accidents as imposing massive instantaneous financial penalties, such as remediation and compensation costs subtracted directly from the period's cash flow, rather than dynamically altering the state space to permanently terminate the asset's life.

To formulate the investment decision as a stochastic control problem, we first define the complete state vector that captures all sources of uncertainty in the physical and financial environment:

$$X_t = [P_t, C_t, K_t, M_t, \theta_t]^\top. \tag{4.24}$$

Next, we define the set of admissible controls. The firm continuously controls the instantaneous investment expenditure rate $I(t)$, chosen from a compact interval representing physical construction constraints, such that $I(t) \in [0, I_{\max}]$.

[1] If the PTC is non-taxable in a specific jurisdiction, it should be excluded from the corporate tax base accordingly.

The firm's objective is to maximize the expected Net Present Value (NPV) by choosing the optimal investment policy $I \in \mathcal{A}$ over the project horizon $[0, T_2]$. The value function $V(t, X_t)$ is formally defined as the supremum of the expected discounted cash flows:

$$V(t, X_t) = \sup_{I \in \mathcal{A}} \mathbb{E}_{\mathbb{Q}} \left[\underbrace{\int_{\tau}^{T_2} e^{-r(s-t)} CF(s)\, ds \cdot \mathbf{1}_{\{\tau < T_2\}}}_{\text{PV of Operational Cash Flows}} - \underbrace{\int_{t}^{\tau \wedge T_2} e^{-r(s-t)} I(s)\, ds}_{\text{PV of Investment Costs}} \,\middle|\, \mathcal{F}_t \right]. \quad (4.25)$$

Theoretically, during the construction phase (where $K_t > 0$), applying Bellman's principle of optimality implies that the value function $V(t, X_t)$ is the viscosity solution to the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\frac{\partial V}{\partial t} + \mathcal{L}_0 V + \sup_{I \in [0, I_{\max}]} \left\{ -I - I \frac{\partial V}{\partial K} + \frac{1}{2} \beta^2 I K \frac{\partial^2 V}{\partial K^2} \right\} - rV = 0, \quad (4.26)$$

where $\mathcal{L}_0$ represents the differential operator for the uncontrolled market and climate variables (P, C, M, θ). The term inside the supremum captures the dynamics controlled by the investment rate I, including the direct capital outflow and the reduction of the remaining budget K. Note that the operational accident jump term U_t does not appear in Eq. (4.26) as it applies only after the plant is commissioned (i.e., when $K_t = 0$).

Crucially, because the HJB equation is linear in the control variable I_t and lacks continuous convex adjustment costs, the optimization yields a "bang-bang" control solution. Consequently, the continuous control problem over I_t simplifies into a sequential optimal stopping problem. At each decision epoch, the firm simply chooses between investing at maximum capacity ($I_{\max}$) or waiting ($I_t = 0$). This fundamental transformation theoretically justifies our use of the Least Squares Monte Carlo (LSM) algorithm, which is inherently designed for optimal stopping problems, to solve this ostensibly continuous investment control problem.

Solving Eq. (4.26) directly involves overcoming two major hurdles: the "curse of dimensionality" inherent in the 5-factor state space, and the structural discontinuity in the coefficients induced by the tipping point mechanism in $\pi_t(M_t)$. To address these challenges robustly, we adopt a simulation-based backward induction framework rooted in the Modified Least Squares Monte Carlo (LSM) method. However, recognizing that standard polynomial basis functions may smooth over the critical regime switch, we extend the classical LSM approach by integrating and comparing three advanced machine learning estimators: Regularized Random Forest, eXtreme Gradient Boosting (XGBoost), and Mixture of Experts (MoE). This comparative approach allows us to rigorously evaluate how different regression architectures approximate the non-smooth value surface and capture the "fear premium" associated with the climate tipping point.

4.3.1 Modified Longstaff–Schwartz Method (LSM)

As discussed in Section 3, the value function $V(t, X_t)$ is theoretically governed by the HJB equation (Eq. 4.26). However, solving this five-dimensional PDE using traditional finite difference methods is computationally intractable due to the curse of dimensionality. To overcome this, we employ a Modified Least Squares Monte Carlo (LSM) algorithm. This method relies on the probabilistic representation of the PDE solution via the Feynman-Kac theorem, approximating the continuous-time optimal stopping problem through discrete-time simulation and regression.

Path Simulation with Extended Horizon

Consistent with the dual-horizon framework, we simulate M paths over the extended valuation horizon T_2 with time discretization Δt. Let N_1 correspond to the decision horizon T_1 (i.e., $t_{N_1} = T_1$) and N_2 correspond to the valuation horizon T_2 (i.e., $t_{N_2} = T_2$).

The simulation evolves the coupled SDEs described in Section 2 up to T_2, which correspond to the stochastic processes driving the differential operator $\mathcal{L}_0$ in the HJB equation. Specifically, we simulate the correlated geometric Brownian motions for the electricity price, generating cost, investment budget, pollutant stock, and the economic variable θ_t.

Crucially, the carbon price π_t is not simulated via an independent stochastic differential equation. Instead, at each time step t, it is derived endogenously from the simulated state variables according to the pricing function defined in the model:

$$\pi_t = \pi(M_t, \theta_t). \tag{4.27}$$

This approach ensures that the simulated carbon price trajectories inherently reflect the nonlinear regime switch dynamics triggered when the pollutant stock M_t approaches the tipping point.

The state vector X_t used for the regression basis during the decision-making phase (up to t_{N_1}) must explicitly match the state space defined in the HJB equation:

$$X_t = [P_t, C_t, K_t, M_t, \theta_t]'. \tag{4.28}$$

Including both M_t and θ_t is essential to capture the correlation ρ between environmental degradation and economic fluctuations. Note that the decay rate δ_t is omitted from the state vector as it is a deterministic function of M_t.

Valuation Procedure

The valuation proceeds by backward induction starting from the end of the decision horizon T_1 (step N_1), rather than the end of the simulation T_2. This backward recursion serves as the discrete-time numerical equivalent to integrating the HJB equation backward from the terminal condition.

Let $V(t_n, X_{t_n})$ denote the value of the investment option at time t_n (for $n \leq N_1$). The discrete-time Bellman equation approximates the continuous supremum operator in Eq. (4.26):

$$V(t_n, X_{t_n}) = \max \left\{ VS(t_n, X_{t_n}),\ e^{-r\Delta t} \mathbb{E}_{\mathbb{Q}} \left[V(t_{n+1}, X_{t_{n+1}}) \mid \mathcal{F}_{t_n} \right] \right\}. \tag{4.29}$$

Immediate Exercise Value (VS) with Extended Cash Flows: If the project investment is completed at t_n (where $t_n \leq T_1$), the investor receives operational cash flows until the end of the valuation horizon T_2. Thus, the exercise value sums the discounted cash flows from t_n to N_2:

$$VS(t_n, X_{t_n}) \approx \mathbb{E}_{\mathbb{Q}} \left[\sum_{k=n}^{N_2} e^{-r(t_k - t_n)} CF(t_k, X_{t_k}) \right] - K_{t_n}. \tag{4.30}$$

This formulation explicitly captures the long-term value of the nuclear asset beyond the decision window, aligning with the operational phase described in the model solution.

Continuation Value (VC): The continuation value represents the expected option value if investment is deferred (equivalent to choosing $I = 0$ in the HJB control set for the current step). This is calculated only within the decision window ($n < N_1$):

$$VC(t_n, X_{t_n}) = e^{-r\Delta t} \mathbb{E}_{\mathbb{Q}} \left[V(t_{n+1}, X_{t_{n+1}}) \mid X_{t_n} \right]. \tag{4.31}$$

At the boundary $t_{N_1} = T_1$, if the option has not been exercised, it expires, so $V(T_1, X_{T_1}) = \max(VS(T_1, X_{T_1}), 0)$.

Estimation via Regression on Physical States

The conditional expectation in Eq. (4.31) is estimated by regressing the realized future payoffs against a set of basis functions. This step is numerically equivalent to approximating the spatial derivatives (gradient and Hessian) in the HJB equation.

Since the carbon price $\pi_t(M_t, \theta_t)$ exhibits structural breaks at the tipping point and depends on the joint evolution of the climate and economy, standard polynomial basis functions of only financial variables are insufficient. To capture the full correlation structure, we define the basis set $\mathcal{L}$ using a full quadratic specification. This explicitly includes the state variables $X_{t_n} = [P_{t_n}, C_{t_n}, K_{t_n}, M_{t_n}, \theta_{t_n}]^\top$ along with all their squares and pairwise cross-products.

The regression approximation is formally given by:

$$\widehat{VC}_{t_n} \approx \beta_0 + \sum_{i=1}^{5} \beta_i X_{i,t_n} + \sum_{i=1}^{5} \beta_{ii} X_{i,t_n}^2 + \sum_{i=1}^{5} \sum_{j=i+1}^{5} \beta_{ij} X_{i,t_n} X_{j,t_n}. \tag{4.32}$$

This specification ensures that all quadratic cross-terms are included, thereby capturing the complete interdependence of the state variables. For instance, the cross-term between the pollutant stock and the economic state, $M_{t_n} \cdot \theta_{t_n}$, is vital as it mathematically represents the correlation ρ between physical climate risks and economic innovations (corresponding to the $\frac{\partial^2 V}{\partial M \partial \theta}$ term in the HJB operator). By including the full set of interactions, the model effectively prices the joint fear premium associated with approaching the tipping point under varying economic conditions, rather than treating these risks in isolation.

4.3.2 Machine Learning-Enhanced Regression Schemes

In the standard Modified LSM framework described in the section above, the continuation value is approximated using a global polynomial basis. While computationally efficient, global polynomials enforce a smoothness assumption across the entire state space. This assumption is fundamentally at odds with the physics-informed nature of our model, where the carbon price $\pi_t(M_t, \theta_t)$ undergoes a discrete structural jump when the pollutant stock M_t breaches the tipping threshold Λ. Furthermore, global polynomials struggle to adequately capture the discontinuous, massive financial penalties (U_t) associated with rare severe nuclear accidents. Polynomials typically fail to fit such step-like discontinuities, resulting in oscillations known as the Gibbs phenomenon that can inadvertently smooth out the "fear premium" associated with the tipping point. To address this, we replace the polynomial regression step in the backward induction with three advanced machine learning estimators selected specifically for their ability to partition the state space nonlinearly.

The first alternative we employ is the Random Forest algorithm (Breiman, 2001). Unlike polynomial regression, which fits a continuous curve, Random Forest aggregates predictions from an ensemble of decision trees. Each tree recursively splits the state space X_t into rectangular regions based on binary rules. This logic allows tree-based models to naturally approximate discontinuous value functions without imposing artificial smoothness. However, a direct application of standard Random Forest faces a challenge in our simulation: the post-tipping regime ($M_t > \Lambda$) represents a tail risk event. In many simulated paths, the system may never reach this state, leading to data sparsity in the region of interest. Standard decision trees, when faced with sparse data, tend to isolate individual paths in leaf nodes, effectively memorizing the Monte Carlo noise rather than learning the structural trend. To mitigate this overfitting, we implement a regularized Random Forest by enforcing a strict minimum leaf size constraint. We require that every terminal node in the tree must contain at least $N_{\min}$ simulated paths (e.g., $N_{\min} \geq 100$). The continuation value is estimated as the average prediction

of B trees:

$$\widehat{VC}^{\mathrm{RF}}(t, X_t) = \frac{1}{B}\sum_{b=1}^{B} T_b(X_t; \Theta_b), \quad \text{subject to } \forall \text{ leaf } l \in T_b,\ |S_l| \geq N_{\min}, \tag{4.33}$$

where T_b is the b-th tree, and $|S_l|$ is the number of samples in leaf l. By forcing the tree to average over a large cluster of paths even in the high-risk regime, this constraint acts as a local smoothing filter, ensuring that the model captures the expected impact of the tipping point rather than the idiosyncratic fluctuations of a single extreme simulation path.

To further refine the approximation, we utilize eXtreme Gradient Boosting (XGBoost) (Chen and Guestrin, 2016). While Random Forest builds trees in parallel to reduce variance, XGBoost builds trees sequentially, where each new tree attempts to correct the residual errors of the previous ensemble. This boosting mechanism is particularly effective at capturing complex, nonlinear multiplicative interactions within the state space—most notably, the compounding effect of the macroeconomic multiplier θ_t amplifying the carbon crisis price when $M_t > \Lambda$. Given the high variance of the post-tipping trajectories, standard boosting can be prone to overfitting outliers. In our implementation, we explicitly control model complexity by incorporating a hybrid Elastic Net regularization into the objective function. At each step k, the algorithm minimizes a regularized loss function where the penalty term $\Omega(f_k)$ includes both L_1 and L_2 regularization. The L_1 term promotes sparsity, helping the model ignore irrelevant feature interactions, while the L_2 term prevents any single leaf weight from becoming too large, stabilizing the prediction in the volatile post-tipping regime. Additionally, we enforce a minimum child weight constraint, ensuring that splits are based on a sufficient sum of instance weights. This configuration allows the model to precisely locate the optimal cut at the tipping point Λ while limiting overfitting in sparsely populated regions of the simulated state space.

Finally, we introduce a tailored regression scheme designed to directly leverage the known theoretical structure of our problem, termed Physics-Informed Threshold Regression. Conceptually, this can be viewed as a hard-gated Mixture of Experts (MoE) model. In a traditional MoE neural network, a gating network learns probabilities to assign input data to different expert networks. However, in our context, we possess prior physical knowledge: we know exactly where the structural break occurs, at $M_t = \Lambda$, because it is defined by the underlying physics module. Therefore, instead of asking a machine learning model to estimate this boundary, which introduces unnecessary statistical error, we impose it as a strict constraint. We construct two specialized expert regression models: a stable regime expert trained exclusively on data where $M_t \leq \Lambda$, and a crisis regime expert trained exclusively on data where $M_t > \Lambda$. The continuation value is then estimated via a deterministic switching function:

$$\widehat{VC}^{\mathrm{TR}}(t, X_t) = \begin{cases} \beta_{\text{stable}}^{\top}\Phi(X_t) & \text{if } M_t \leq \Lambda \\ \beta_{\text{crisis}}^{\top}\Phi(X_t) & \text{if } M_t > \Lambda \end{cases} \tag{4.34}$$

where $\Phi(X_t)$ represents the full quadratic polynomial basis features of the state variables, and β denotes the regression coefficients for each respective regime. This approach is theoretically robust for our specific problem structure as it completely eliminates the smoothing bias inherent in global polynomials. By explicitly separating the stable world from the crisis world, the model allows the option value to be priced based purely on the high-severity dynamics of the post-tipping regime, without being mathematically diluted by data from the stable regime.

4.3.3 Implementation and Comparative Framework

The four distinct valuation methodologies, namely the Modified Least Squares Monte Carlo reliant on polynomials, Regularized Random Forest, Regularized XGBoost, and Threshold Regression informed by physics, are implemented using an identical set of 50,000 Monte Carlo simulation paths. To ensure that any observed divergence in valuation is strictly attributable to the regression architecture rather than sampling artifacts, the variance reduction technique of common random numbers (CRN) is employed across all models. The regression basis at each time step is constructed from the five-dimensional state vector $X_{t_n} = [P_{t_n}, C_{t_n}, K_{t_n}, M_{t_n}, \theta_{t_n}]'$, comprising the spot electricity price, the unit generating cost, the remaining investment budget, the atmospheric GHG concentration, and the macroeconomic state variable, respectively. The explicit inclusion of the pollutant stock M_{t_n} is particularly critical, as it fundamentally drives the regime-switching carbon price dynamics described in the physics-informed pricing kernel.

Unlike neural network-based approaches, the tree-based and polynomial regressors utilized in this framework are invariant to monotonic transformations of the input features; consequently, input standardization is not required. Furthermore, rather than relying on automated cross-validation, which risks overfitting the high-variance noise characteristic of the sparse post-tipping regime, a robust regularization strategy is adopted for hyperparameter calibration. For the Random Forest algorithm, a strict minimum leaf size constraint is enforced to smooth out Monte Carlo noise in the tail regions of the distribution. Similarly, the XGBoost implementation incorporates L_1 and L_2 regularization terms alongside a minimum child weight constraint, preventing the model from isolating single outlier paths. In the case of Threshold Regression, the split point is fixed deterministically at the physical tipping point Λ, thereby eliminating parameter uncertainty regarding the regime boundary.

The performance and economic implications of these models are evaluated along three primary dimensions. First, the expected project value at the initial time is analyzed to compare the polynomial benchmark against the tree-based and threshold methods. Since the polynomial approach tends to artificially smooth discontinuities, the valuation divergence observed in the machine learning estimates serves as a quantitative proxy for the "fear premium"—the additional value derived from reacting precisely to the tipping point mechanism. Second, the

mean optimal stopping time is calculated across all exercised paths to determine how different algorithmic structures alter the investment urgency in response to the looming climate threshold. Finally, estimation stability is assessed via the standard error of the mean and the variance of the valuation distribution, providing a measure of the estimator's robustness against the inherent stochasticity of the high-volatility post-tipping regime. This comparative analysis empirically verifies which architecture best balances the need to capture structural discontinuities against the risk of overfitting numerical noise.

4.4 Model Calibration

The parameters used in the model are calibrated to reflect typical characteristics of nuclear power plants of different scales (small, medium, and large). Table 4.1 summarizes the key input values grouped into five categories: plant size and cost parameters, market and simulation parameters, convenience yield parameters, physics-informed carbon parameters, and operational parameters.

For plant size and cost parameters, the installed capacity is set to 50,000 kW, 300,000 kW, and 1,000,000 kW to represent small modular reactors (SMRs), medium reactors, and large-scale plants, respectively. The corresponding initial investment costs (K_0) are scaled accordingly. The investment cost uncertainty factor β is fixed at 0.5, consistent with previous real option analyses of large infrastructure projects, reflecting the high variance in construction outlays.

A critical update in this study is the adoption of the Modified LSM framework, where the time horizons are explicitly calibrated to the project scale. For the Small scenario (SMRs), we define a **decision horizon** T_1 of 5 years and a **valuation horizon** T_2 of 35 years. For the Medium scenario, these horizons extend to $T_1 = 10$ years and $T_2 = 40$ years. For the Large-scale scenario, reflecting the significant complexity and long-term planning required, we assume $T_1 = 15$ years and $T_2 = 45$ years. To ensure the statistical convergence of the machine learning estimators, specifically for the high-variance post-tipping regimes, we employ $M = 50,000$ Monte Carlo simulation paths across all scales. The risk-free rate is fixed at 5% in line with long-term government bond yields. The electricity price dynamics follow the Schwartz (1997) two-factor model, with spot price volatility σ_P calibrated at 20% and mean-reversion rate κ_P at 0.15. Convenience yield parameters are included to capture the strategic value of dispatchable baseload capacity.

Policy and carbon parameters represent the most significant methodological shift. Instead of modeling carbon prices as a simple exogenous geometric Brownian motion, we calibrate the physics-informed model based on the atmospheric GHG concentration M_t. The initial concentration M_0 is set to 420 ppm, with a critical tipping threshold Λ parameterized at 450 ppm. The carbon pricing function is calibrated to capture the "fear premium" via a regime-

Table 4.1: Calibration of Model Parameters for Different Nuclear Plant Scales

Parameter	Description	Small	Medium	Large
A. Plant Size & Cost Parameters				
$M_{\text{installed}}$	Installed capacity (kW)	50,000	300,000	1,000,000
K_0	Initial investment cost ($M)	350	2,000	6,000
$I_{\max}$	Max yearly investment ($M)	117	400	857
C_0	Initial unit cost ($/kWh)	0.070	0.060	0.055
β	Investment cost uncertainty	0.5	0.5	0.5
B. Market & Simulation Parameters				
T_1	**Decision horizon (years)**	5	10	15
T_2	**Valuation horizon (years)**	35	40	45
M	**Monte Carlo simulation paths**	**50,000**	**50,000**	**50,000**
r	Risk-free rate	0.05	0.05	0.05
P_0	Initial electricity price ($/kWh)	0.10	0.10	0.10
κ_P	Price mean reversion rate	0.15	0.15	0.15
σ_P	Volatility of electricity price	0.20	0.20	0.20
α_C	Drift of generating cost	0.01	0.01	0.01
σ_C	Volatility of generating cost	0.06	0.06	0.06
C. Convenience Yield Parameters				
δ_0	Initial convenience yield	0.01	0.01	0.01
κ	Mean reversion rate	0.25	0.25	0.25
α	Long-run mean of δ_t	0.02	0.02	0.02
σ_δ	Volatility of δ_t	0.20	0.20	0.20
ρ	Correlation: $dW^{(P)}, dW^{(\delta)}$	0.5	0.5	0.5
D. Physics-Informed Carbon Parameters				
M_0	Initial GHG concentration (ppm)	420	420	420
Λ	Tipping point threshold (ppm)	450	450	450
π_{base}	Pre-tipping carbon price ($/tCO$_2$)	20	20	20
π_{crisis}	Post-tipping crisis price ($/tCO$_2$)	100	100	100
α_1	Pre-tipping sensitivity	0.001	0.001	0.001
α_2	Post-tipping sensitivity	2.0	2.0	2.0
E_{global}	Global emission rate factor	5.0	5.0	5.0
σ_M	Climate volatility	1.0	1.0	1.0
s_0	PTC per MWh ($)	33	33	33
E. Operational Parameters				
τ_{tax}	Corporate income tax rate	0.25	0.25	0.25
h	Annual operating hours	8,760	8,760	8,760
cap	Capacity factor	0.95	0.95	0.95
avai	Availability rate	0.90	0.90	0.90
c_{waste}	Waste disposal cost ($/kWh)	0.02	0.02	0.02

switching sensitivity model. The baseline pre-tipping price π_{base} is set at 20 USD/tCO_2 with a low sensitivity ($\alpha_1 = 0.001$) to concentration changes. However, if M_t breaches the threshold Λ, the model triggers a structural break: the price jumps to a crisis level π_{crisis} of 100 USD/tCO_2, and the sensitivity to further emissions increases drastically ($\alpha_2 = 2.0$). The global baseline emission rate E_{global} and climate volatility σ_M are calibrated to ensure that the system has a non-trivial probability of hitting the tipping point within the valuation horizons.

Finally, regarding operational parameters, the corporate income tax rate is fixed at 25%. The plant is assumed to operate 8,760 hours per year with a 95% capacity factor and 90% availability rate, reflecting the high reliability of modern nuclear designs. The nuclear waste disposal cost is set at 0.02 USD/kWh. Additionally, we incorporate the Production Tax Credit (PTC) at a rate of $s_0 = 33$ USD/MWh for the first 10 years of operation, consistent with recent federal incentives for advanced nuclear deployment.

4.5 Model Results and Discussion

Before analyzing the numerical outcomes across the varying nuclear plant scales, it is essential to rigorously define the two primary metrics used to evaluate investment performance: the **Mean Option Value** and the **Mean Optimal Stopping Time ($\bar{\tau}$)**. These metrics collectively quantify not just the financial worth of the project, but the strategic behavior of the rational investor under physical climate constraints.

The Mean Option Value represents the expected risk-neutral value of the investment opportunity at the initial time $t = 0$. Unlike a traditional static Net Present Value (NPV) calculation, which assumes a fixed "now-or-never" investment decision, the Mean Value captures the full value of managerial flexibility. It is calculated as the average payoff across all Monte Carlo simulation paths, discounted back to the present. Crucially, this metric aggregates two distinct types of outcomes: the high-value paths where the option is optimally exercised (construction begins), and the zero-value paths where the project is optimally abandoned due to unfavorable market or climate conditions. Therefore, a high Mean Value reflects not only the project's profitability potential but also the strategic value of the "right to walk away" from downside risks.

The Mean Optimal Stopping Time ($\bar{\tau}$) serves as a proxy for the optimal investment timing strategy. It measures the average duration, in years, that the firm waits before exercising the option to build. Economically, τ represents the equilibrium point where the marginal benefit of waiting for new information equals the marginal cost of delaying revenue. In our physics-informed framework, the "information" being awaited is primarily the resolution of uncertainty regarding the climate tipping point. A larger $\bar{\tau}$ implies that the model assigns a high value to the "wait-and-see" strategy, advising the investor to delay irreversible capital commitment until the stochastic trajectory of the pollutant stock (M_t) clarifies whether a

carbon price crisis will occur. Conversely, a near-zero $\bar{\tau}$ indicates a "rush-to-build" signal, suggesting that the opportunity cost of delaying operational cash flows outweighs the value of acquiring further information.

Table 4.2 summarizes the valuation results derived from the four numerical methods: Modified LSM, Random Forest, XGBoost, and Physics-Informed Threshold Regression. The mean option value exhibits a robust scaling with project size, ranging from approximately $595 million for the Small scenario to over $23.1 billion for the Large scenario.

Table 4.2: Comparison of Real Option Valuation Results (Extended Horizon)

Scenario	Method	Mean Value ($M)	Std Dev ($M)	Mean τ (years)	Ex. Prob	Runtime (s)
Small	Polynomial (LSM)	598.87	98.18	3.42	11.34%	23.05
	Random Forest	596.43	98.36	3.41	11.34%	204.07
	XGBoost	594.86	98.47	3.45	11.34%	32.25
	Threshold Reg.	598.87	98.18	3.42	11.34%	17.02
Medium	Polynomial (LSM)	5,504.94	652.30	7.75	15.14%	23.50
	Random Forest	5,508.46	652.07	7.72	15.14%	402.10
	XGBoost	5,520.72	651.27	7.92	15.14%	58.81
	Threshold Reg.	5,504.94	652.30	7.75	15.14%	26.90
Large	Polynomial (LSM)	22,787.00	2,397.10	12.15	18.62%	30.87
	Random Forest	22,945.18	2,387.27	12.25	18.61%	612.23
	XGBoost	23,124.99	2,376.18	12.62	18.61%	84.54
	Threshold Reg.	22,787.00	2,397.10	12.15	18.62%	38.48

A central observation from Table 4.2 is the broad convergence of the baseline valuations across four distinct approximation methods. The fact that global polynomial smoothing, binary partitioning utilizing decision trees (Random Forest and XGBoost), and strictly gated regression informed by physics yield comparable magnitudes, ranging from approximately $595 million for Small Modular Reactors (SMRs) to over $23.1 billion for projects of immense scale, serves as a powerful robustness check across multiple methods. It empirically demonstrates that the immense value of the nuclear project is driven by robust economic fundamentals and the underlying physical climate dynamics, rather than being an artificial numerical artifact of the highly complex state space comprising five dimensions.

Furthermore, the structural volatility has stabilized. While the standard deviations remain substantial in absolute terms (e.g., approximately $2.38 billion for the Large scale), the coefficient of variation has normalized compared to shorter-horizon simulations. This confirms that while the physics-informed climate model introduces significant “fat-tailed” risks associated with the tipping point ($M_t > \Lambda$), extending the operational horizon allows the project cash flows to better absorb these shocks, transforming extreme volatility into manageable strategic risk.

A critical economic insight emerges from the analysis of the Small scenario (SMRs). With the alignment of operational lifecycles, the mean option value for SMRs has stabilized at approximately $599 million, with an exercise probability of 11.34%. Crucially, the mean optimal stopping time $\bar{\tau}$ has shifted from near-zero in previous short-horizon literature to approximately 3.4 years. This shift indicates that the SMR calibration retains meaningful

strategic flexibility. SMRs are no longer "now-or-never" distress assets; rather, they function as highly viable real options. The 3.4-year waiting period suggests that even for smaller units, there is significant strategic value in observing the resolution of early-stage climate policy uncertainty before committing capital.

The main value of comparing the alternative estimators lies not in their baseline convergence, but in how they differ near the margins of the non-smooth regime. In the Large scenario, while the Polynomial benchmark and Threshold Regression yield identical valuations of approximately \$22.79 billion, the XGBoost model identifies a higher option value of \$23.12 billion. In traditional corporate finance, a 1.5% valuation difference might be dismissed as marginal; however, in the context of climate irreversibility, this \$338 million discrepancy illustrates the valuation effect of representing the crisis-region convexity differently across numerical methods. Global polynomial bases mathematically smooth out the precipitous structural break at the climate tipping point. In contrast, the gradient boosting algorithm, leveraging its sequential binary splits, precisely isolates the high-value convexity of the crisis paths where $M_t > \Lambda$, capturing the extreme tail value that traditional methods systematically underestimate.

The valuation difference is also accompanied by a change in the implied investment timing. In the Large scenario, XGBoost advises waiting significantly longer ($\bar{\tau} \approx 12.62$ years) compared to the Polynomial benchmark ($\bar{\tau} \approx 12.15$ years). This divergence of nearly half a year shows that model choice can affect the implied timing recommendation. By better capturing the immense upside convexity of the payoff function during a future carbon crisis, XGBoost places a higher premium on the "option to wait" for information. It warns the decision-maker that the uncertainty surrounding the tipping point is too severe, advising a delay in irreversible capital commitment until the physical climate trajectory becomes clearer. Conversely, the smoothing models underestimate the crisis regime's upside potential and consequently recommend premature market entry.

Finally, the exercise probability demonstrates a logical scaling effect, rising from 11.34% in the Small scenario to 18.62% in the Large scenario. This pattern indicates that the exercise frequency rises with project scale in the present calibration; it should not, by itself, be interpreted as proving that the largest plants are universally the most effective hedge against climate risk. The massive economies of scale associated with these larger reactors provide a stronger buffer against baseline market fluctuations, allowing them to capitalize more frequently on the fear premium generated by pricing mechanisms informed by physics. Furthermore, the success of the Threshold Regression informed by physics, which perfectly matches the stable baseline by imposing the physical boundary ($M_t = 450$ ppm) as a strict mathematical constraint, suggests that explicitly incorporating known physical thresholds can provide a computationally stable and economically interpretable benchmark.

Figure 4.1 reports the sensitivity of the nuclear option value to volatility and the discount

rate. The figure complements the baseline valuation results by showing how the option value responds to changes in two economically important parameters. Because the Large scenario uses a 45-year valuation horizon in the baseline calibration, references to a separate 30-year greenhouse-gas-path figure are omitted here unless such a figure is added explicitly.

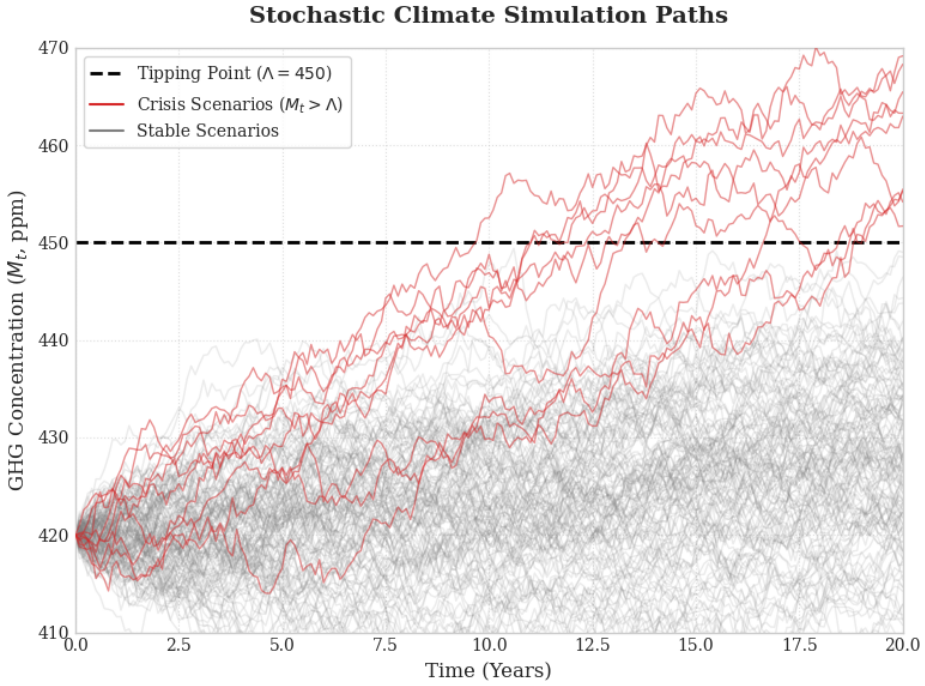


Figure 4.1: Sensitivity of Nuclear Option Value to Volatility and Discount Rate

4.6 Conclusion

This paper has presented a novel real options framework informed by physics for valuing nuclear power investments in an era of intensifying climate change. By moving beyond traditional exogenous commodity models and appropriately treating the investor as a price taker facing a social cost of carbon contingent upon the environmental state, we explicitly linked the economic value of assets with low carbon emissions to the physical risk of crossing critical greenhouse gas thresholds.

Our numerical experiments lead to three primary conclusions regarding the economics of nuclear power under physical climate constraints. First, the inclusion of physical tipping points introduces a substantial "fear premium." The resulting real option valuations are significantly higher than those derived from traditional exogenous models but are accompanied by extreme structural volatility. The fat-tailed distribution of outcomes, driven by the binary risk of staying in a stable regime versus jumping to a crisis regime, underscores the critical strategic role of large-scale nuclear infrastructure. Furthermore, by incorporating the macroeconomic multiplier (θ_t) and massive instantaneous financial penalties for operational accidents (U_t), the model realistically captures the nonlinear variances that make carbon-free baseload capacity an essential hedge against catastrophic climate and market scenarios.

Second, the study reveals the importance of aligning decision and operational horizons across project scales. In the Small Modular Reactor (SMR) calibration, the mean optimal waiting time is approximately 3.4 years across the four estimators, indicating a meaningful option to wait before committing capital. Nonetheless, large-scale nuclear projects benefit the most from the option to wait, leveraging their immense economies of scale and longer operational lifespans to capitalize fully on the resolution of climate uncertainty.

Third, the systematic methodological comparison across four distinct valuation algorithms, namely the Modified Least Squares Monte Carlo, Regularized Random Forest, XGBoost, and Threshold Regression informed by physics, provides an endogenous validation of the structural complexity embedded within the model. Because the continuous investment control problem mathematically reduces to an optimal stopping problem, applying backward induction reliant on the Least Squares Monte Carlo method is theoretically justified. Within this framework, while Threshold Regression and Polynomial estimators yielded similar baseline valuations, the eXtreme Gradient Boosting algorithm, which utilizes decision trees, produced a distinct response to the convexity around the tipping point. Unlike polynomial approximators that tend to artificially smooth the value surface across the transition, XGBoost effectively isolated the crisis paths associated with exceptional value. In the Large scenario, XGBoost produces an option value approximately $338 million above the polynomial benchmark. This difference should be interpreted as evidence that the treatment of the non-smooth crisis regime matters for valuation, rather than as a separately identified causal premium.

Finally, this study highlights the strategic implications of model selection on investment timing. The gradient boosting estimator consistently recommended a longer waiting period before investment compared to the benchmark smoothing models (e.g., advising an additional half-year delay in the large-scale scenario). This divergence suggests that correctly quantifying the sharpness of the carbon price jump reveals a fundamentally higher value in delaying irreversible commitments. Models that rely on smooth approximations tend to underestimate the upside potential of the crisis regime and consequently recommend premature capital deployment. Ultimately, this study bridges the gap between physical climate science and financial engineering, showing that alternative machine-learning estimators can be useful complements to conventional continuation-value approximations when the state space contains nonlinear regime boundaries.

Chapter 5

Organizational Capital and Stock Price Crash Risk

5.1 Introduction

Organizational capital has been a significant topic in management and economics research for nearly five decades (Spence, 1977). However, despite its important implications for corporations, the concept has received relatively limited attention in the mainstream finance literature until recently. Broadly defined, organizational capital (OC) represents the accumulation of exceptional firm-specific knowledge, distinct from physical capital, that is embedded in the firm's structures, processes, culture, and employee networks (Lev and Radhakrishnan, 2005; Lev et al., 2009). It serves as an "enabling asset" that facilitates coordination and transforms individual labor and technological inputs into productive capabilities beyond what tangible assets alone can achieve (Bloom et al., 2012; Hasan and Cheung, 2018a).

Crucially, unlike physical assets which can be easily transferred or sold, organizational capital is often "embodied" in the firm's key talent and workforce (Eisfeldt and Papanikolaou, 2013). It represents a form of firm-specific human capital that creates a distinct competitive moat but also introduces unique risks associated with employee retention and managerial incentives. This intangible asset encompasses managerial routines, proprietary communication systems, and unique coordination mechanisms (Atkeson and Kehoe, 2005). By enhancing decision-making efficiency and information processing, organizational capital acts as a production factor that unlocks sustainable competitive advantage (Bloom et al., 2012; Bauer et al., 2025). Recent empirical evidence further substantiates that firms with high organizational capital exhibit superior operational performance, more informed investment decisions, and increased innovation output, ultimately delivering significant value to long-term investors (Lev et al., 2009; Bauer et al., 2025).

Recent corporate episodes illustrate how operational failures, weak internal systems, and the release of adverse information can coincide with substantial losses in firm value. For example, Intel reported large losses in its foundry operations and experienced a sharp de-

cline in market value (Reuters, 2025), while Boeing faced delivery disruptions and material share-price declines amid continuing production problems (Reuters, 2023). Such episodes motivate the broader question examined in this chapter: whether firm-specific organizational capabilities and information systems are associated with a lower likelihood of extreme negative stock-price outcomes. These examples are illustrative rather than causal evidence of organizational capital itself.

We build on prior research to investigate an important question: do firms with stronger organizational capital face a lower risk of future stock price crashes? Our study empirically explores this relationship.

Research on stock price crash risk has identified agency conflicts between corporate insiders and outside investors as a key factor influencing managerial behaviour. These conflicts often incentivize managers to withhold negative information to protect stock prices, leading to the accumulation of unfavourable data that external investors are unaware of (Jin and Myers, 2006; Bleck and Liu, 2007; Hutton et al., 2009). As insiders misappropriate company resources and manipulate or conceal bad news, this buildup can reach a tipping point. Once the concealed information becomes sufficiently significant or management's motivation to hide it fades, the negative news is released to the market, resulting in substantial stock price declines (Ball, 2009).

Recent literature highlights several intangible factors that contribute to reducing crash risk in organisations. Hasan et al. (2022) reveal that brand capital diminishes crash risk by fostering stakeholder trust and enhancing transparency. Furthermore, Shi and Xia (2023) demonstrate that the quality of human capital can alleviate the impacts of bad-news hoarding and overinvestment. Lastly, Jebran et al. (2022) suggest that the social capital of external board members mitigates crash risk by improving monitoring practices and the flow of information.

While numerous studies have examined the factors contributing to stock price crashes, little empirical evidence exists on how corporate intangible assets, particularly organizational capital, influence these crashes. Chen et al. (2024) suggest that organizational capital may increase the risk of stock price crashes, contradicting earlier studies which claim that intangible assets like brand equity, human capital, and social capital improve firms' information environments and governance quality, ultimately reducing the likelihood of crashes (Hasan et al., 2022; Shi and Xia, 2023; Jebran et al., 2022). This inconsistency may reflect institutional differences between the Chinese and U.S. settings or differences in how organizational capital interacts with managerial opacity and governance. The conflicting evidence therefore motivates a re-examination of the relation between organizational capital and crash risk in a setting with comparatively developed disclosure and monitoring institutions.

We develop two complementary arguments for why organizational capital may reduce future stock price crash risk. The first is grounded in information asymmetry theory: organizational capital may improve internal information processing and transparency, thereby

reducing managers' ability or incentive to hoard bad news. Information asymmetry theory posits that the likelihood of a stock price crash arises when managers possess private information withheld from investors (Jin and Myers, 2006). Managers typically hide adverse information to protect their private benefits or maintain firm reputation, which gives rise to a gradual accumulation of hidden bad news that is released all at once, thus creating a crash in stock price.

The finance literature identifies both internal and external determinants of crash risk. Internal factors include managerial incentives (Kim et al., 2011b), accounting conservatism and CEO overconfidence (Kim et al., 2016), financial-statement comparability (Kim et al., 2011a), CSR disclosure (Kim et al., 2014), and board effectiveness (Anderson et al., 2009). External factors include product-market competition, stock liquidity, and local religiosity (Li et al., 2017; Chen et al., 2017; Fang et al., 2017). Within this setting, organizational capital may operate as an internal mechanism that reduces information asymmetry. A firm with high organizational capital may possess more systematized structures, open communication channels, and stronger managerial controls, which can collectively facilitate timely information dissemination (Lev et al., 2009; Eisfeldt and Papanikolaou, 2013; Bloom et al., 2012; Hasan and Cheung, 2018a) and, in turn, reduce managerial opportunism. By improving the internal and external flow and quality of information, organizational capital may reduce both the incentive and the opportunity for managers to withhold bad news, thereby lowering the likelihood of a sudden stock-price crash (Hutton et al., 2009).

The second argument draws on the resource-based view (RBV), which proposes that firms can obtain sustainable competitive advantages from valuable, rare, inimitable, and non-substitutable resources (Barney, 1991). Organizational capital embodies firm-specific knowledge systems, managerial routines, and employee capabilities that may strengthen operational resilience and strategic adaptability. Prior research links organizational capital to production efficiency and productivity (Black and Lynch, 2005), long-run value creation (Lev et al., 2009), skill development and innovation (Caroli and Van Reenen, 2001), employee retention and skill diversity (Carlin et al., 2012), stock returns (Eisfeldt and Papanikolaou, 2013), and investment–cash-flow sensitivity (Attig and Cleary, 2014). Organizational capital can also be risky because its value is partly embodied in key employees and firm-specific capabilities (Eisfeldt and Papanikolaou, 2014). On balance, however, stronger organizational capabilities may reduce idiosyncratic risk (Luo and Bhattacharya, 2009) and lower the probability of extreme negative outcomes.

To examine these arguments, we use 13,489 firm-year observations for U.S. listed firms over 1999–2022. Organizational capital is measured as the accumulated value of a proportion of past selling, general, and administrative (SG&A) expenditures using the perpetual-inventory method (Peters and Taylor, 2017). Following prior studies (Hutton et al., 2009; Jin and Myers, 2006), the main analysis uses two firm-specific crash-risk measures derived from idiosyncratic weekly returns. Negative conditional skewness (NCSKEW) captures asymme-

try in the distribution of firm-specific returns, while down-to-up volatility (DUVOL) compares the volatility of firm-specific returns in down weeks with that in up weeks. The empirical results reveal a significant negative association between organizational capital and both measures of future crash risk.

Our study contributes to the literature in three important ways. First, we extend the growing body of research on the determinants of stock price crash risk. While prior studies have examined factors such as financial reporting opacity (Hutton et al., 2009), tax avoidance (Kim et al., 2011b), and general intangible intensity (Wu and Lai, 2022), the specific impact of organizational capital remains underexplored. Unlike recent arguments that intangibles exacerbate crash risk due to valuation complexity, our findings provide novel evidence that organizational capital acts as a stabilizing force. By enhancing internal information flows and coordination, organizational capital reduces the likelihood of extreme negative stock returns.

Second, we contribute to the burgeoning literature on the economic consequences of organizational capital. Previous research documents that organizational capital enables firms to obtain high stock returns (Eisfeldt and Papanikolaou, 2013), reduces the cost of equity (Attig and El Ghoul, 2018), and facilitates better acquisition performance (Li et al., 2018). We expand this stream by documenting a distinct "governance value" of organizational capital. Our results suggest that organizational capital does not merely represent a risky investment; rather, it serves as a critical intangible resource that aligns internal processes and mitigates agency conflicts related to bad news hoarding.

Finally, our study contributes to the intersection of human capital and corporate finance by highlighting the conditioning role of managerial characteristics. We show that the efficacy of organizational capital is contingent upon the CEO's skill set and incentive structure. Specifically, we find that Generalist CEOs—who possess transferable skills and broader cognitive frameworks—are better equipped to leverage organizational capital to reduce downside risk. Similarly, appropriate compensation incentives are essential for motivating managers to utilize these intangible resources for risk management rather than opportunism. These findings offer practical implications for corporate boards and investors, suggesting that the value of intangible assets is maximized when matched with capable leadership and aligned incentives.

The remainder of this paper is organized as follows. Section 2 reviews the related literature and develops our hypotheses. Section 3 describes the data, variable measurement, and empirical strategy. Section 4 presents the empirical results and mechanism analyses. Section 5 discusses robustness checks and endogeneity tests. Section 6 concludes the paper.

5.2 Literature Review

5.2.1 Organizational Capital

Organizational capital refers to the accumulation of firm-specific knowledge that enhances operational efficiency, investment success, and innovative performance. This includes technologies such as business practices, processes, and designs (Lev et al., 2009). Recent studies emphasize the crucial role of organizational capital in the growth of both corporate and global capital markets. For example, research by Peters and Taylor (2017) and Eisfeldt and Papanikolaou (2014) demonstrates its growing significance in the U.S. and global capital markets. Furthermore, Lev and Radhakrishnan (2005) and Lev et al. (2009) show that investing in organizational capital is essential for maintaining competitive advantages and achieving superior performance.

5.2.2 Stock Price Crash Risk

Stock price crash risk involves a rapid and uneven decline in a company's stock price following the unexpected disclosure of previously hidden negative information (Jin and Myers, 2006; Hutton et al., 2009; Kim et al., 2011b). Therefore, stock price crash risk is the likelihood that individual stocks will experience significant future price declines, which can have far-reaching effects on capital markets. Prior research underscores the significance of reducing the information asymmetry between managers and external investors to mitigate the risk of stock price crashes. Empirical evidence supports this viewpoint, demonstrating that improvements in firm-level information factors such as enhanced financial reporting transparency (Kim et al., 2014), better accounting comparability (Kim et al., 2016), heightened accrual quality (Hutton et al., 2009), effective tax planning (Kim et al., 2011a), cultural background (Fu and Zhang, 2019), brand, social and human capital (Hasan et al., 2022; Shi and Xia, 2023; Jebran et al., 2022) and increased analyst coverage (Kim et al., 2019) are associated with a reduced likelihood of firm-specific crash events.

5.2.3 Theoretical Underpinning and Hypotheses Development

Organizational capital (OC) combines a company's accumulated intangible assets, such as embedded knowledge, technological skills, management expertise, cultural values, operational routines, and systemic frameworks. These are combined with human skills and physical resources to boost efficiency and competitiveness.

We offer two theoretical arguments to establish a negative relation between organizational capital and stock price crash risk. Initially, the information asymmetry theory suggests that

stock price crashes occur when managers withhold negative information until it accumulates and is eventually released all at once (Jin and Myers, 2006). Although the existing literature highlights the positive aspects of organizational capital, it also shows that organizational capital mitigates this mechanism through strong reporting practices, cross-functional information sharing, and reliable communication channels that improve the timeliness and verifiability of firm information (Hutton et al., 2009), high OC to lower equity capital costs through better management practices (Attig and El Ghoul, 2018), improved post-acquisition performance when acquirers integrate OC from targets (Li et al., 2018), variation across different firm life-cycle stages (Hasan and Cheung, 2018a), and more generous payout policies (Hasan and Uddin, 2020). These findings support the idea that OC enhances information credibility, cash-flow reliability, and governance. In asset valuation, OC is an economically significant yet risky intangible asset (Eisfeldt and Papanikolaou, 2013), with its fragility recognised as a distinct risk factor (Boguth et al., 2022). Overall, these studies indicate that OC promotes coordination, monitoring, and learning, reducing the incentives for managerial bad-news hoarding- a key factor behind stock price crashes (Jin and Myers, 2006; Hutton et al., 2009).

Second, the resource-based view (RBV) defines organizational capital (OC) as a valuable, rare, and inimitable asset that enhances adaptive learning and organisational stability (Barney, 1991; Eisfeldt and Papanikolaou, 2013). From the perspective of intangible capital, these assets positively contribute to stock value (Crouzet et al., 2022). OC boosts a firm's competitiveness, resulting in higher stock returns (Chan et al., 2022; Francis et al., 2021), and Marwick et al. (2020) observe that firms with more organizational capital tend to hold larger cash reserves. Recent studies also shed light on various intangible factors that mitigate crash risk in organisations. Hasan et al. (2022) find that brand capital reduces crash risk by building stakeholder trust and increasing transparency. Similarly, Shi and Xia (2023) show that high-quality human capital can mitigate the effects of bad-news hoarding and overinvestment. Lastly, Jebran et al. (2022) suggest that social capital, especially among external board members, lowers crash risk by enhancing monitoring and information flows, aligning with the idea that greater internal organizational capital reduces downside risk and the likelihood of stock price crashes (Luo and Bhattacharya, 2009).

A related stream of research examines determinants of crash risk including financial reporting quality (Hutton et al., 2009), local religiosity (Callen and Fang, 2015), product-market competition (Ma et al., 2024), audit-industry specialization (Robin and Zhang, 2015), equity incentives and short-term compensation (Callen and Fang, 2015; Kim et al., 2011b), tax planning (Kim et al., 2011a), CEO overconfidence (Kim et al., 2016), and internal-control systems (Chen et al., 2017). Improving coordination and internal trust enhances organizational capital, aligning with the resource-based view (RBV). This strengthens a firm's adaptive and integrative capabilities and reduces the risk of considerable negative outcomes, such as future stock price crashes. Therefore, we predict that increased organizational capital improves information quality and decreases stock price crash risk. Based on the above arguments, we develop the following hypothesis:

H1a. Organizational capital is negatively associated with future stock price crash risk.

In contrast to the above arguments, we present reasons to argue that firms with high levels of organizational capital are positively associated with stock price crash risk. Expenditures on organizational capital (OC), often categorised under selling, general, and administrative expenses (SG&A), can disguise their strategic purpose and increase information asymmetry (Aboody and Lev, 2000; Barth et al., 2001). These investments might concentrate value among key individuals, heightening agency conflicts (Eisfeldt and Papanikolaou, 2013). Additionally, firms with complex organizational structures are more capable of hiding negative information, and high levels of intangible assets can facilitate earnings manipulation, leading to sudden market adjustments (Lee et al., 2007; Kim et al., 2011b). Overall, greater reliance on intangible assets tends to increase information asymmetry (Wu and Lai, 2022). Based on the above arguments, we develop the following hypothesis:

H1b. Organizational capital is positively associated with future stock price crash risk.

CEOs play a crucial role in a firm's decision-making and in managing operations and resource efficiency. CEOs with broad experience across multiple functions, industries, or companies tend to have broader cognitive frameworks, diverse social networks, and improved information-processing skills (Custódio et al., 2013). CEOs who develop transferable skills across various firms and industries are less likely to conceal bad news and may adapt better to volatile economic conditions. Furthermore, research indicates a negative association between CEOs' general ability and the likelihood of future stock price crashes (Fang et al., 2023). This study extends the existing literature by examining how a generalist CEO influences the negative relationship between organizational capital and stock price crash risk, driven by managers hoarding bad news. CEOs with higher or superior educational qualifications are associated with environments characterised by low information transparency and greater capacity to raise capital (Gounopoulos et al., 2021).

From the Knowledge-Based View, Grant (1996) posits that knowledge is found within individuals and organisational structures. CEOs with generalist backgrounds act as boundary-spanners, linking individual and organizational knowledge to coordinate action (Kor and Mesko, 2013). Since generalist CEOs possess extensive experience, knowledge, and skills applicable across different firms and sectors, they are more likely to seek new opportunities, initiate innovative projects, and undertake higher-risk ventures (Xu, 2023). Betzer et al. (2020) found that greater skill disparities between incoming and outgoing CEOs are associated with higher abnormal stock returns at the time of the turnover announcement and better operating performance afterwards. Therefore, generalist CEOs' dynamic, innovation-driven strategies often demand complex structures and diverse operations (Xu, 2023). Therefore, generalist executives improve firm adaptability and balanced disclosure (Custódio et al., 2013), thereby

strengthening transparency within organizational capital. Building on these CEO experiences, they can better leverage organizational capital to increase transparency, streamline information flows, and reduce undisclosed negative information. Consequently, firms with strong organizational capital led by generalist CEOs are better positioned to communicate relevant information effectively, which can further reduce crash risk. Based on the above arguments, we develop the following hypothesis:

H2. CEO generalist experience strengthens the negative association between organizational capital and the likelihood of future stock-price crashes.

Mechanism of the CEO's Incentive role: Compensation structure for CEOs, whether equity- or cash-based, plays a central role in shaping managerial behaviour regarding information disclosure and risk-taking. Drawing on the Resource-Based View (Barney, 1991), this study argues that CEO incentives enhance the effectiveness of organizational capital in mitigating stock price crash risk. While stock prices are generally expected to reflect a company's fundamental value (Kim et al., 2011b), managers' control over information flow can cause prices to deviate from their intrinsic value for extended periods (Kim et al., 2011b; Peng and Roll, 2008). Previous studies have found managerial equity incentives to be associated with earnings manipulation (Bergstresser and Philippon, 2006; Burns and Kedia, 2006; Jiang et al., 2010), and found that greater sensitivity to stock options is associated with an increased risk of stock price crashes (Benmelech et al., 2010; Kim et al., 2011a). Similarly, Jia (2018) finds that wider pay gaps between CEOs and other executives, reflecting more substantial promotion incentives, are associated with elevated crash risk, though this effect is weaker in firms with strong governance.

In fact, Jin and Myers (2006) model suggests that opacity does not necessarily increase crash risk when managers act ethically and in shareholders' interests. Such behaviour is more likely in firms with higher levels of organizational capital, characterised by knowledge sharing, open communication, and established routines (Eisfeldt and Papanikolaou, 2013). These organizational attributes foster transparency and reduce the incentives to withhold negative information. Well-designed CEO compensation structures, regardless of their form, reinforce these mechanisms by aligning managerial wealth with firm value, thereby enhancing the disciplinary and informational role of organizational capital (Barney, 1991). Both equity-based and cash-based incentives that are effectively aligned with long-term firm performance strengthen the negative association between organizational capital and stock price crash risk. When a firm holds a higher level of organizational capital, which mitigates CEO incentives, addresses stock price crash risk, promotes transparency and sustains investor confidence, thereby reducing the likelihood of crash events (Kim et al., 2011b; Andreou et al., 2016). Based on the above arguments, we develop the following hypothesis:

H3. CEO compensation incentives (both equity- and cash-based) strengthen the negative association between organizational capital and stock price crash risk.

Agency theory suggests that appropriate compensation contracts can align the interests of managers with those of shareholders, thereby mitigating agency conflicts and reducing opportunistic behaviors (Jensen and Meckling, 1976; Holmström, 1979). If organizational capital serves as a governance mechanism that curbs managerial bad news hoarding, as posited in our first hypothesis, then the effectiveness of this mechanism should vary with the strength of managerial incentives.

Prior literature offers two competing views on the relationship between CEO compensation and crash risk. The managerial opportunism view suggests that high equity incentives may induce managers to withhold bad news to support short-term stock prices, thereby increasing crash risk (Kim et al., 2011b). However, the interest alignment view argues that optimal compensation packages encourage managers to focus on long-term value creation and transparency, reducing the likelihood of future crashes (Cheng, 2004). In the context of organizational capital, we conjecture that the alignment effect dominates.

Organizational capital is inherently embodied in key talent and is distinct from physical assets (Eisfeldt and Papanikolaou, 2013). The preservation and accumulation of this intangible capital require managers to commit to the firm's long-term strategy rather than engaging in myopic behaviors. High compensation, whether in the form of equity or cash, acts as a "bonding" mechanism. Eisfeldt and Papanikolaou (2013) argue that firms with high organizational capital must pay higher compensation to retain key talent and compensate them for the risk associated with organization-specific skills. When CEOs are well-compensated (via high equity stakes or cash salary), they are effectively "bonded" to the firm's long-term success. Consequently, these managers are more likely to leverage the governance role of organizational capital to maintain information transparency and avoid the accumulation of bad news that could destroy the value of the organizational capital they manage.

Therefore, we expect that strong compensation incentives (both equity and cash) will reinforce the governance effect of organizational capital. Specifically, for firms with high CEO compensation, the manager's interest is better aligned with the preservation of organizational capital, leading to a more pronounced reduction in crash risk.

5.3 Research methodology

5.3.1 Sample and data

Our sample consists of U.S. publicly listed firms obtained from the Compustat and CRSP databases over the period 1999–2022. We obtain financial statement data from Compustat, which include selling, general, and administrative (SG&A) expenses used to construct organizational capital. Stock prices and returns used to calculate crash risk measures are obtained

from the CRSP database. CEO characteristics and other governance variables are further collected from the ExecuComp database.

Following prior studies on stock price crash risk (e.g., Kim et al. (2011b,a); Chen et al. (2017)), we exclude financial firms (SIC codes 6000–6999) and utilities (SIC codes 4400–4999) because of their distinct regulatory environment and reporting structure. We also exclude firm-year observations with missing data for our dependent, independent, or control variables. To mitigate the influence of outliers, all continuous variables are winsorized at the 1st and 99th percentiles each year. These sampling procedures yield a final unbalanced panel of 13,489 firm-year observations.

Table 5.1 presents the summary statistics of all variables used in our main analysis. Consistent with prior literature, the sample exhibits substantial variation in both crash risk measures (NCSKEW and DUVOL) and organizational capital (OC/TA), providing sufficient cross-sectional and time-series variation to identify the relation between organizational capital and stock price crash risk.

5.3.2 Organizational Capital

Following Peters and Taylor (2017), we measure organizational capital as the accumulated stock of a proportion of past selling, general, and administrative expenses (SG&A) by employing the perpetual inventory method as shown below:

$$OC_{i,t} = (1 - \delta_0)OC_{i,t-1} + SG\&A_{i,t} \times \beta_0, \tag{5.1}$$

where $OC_{i,t}$ indexes firm-specific stock of organizational capital at time t, δ_0 indexes the depreciation rate of organizational capital, $SG\&A_{i,t}$ proxies for the firms' SG&A expenses at time t, and β_0 indexes the percentage of SG&A expense invested in organizational capital.

SG&A expense is used in the measurement of organizational capital because "a large part of SG&A consists of expenses related to labor and IT (white collar wages, training, consulting, and IT expenses), consistent with the idea that any accrued value will be somewhat firm specific and must be shared with key talent. SG&A contains the part of labor expenses that cannot be directly attributed to a particular unit of output. Hence, any spending on the part of the firm to increase its organizational capital will be included in SG&A expenses" (Eisfeldt and Papanikolaou, 2013). Lev et al. (2009) also argue that SG&A expenses consist of costs linked to developing employee training, information systems, R&D, consultant fees and brand promotion, all of which can contribute to building the firm's organizational capital. Finally, expenditures on worker training, advertising, and related organizational investments are generally recorded within SG&A (Peters and Taylor, 2017).

The following estimates the initial stock of organizational capital:

$$OC_{i,t0} = \frac{SG\&A_{i,t0} \times \beta_0}{g + \delta_0}, \tag{5.2}$$

where g indexes the growth in the flow of organizational capital, calculated as the mean growth of firm-level SG&A expenditure. Following prior literature (Eisfeldt and Papanikolaou, 2013; Peters and Taylor, 2017; Hasan and Uddin, 2020), we use 0.30 of SG&A in estimating the initial stock of organizational capital. Additionally, we use a depreciation rate of 0.20 (Peters and Taylor, 2017). The final organizational capital estimate is scaled by the book value of total assets (OC/TA).

5.3.3 Stock Price Crash Risk

We obtain stock prices and returns from the Centre for Research in Security Prices (CRSP) to construct measures of stock price crash risk. Following Hutton et al. (2009) and Jin and Myers (2006), we use two firm-specific measures of crash risk rather than relying on broad market movements. To derive firm-specific weekly returns, we first estimate the following regression model (Hutton et al., 2009):

$$r_{j,\tau} = \alpha_j + \beta_{1,j} r_{m,\tau-2} + \beta_{2,j} r_{m,\tau-1} + \beta_{3,j} r_{m,\tau} + \beta_{4,j} r_{m,\tau+1} + \beta_{5,j} r_{m,\tau+2} + \epsilon_{j,\tau}, \tag{5.3}$$

where $r_{j,\tau}$ represents the return of firm j in week τ and $r_{m,\tau}$ denotes the market return based on the CRSP value-weighted index in week τ. In addressing non-synchronous trading effects, we adopt the methodology proposed by Dimson (1979) and Chen et al. (2001). This approach involves incorporating lagged and leading terms of the market index into our analysis to capture market information that may be incorporated into stock prices with a delay or lead due to trading frictions. Specifically, the inclusion of leading terms (returns from future weeks) ensures that the estimated residual is orthogonal to market movements even if the firm's price reaction is not perfectly synchronized with the market index.

The firm-specific weekly returns, denoted as $W_{j,\tau}$, are calculated as the natural logarithm of one plus the estimated residual return. This step effectively filters out systematic market-wide variations, isolating the idiosyncratic component of returns unique to the firm. This is expressed in the equation: $W_{j,\tau} = \ln(1 + \epsilon_{j,\tau})$, where $\epsilon_{j,\tau}$ represents the estimated residual return for firm j in week τ.

Our first crash risk measure is the negative conditional skewness of firm-specific weekly returns (NCSKEW), which captures the asymmetry in the return distribution. We follow the methodologies of Chen et al. (2001) and Kim et al. (2011a). The calculation of NCSKEW for each firm-year involves first determining the ratio of the third moment of the firm-specific

weekly return, $W_{j,\tau}$, to the cube of the standard deviation of these returns. This ratio is then multiplied by -1. The formula for NCSKEW is given as:

$$NCSKEW = -[n(n-1)^{3/2} \sum W_{j,\tau}^3]/[(n-1)(n-2)(\sum W_{j,\tau}^2)^{3/2}] \tag{5.4}$$

where $W_{j,\tau}$ has been previously defined as the firm-specific weekly returns, and n represents the total number of observed trading weeks for a given firm in a year. A higher NCSKEW indicates more negative skewness, implying greater crash risk.

The second measure is the down-to-up volatility (DUVOL). For each firm, we divide weeks into "down" weeks (returns below the annual mean) and "up" weeks (returns above the mean). DUVOL is calculated as the logarithm of the ratio of the standard deviation of returns in down weeks to that in up weeks. DUVOL for firm j in week τ is then calculated as the natural logarithm of this ratio:

$$DUVOL_{j,\tau} = \log\left(\frac{(n_u - 1)\sum_{Down} W_{j,\tau}^2}{(n_d - 1)\sum_{Up} W_{j,\tau}^2}\right) \tag{5.5}$$

where n_u and n_d represent the number of up and down weeks each year, respectively. A higher DUVOL suggests a greater likelihood of a crash. Unlike skewness-based measures, DUVOL is less sensitive to extreme return observations.

5.3.4 Control Variable

We control for several variables that prior studies indicate are associated with crash risk (Chen et al., 2017; Chang et al., 2017). To account for the potential serial correlation of NCSKEW or DUVOL, we include the lagged conditional firm-specific weekly return skewness ($NCSKEW$). To capture the intensity of variations in investor opinion that Chen et al. (2001) link to crash risk, we use $DTURN$, defined as the change in average monthly stock turnover. $SIGMA$ proxies for stock price volatility, measured as the standard deviation of firm-specific weekly returns over the fiscal year. We further control for average firm-specific weekly returns over the fiscal year (RET), since high prior returns are more likely to lead to crashes. Operating performance (ROA) is measured as pre-tax income scaled by total assets, following earlier studies which suggest that stronger performance reduces the crash risk. Firm size ($SIZE$), defined as the natural logarithm of market value of equity, is included as larger firms are generally more exposed to crash risk. Growth opportunities are captured by the market-to-book ratio (MTB). Financial leverage (LEV) is added because prior evidence indicates that leverage negatively correlates with crash risk. We also control accounting opacity using discretionary accruals ($ACCM$), measured as the absolute value

of residuals from the modified Jones model, since firms with greater accruals-based earnings management are more prone to crashes. In addition, we include MA to control for Managerial Ability, following the methodology of Demerjian et al. (2012). This measure captures the efficiency of managers in generating revenues from corporate resources using Data Envelopment Analysis (DEA). Finally, we control industry and year fixed effects. Detailed definitions of all variables are provided in the Appendix.

5.3.5 Empirical models

We examine our baseline hypothesis by estimating the following ordinary least squares (OLS) specification:

$$Crash_{i,t+1} = \alpha + \beta_1 OC_{i,t} + \gamma' Controls_{i,t} + YearFE + IndustryFE + \epsilon_{i,t}, \tag{5.6}$$

where $Crash$ represents firm-specific stock price crash risk, measured alternatively by NCSKEW or DUVOL as defined above. The key independent variable is organizational capital (OC/TA), also defined above. Control variables include firm-level characteristics that prior studies have shown to influence crash risk (e.g., size, leverage, profitability, volatility, and accruals). All regressions include year and two-digit SIC industry fixed effects to capture time-varying and industry-specific shocks. Standard errors are clustered at the firm level. If organizational capital intensifies agency conflicts and encourages bad news hoarding, we expect β_1 to be positive and significant. Conversely, if organizational capital constrains opportunism by improving monitoring or operational efficiency, then β_1 may be insignificant or negative.

5.4 Empirical analyses

5.4.1 Descriptive Statistics

Table 5.1 presents the summary statistics for the variables used in the main analysis. The sample comprises 13,489 firm-year observations. The mean values of our two primary crash risk measures, NCSKEW and DUVOL, are 0.205 and 0.133, respectively. These statistics are comparable to those reported in prior studies (e.g., Kim et al. (2011a)), confirming the presence of negative return skewness in our sample. The key independent variable, organizational capital (OC/TA), has a mean of 0.756 with a standard deviation of 0.941, indicating substantial cross-sectional variation in the accumulation of intangible organizational assets

across firms. Regarding control variables, the average firm in the sample has a return on assets (ROA) of 10.6%, a market-to-book ratio (MTB) of 2.971, and a leverage ratio (LEV) of 17.3%. Firm size (SIZE) has a mean of 6.995, reflecting a sample dominated by mid- to large-sized firms. Other variables, including discretionary accruals (ACCM) and stock return volatility (SIGMA), also exhibit wide variation. Overall, the distribution of the variables suggests sufficient heterogeneity to support the empirical analyses.

Table 5.1: **Summary Statistics**

Variable	N	Mean	SD	Min	Q1	Median	Q3	Max
NCSKEW	13489	0.205	0.928	-3.998	-0.298	0.109	0.569	5.179
DUVOL	13489	0.133	0.638	-2.096	-0.274	0.105	0.505	2.611
OC_TA	13489	0.756	0.941	-0.519	0.166	0.468	1.007	6.282
DTURN	13489	0.000	0.046	-0.130	-0.025	0.000	0.025	0.126
RET	13489	31.280	26.950	0.760	11.900	23.380	42.650	185.300
SIGMA	13489	3.846	3.420	0.140	1.594	2.865	4.894	28.100
SIZE	13489	6.995	1.843	1.783	5.643	6.847	8.237	11.790
MTB	13489	2.971	1.873	0.648	1.632	2.462	3.720	11.360
LEV	13489	0.173	0.158	0.000	0.010	0.155	0.287	0.653
ROA	13489	0.106	0.109	-0.949	0.074	0.120	0.163	0.347
ACCM	13489	0.225	0.201	0.024	0.093	0.158	0.279	1.288
MA	13489	-0.002	0.128	-0.271	-0.076	-0.029	0.035	0.648

This table reports the summary statistics for our measures of organizational capital and stock price crash risk variables, as well as the main control variables used in this study. See Appendix for variable definitions.

5.4.2 Regression Analysis

Test of H1 - Organizational Capital and Future Stock Price Crash Risk

Table 5.2 reports the baseline regression results testing the relationship between organizational capital and future stock price crash risk. Column (1) presents the results using negative conditional skewness (NCSKEW) as the dependent variable, while Column (2) uses down-to-up volatility (DUVOL).

Table 5.2: **Baseline Regression: Stock Price and Organizational Capital**

	(1) $\mathbf{NCSKEW}_{t+1}$	(2) $\mathbf{DUVOL}_{t+1}$
OC_TA	-0.042***	-0.029***
	(-3.050)	(-3.094)
NCSKEW	0.001	0.000
	(0.118)	(0.009)
DTURN	0.074	0.056
	(0.399)	(0.452)
RET	0.001	0.000
	(1.246)	(1.227)
SIGMA	0.008**	0.005**
	(2.176)	(2.041)
SIZE	0.023***	0.023***
	(3.882)	(5.639)
MTB	0.036***	0.024***
	(6.968)	(6.779)
LEV	-0.129**	-0.103**
	(-2.011)	(-2.362)
ROA	0.502***	0.348***
	(5.353)	(5.691)
ACCM	0.071	0.034
	(1.477)	(1.048)
MA	-0.257***	-0.176***
	(-3.764)	(-3.834)
Year Fixed Effects	YES	YES
Industry Fixed Effects	YES	YES
Cluster	Firm	Firm
Observations	13,489	13,489
Adjusted R-squared	0.043	0.047

Table 5.2 reports the baseline regression results examining the association between organizational capital and future stock price crash risk. Column (1) presents the results using negative conditional skewness (NCSKEW) as the dependent variable, while Column (2) uses down-to-up volatility (DUVOL). The sample includes 13,489 firm-year observations over 1999–2022. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

Across both specifications, the coefficient on OC/TA is negative and statistically significant at the 1% level. Specifically, the coefficient is -0.042 ($t = -3.050$) for NCSKEW and

-0.029 ($t = -3.094$) for DUVOL. To gauge the economic magnitude of these estimates, we consider the standard deviation of organizational capital in our sample (0.941). A one-standard-deviation increase in organizational capital is associated with approximately 0.040 lower NCSKEW (0.941×0.042), equivalent to about 19.5% of the sample mean of 0.205. This indicates that the negative association is economically meaningful as well as statistically significant.

These findings are consistent with Hypothesis 1a and with the resource-based view (Barney, 1991). One interpretation is that stronger internal communication, monitoring, and operational routines may be associated with less managerial bad-news hoarding, as suggested by prior literature (Jin and Myers, 2006; Hutton et al., 2009). However, the baseline regressions themselves establish an association rather than a causal mechanism.

Regarding the control variables, the results are largely consistent with prior literature. Firms with larger size (SIZE), higher growth opportunities (MTB), and higher profitability (ROA) exhibit greater crash risk in these specifications, while leverage (LEV) is negatively related to crash risk. Because the sign on ROA is positive in both columns, we do not interpret profitability as a mitigating factor in this sample. The coefficient on Managerial Ability (MA) is negative and significant, which is consistent with Demerjian et al. (2012). Overall, the baseline regressions document a statistically and economically meaningful negative association between organizational capital and future stock price crash risk.

Test of H2 - The Role of CEO Generalist Experience

We hypothesize in Hypothesis 2 (H2) that the negative association between organizational capital and stock price crash risk is strengthened when firms are led by Generalist CEOs. Generalist CEOs, characterized by transferable skills across firms and industries, are expected to possess broader cognitive frameworks and superior information-processing capabilities (Custódio et al., 2013). These attributes enable them to better leverage organizational capital, such as internal communication systems and adaptive routines, to enhance transparency and mitigate the accumulation of bad news.

To test this hypothesis, we employ the General Ability Index (GAI) developed by Custódio et al. (2013). The GAI is a composite index based on five aspects of a CEO's professional background: past number of positions, number of firms, number of industries, CEO experience at other firms, and conglomerate experience. We partition our sample based on the annual median of the GAI. Firms with a CEO GAI score above the industry-year median are classified as the High GAI (Generalist) subsample, while those below are classified as the Low GAI (Specialist) subsample.

Table 5.3 presents the regression results for these subsamples. The results provide robust support for our hypothesis. In the High GAI subsample (Columns 1 and 3), the coefficient

on organizational capital (OC_TA) is negative and statistically significant at the 1% level across both crash risk measures. Specifically, for future negative skewness (NCSKEW), the coefficient is -0.042 (t = -2.673), and for down-to-up volatility (DUVOL), the coefficient is -0.029 (t = -2.810). These findings indicate that in firms led by generalist CEOs, a higher stock of organizational capital significantly reduces the likelihood of future stock price crashes.

Table 5.3: **Mechanism of CEO Experience**

	(1) **NCSKEW$_{t+1}$ High GAI**	(2) **NCSKEW$_{t+1}$ Low GAI**	(3) **DUVOL$_{t+1}$ High GAI**	(4) **DUVOL$_{t+1}$ Low GAI**
OC_TA	-0.042***	-0.037	-0.029***	-0.024
	(-2.673)	(-1.344)	(-2.810)	(-1.264)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	9,490	3,996	9,490	3,996
Adjusted R-squared	0.048	0.041	0.051	0.040

Table 5.3 examines the moderating role of CEO generalist experience on the relationship between organizational capital and crash risk. The sample is partitioned into two subsamples based on the General Ability Index (GAI): High GAI (Generalist CEOs) and Low GAI (Specialist CEOs). Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

In sharp contrast, for the Low GAI subsample (Columns 2 and 4), the coefficients on OC_TA are statistically insignificant. The coefficient for NCSKEW is -0.037 (t = -1.344) and for DUVOL is -0.024 (t = -1.264). The lack of significance in the specialist CEO group suggests that organizational capital alone may not be sufficient to curb crash risk without the complementary broad managerial skills of a generalist CEO. This confirms that generalist experience acts as a critical mechanism that amplifies the governance and risk-mitigating benefits of organizational capital.

Test of H3 - The Role of CEO Compensation Mechanism

Hypothesis 3 (H3) posits that CEO compensation incentives reinforce the governance role of organizational capital. High financial incentives align managerial interests with long-term shareholder value, motivating CEOs to utilize organizational capital for efficient information flow rather than for hoarding bad news (Kim et al., 2011b). We examine this mechanism by analyzing two forms of compensation: Equity Incentives (EQ_INC) and Cash Incentives (CASH_INC).

Following Core and Guay (2002) and Kim et al. (2011b), we measure Equity Incentives (EQ_INC) as the natural logarithm of the CEO's portfolio delta. Delta is defined as the dollar change in the value of the CEO's stock and option portfolio for a 1% change in the stock price.

A higher delta implies a stronger alignment between the CEO's wealth and shareholder value. We measure Cash Incentives (CASH_INC) as the natural logarithm of the sum of the CEO's salary and bonus (Core et al., 2003). This captures the short-term performance incentives provided to the CEO.

Table 5.4 reports the regression results partitioned by the median values of equity and cash incentives. Specifically, we classify firms with equity incentives above the sample median into the High Equity Incentive subsample, while those below the median constitute the Low Equity Incentive subsample. We first examine the impact of equity-based compensation. Columns 1, 2, 5, and 6 display the results based on equity incentives. In the High Equity Incentive subsample (Columns 1 and 5), we observe a significant negative relationship between organizational capital and crash risk. The coefficient on OC_TA is -0.044 ($t = -2.057$) for NCSKEW and -0.033 ($t = -2.260$) for DUVOL, both significant at the 5% level. Conversely, in the Low Equity Incentive subsample (Columns 2 and 6), the coefficients on OC_TA are statistically insignificant ($t = -0.311$ for NCSKEW and $t = -0.407$ for DUVOL). This suggests that equity-based alignment is crucial for activating the protective effects of organizational capital.

Table 5.4: **Mechanism of CEO Compensation**

	(1) **NCSKEW$_{t+1}$ High EQ**	(2) **NCSKEW$_{t+1}$ Low EQ**	(3) **NCSKEW$_{t+1}$ High CASH**	(4) **NCSKEW$_{t+1}$ Low CASH**	(5) **DUVOL$_{t+1}$ High EQ**	(6) **DUVOL$_{t+1}$ Low EQ**	(7) **DUVOL$_{t+1}$ High CASH**	(8) **DUVOL$_{t+1}$ Low CASH**
OC_TA	-0.044**	-0.029	-0.045**	-0.059	-0.033**	-0.024	-0.035**	-0.027
	(-2.057)	(-0.311)	(-2.106)	(-0.671)	(-2.260)	(-0.407)	(-2.417)	(-0.476)
Controls	YES	YES	YES	YES	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm	Firm	Firm	Firm	Firm
Observations	7,658	1,119	7,564	1,210	7,658	1,119	7,564	1,210
Adjusted R^2	0.048	0.061	0.049	0.082	0.050	0.053	0.049	0.066

This table investigates whether CEO compensation incentives reinforce the governance role of organizational capital. The regression results are reported for subsamples partitioned by the median values of Equity Incentives (EQ_INC) and Cash Incentives (CASH_INC). Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

We next turn to the analysis of cash-based incentives, reported in Columns 3, 4, 7, and 8. Consistent with the equity results, the negative association between organizational capital and crash risk is concentrated in the High Cash Incentive group. The coefficient on OC_TA is -0.045 (t = -2.106) for NCSKEW and -0.035 (t = -2.417) for DUVOL, both significant at the 5% level. In the Low Cash Incentive group, the coefficients are insignificant (t = -0.671 for NCSKEW and t = -0.476 for DUVOL).

Collectively, these findings are consistent with H3. The cross-sectional pattern suggests that the negative association between organizational capital and future crash risk is stronger when CEO compensation incentives are higher. This evidence is consistent with an incentive-alignment interpretation, but it does not by itself establish the underlying managerial mechanism causally.

5.4.3 Organizational Capital and Stock Price Crash Risk: Agency Theory Perspectives

The Role of External Monitoring Mechanisms

Our analysis begins with institutional investors, who are typically sophisticated market participants with the resources and incentives to monitor management effectively (Callen and Fang, 2013). We measure Institutional Ownership (IO) as the percentage of outstanding shares held by institutional investors. Table 5.5 presents the regression results partitioned by the median level of institutional ownership. In the subsample characterized by High Institutional Ownership (Columns 1 and 3), organizational capital (OC_TA) exhibits a statistically significant negative association with both crash-risk measures. The coefficient is -0.052 (t = -2.528) for NCSKEW and -0.040 (t = -2.842) for DUVOL. In the Low Institutional Ownership group, the NCSKEW coefficient is smaller and statistically insignificant. The table as currently reported, however, also shows a negative and significant DUVOL coefficient in the Low IO subsample. The institutional-monitoring evidence should therefore be interpreted as mixed rather than as uniformly concentrated in the High IO group.

Table 5.5: **Stock Price and Organizational Capital: Institutional Investors' Orientation**

	(1) $\textbf{NCSKEW}_{t+1}$ **High IO**	(2) $\textbf{NCSKEW}_{t+1}$ **Low IO**	(3) $\textbf{DUVOL}_{t+1}$ **High IO**	(4) $\textbf{DUVOL}_{t+1}$ **Low IO**
OC_TA	-0.052**	-0.024	-0.040***	-0.040***
	(-2.528)	(-1.320)	(-2.842)	(-2.842)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	7,031	6,458	7,027	6462
Adjusted R-squared	0.035	0.055	0.037	0.037

This table analyzes the role of external monitoring mechanisms in conditioning the effect of organizational capital. This table reports the regression results for subsamples classified by High versus Low Institutional Ownership (IO). Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

Table 5.6: **Stock Price and Organizational Capital: Analyst Monitoring**

	(1) $\textbf{NCSKEW}_{t+1}$ **High Analyst**	(2) $\textbf{NCSKEW}_{t+1}$ **Low Analyst**	(3) $\textbf{DUVOL}_{t+1}$ **High Analyst**	(4) $\textbf{DUVOL}_{t+1}$ **Low Analyst**
OC_TA	-0.055***	-0.022	-0.038***	-0.015
	(-2.967)	(-1.101)	(-3.067)	(-1.067)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	6,227	7,260	6,227	7,260
Adjusted R-squared	0.060	0.031	0.060	0.036

This table analyses the role of external monitoring mechanisms in conditioning the effect of organizational capital. This table reports the results partitioned by High versus Low Analyst Coverage. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

We extend this analysis by examining financial analyst coverage, which serves as another primary channel for reducing information asymmetry between insiders and outsiders (Yu,

2008). We measure Analyst Coverage as the natural logarithm of one plus the number of analysts issuing earnings forecasts for the firm. Table 5.6 reports the results for firms split by high and low analyst coverage. The empirical patterns mirror those observed for institutional investors. For firms with High Analyst Coverage (Columns 1 and 3), the negative association between organizational capital and crash risk is statistically significant, with a coefficient of -0.055 (t = -2.967) for NCSKEW and -0.038 (t = -3.067) for DUVOL. In contrast, within the Low Analyst Coverage subsample (Columns 2 and 4), the coefficients drop and become statistically indistinguishable from zero. These findings underscore the importance of analysts as information intermediaries who facilitate the transparency benefits inherent in organizational capital.

Taken together, Tables 5.5 and 5.6 are broadly consistent with a complementary role for external monitoring, although the institutional-ownership split is not uniform across both crash-risk measures as currently reported. The analyst-coverage results provide the cleaner evidence that the negative association is concentrated in the more intensively monitored subsample.

The Role of Financial Reporting Transparency

We further explore the mechanism through which organizational capital influences crash risk by considering the firm's information environment. Agency theory posits that opaque financial reporting facilitates the accumulation of bad news by managers, whereas transparent reporting environments naturally curb such opportunistic behavior (Hutton et al., 2009). To test whether the efficacy of organizational capital depends on this environment, we measure Financial Opacity following Hutton et al. (2009), using the absolute value of discretionary accruals estimated from the modified Jones model (Dechow et al., 1995). We partition our sample based on the median of this opacity measure.

Table 5.7 reports the results. We find that the negative relationship between organizational capital and stock price crash risk is statistically significant exclusively in the Low Opacity (i.e., High Transparency) subsample. Specifically, the coefficient on organizational capital for this group is -0.061 (t = -3.408) for NCSKEW and -0.043 (t = -3.459) for DUVOL, both significant at the 1% level. In contrast, for firms characterized by High Opacity (low transparency), the estimated coefficients are insignificant (-0.027 and -0.018). This disparity is consistent with a complementary relationship in which the negative association between organizational capital and crash risk is stronger in more transparent financial reporting environments than in opaque firms.

Table 5.7: **Stock Price and Organizational Capital: Financial Reporting Transparency**

	(1) **NCSKEW$_{t+1}$ High Opacity**	(2) **NCSKEW$_{t+1}$ Low Opacity**	(3) **DUVOL$_{t+1}$ High Opacity**	(4) **DUVOL$_{t+1}$ Low Opacity**
OC_TA	-0.027	-0.061***	-0.018	-0.043***
	(-1.234)	(-3.408)	(-1.263)	(-3.459)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	6,564	6,925	6,563	6,926
Adjusted R-squared	0.046	0.045	0.049	0.047

This table explores the mechanism of financial reporting transparency. The sample is split into High Opacity and Low Opacity subsamples based on the firm's financial reporting quality. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

The Role of Managerial Risk Perception

Finally, we investigate whether the risk-mitigating effect of organizational capital is conditional upon the manager's perception of risk. We posit that managers who are more cognizant of potential risks are better positioned to utilize the firm's organizational capabilities to identify and neutralize threats before they precipitate a stock price crash. For this analysis, we construct a Managerial Risk Perception index within our sample using firm characteristics intended to capture a more risk-averse managerial posture, including lower R&D intensity, lower leverage, and higher cash holdings. We then divide the sample into High-Risk Perception and Low-Risk Perception groups based on this index. Because this measure is constructed for the present analysis, its results should be interpreted as a cross-sectional conditioning test rather than as a validated standalone measure of managerial risk perception.

Table 5.8 presents the regression results for this analysis. In the subsample of managers with High-Risk Perception (Columns 1 and 3), we observe a significant negative association between organizational capital (OC_TA) and crash risk. The coefficient on OC_TA is -0.043 (t = -2.544) for NCSKEW and -0.028 (t = -2.400) for DUVOL, both significant at the 5% level. These results indicate that the negative association between organizational capital and future crash risk is statistically detectable in the High-Risk Perception subsample. Conversely, in the Low-Risk Perception subsample (Columns 2 and 4), the coefficients on OC_TA are statistically insignificant (-0.042 and -0.020). The difference across subsamples is consistent with the possibility that the relationship varies with managerial risk orienta-

tion, although this split-sample evidence should not be interpreted as establishing a causal mechanism.

Table 5.8: **Stock Price and Organizational Capital: Managerial Risk Perception**

	(1) **NCSKEW$_{t+1}$ High Risk Perception**	(2) **NCSKEW$_{t+1}$ Low Risk Perception**	(3) **DUVOL$_{t+1}$ High Risk Perception**	(4) **DUVOL$_{t+1}$ Low Risk Perception**
OC_TA	-0.043**	-0.042	-0.028**	-0.020
	(-2.544)	(-1.515)	(-2.400)	(-1.331)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	6,623	6,866	6,620	6,869
Adjusted R-squared	0.040	0.047	0.044	0.048

This table presents the regression results conditional on managerial risk perception. The analysis partitions the sample into High-Risk Perception and Low Risk Perception groups. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

5.4.4 Exclusion of computers, software, and electronic equipment firms (business equipment industry)

Following Hasan and Cheung (2018b), we note that 17.1% of our sample belongs to the business equipment industry (computers, software, and electronic equipment firms). Technology firms usually start business with less physical capital and incur more expenses on intangibles (e.g., R&D). One may argue that our documented association is driven mainly by the high-tech nature of the business equipment industry. To alleviate this concern, we re-run the regressions after excluding the business equipment industry from the sample.

Table 5.9 reports the regression results for this reduced sample. As shown in Columns (1) and (2), the coefficients on OC_TA remain negative and statistically significant at the 1% level (-0.0843 and -0.0579, respectively). These results indicate that the negative relationship between organizational capital and stock price crash risk is robust and not solely driven by the specific characteristics of high-technology firms.

Table 5.9: **Exclusion of computers, software, and electronic equipment firms**

	(1)	(2)
	NCSKEW$_{t+1}$	**DUVOL**$_{t+1}$
OC_TA	-0.0843***	-0.0579***
	(-3.250)	(-3.190)
Controls	YES	YES
Year Fixed Effects	YES	YES
Industry Fixed Effects	YES	YES
Cluster	Firm	Firm
Observations	11,168	11,182
Adjusted R^2	0.041	0.043

This table reports the regression results after excluding firms in the business equipment industry (computers, software, and electronic equipment). This subsample analysis ensures our findings are not driven by high-tech firms. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

5.4.5 Alternative Measures of Organizational Capital

To ensure the reliability and validity of our baseline findings, we conduct robustness tests using alternative specifications for our key independent variable. In our main analysis, we scaled organizational capital by total assets. However, to rule out the possibility that our

results are driven by the specific scaling method or functional form, we re-estimate the baseline regression model using two alternative measures: the natural logarithm of organizational capital (OC_ln) and organizational capital scaled by the market value of equity (OC_MVE).

Table 5.10 reports the regression results for these alternative measures. Columns 1 and 2 present the results using the natural logarithm of organizational capital. The coefficient on OC_ln is negative and statistically significant at the 1% level across both crash risk measures. Specifically, when the dependent variable is $NCSKEW_{t+1}$, the coefficient is -0.044 with a t-statistic of -3.434. Similarly, when the dependent variable is $DUVOL_{t+1}$, the coefficient is -0.028 with a t-statistic of -3.171. These results indicate that the inverse relationship between organizational capital and crash risk remains robust when controlling for the skewness of the variable distribution via logarithmic transformation.

Columns 3 and 4 of Table 5.10 present the results using organizational capital scaled by the market value of equity. This alternative scaling captures the relative importance of organizational capital to the firm's market valuation rather than its book assets. Consistent with our primary findings, the coefficient on OC_MVE is negative and highly significant at the 1% level. For $NCSKEW_{t+1}$, the coefficient is -0.062 with a t-statistic of -6.086, and for $DUVOL_{t+1}$, the coefficient is -0.042 with a t-statistic of -5.758. These findings provide strong confirmatory evidence that the negative association between organizational capital and stock price crash risk is not sensitive to the specific measurement or scaling choice of the independent variable.

Table 5.10: **Alternative Measure: Organizational Capital**

	(1) $\mathbf{NCSKEW}_{t+1}$	(2) $\mathbf{DUVOL}_{t+1}$	(3) $\mathbf{NCSKEW}_{t+1}$	(4) $\mathbf{DUVOL}_{t+1}$
OC_ln	-0.044***	-0.028***		
	(-3.434)	(-3.171)		
OC_MVE			-0.062***	-0.042***
			(-6.086)	(-5.758)
Controls	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Industry Fixed Effects	YES	YES	YES	YES
Cluster	Firm	Firm	Firm	Firm
Observations	11,873	11,873	12,523	12,523
Adjusted R-squared	0.045	0.051	0.047	0.048

This table presents robustness tests using alternative measures of organizational capital. We re-estimate the baseline model using the natural logarithm of organizational capital (OC_ln) in Columns (1) and (2), and organizational capital scaled by the market value of equity (OC_MVE) in Columns (3) and (4). The estimation sample varies across specifications because of data availability. Industry fixed effects are based on the two-digit SIC codes. Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of other variables.

5.4.6 An IV approach with heteroskedasticity-based instruments

A critical challenge in interpreting the relationship between organizational capital and stock price crash risk causally is endogeneity, arising from unobservable firm-specific characteristics and potential reverse causality. Finding credible external instrumental variables (IVs) for organizational capital is particularly difficult because macroeconomic or industry-level shocks that affect a firm's SG&A investment may also influence stock price crash risk through other channels.

As a complementary endogeneity check, we employ the heteroskedasticity-based IV approach developed by Lewbel (2012). The method uses heteroskedasticity in the first-stage residuals to generate additional instruments from mean-centered exogenous covariates. The presence of heteroskedasticity supports the relevance condition required for constructing these instruments. However, the exclusion restriction remains an identifying assumption rather than an empirically established condition. We therefore interpret the Lewbel estimates as complementary robustness evidence rather than as definitive evidence of causality. To assess sensitivity to possible weak-instrument concerns, we also report the identification-robust confidence-set approach of Andrews et al. (2018), which is designed to limit coverage distortions when instruments are not strongly identified (Stock and Yogo, 2005).

Table 5.11 presents the 2SLS regression results. The Hansen J-statistics reported at the bottom of the table provide an important diagnostic. The p-values of 0.0147 and 0.0049 imply rejection of the overidentifying restrictions at conventional significance levels, which raises concerns about whether the generated instruments are strictly exogenous. The estimated coefficients on OC_TA nevertheless remain negative and statistically significant in both specifications. Given the Hansen J results, these estimates should be viewed as supporting the robustness of the negative association rather than as establishing a causal effect.

More broadly, the empirical analysis in this chapter is fundamentally observational and reduced-form in nature. The Lewbel IV estimates, Heckman selection model, Propensity Score Matching, entropy balancing, and other robustness checks strengthen confidence that the baseline association is not driven by a single specification or an obvious source of selection. However, these procedures cannot fully rule out all forms of endogeneity or establish definitive causality. Accordingly, the evidence in this chapter should be interpreted as robust evidence of a negative association between organizational capital and future stock price crash risk, rather than as proof of a causal relationship.

Table 5.11: **Endogeneity Test: Lewbel (2012)**

	(1)	(2)
	NCSKEW$_{t+1}$	**DUVOL**$_{t+1}$
OC_TA	-0.041***	-0.028***
	(-3.06)	(-3.10)
Controls	YES	YES
Year Fixed Effects	YES	YES
Industry Fixed Effects	YES	YES
Cluster	Firm	Firm
Observations	13,490	13,490
Hansen J p-value	0.0147	0.0049
Hansen J Stat	17.46**	18.004***

This table presents the regression results using the Lewbel (2012) heteroscedasticity-based instrumental variable approach to address endogeneity. The sample includes 13,490 firm-year observations. The dependent variables are one-year-ahead negative conditional skewness (NCSKEW$_{t+1}$) in Column (1) and down-to-up volatility (DUVOL$_{t+1}$) in Column (2). Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of variables.

5.4.7 Heckman selection IV model

To examine potential endogeneity arising from sample self-selection, we employ a Heckman two-stage selection model (Heckman, 1979). In the selection equation, we use the industry-average organizational capital of other firms in the same industry-year (*Ind_Avg_OC*) as an exclusion variable. This specification relies on two conditions: the variable must be relevant for a firm's organizational-capital status, and its effect on future crash risk should operate primarily through the focal firm's organizational capital rather than through an independent direct channel.

The relevance argument is that firms' investments in organizational capabilities are influenced by common industry-level structural trends and competitive conditions, so the organizational-capital accumulation of peer firms is informative about the focal firm's propensity to maintain high organizational capital. The proposed exclusion restriction is motivated by the firm-specific bad-news-hoarding mechanism underlying crash risk (Jin and Myers, 2006; Hutton et al., 2009): peer firms' organizational capital should not directly determine the focal firm's internal information environment or managerial concealment incentives. This argument provides the economic rationale for using *Ind_Avg_OC* as an exclusion variable, although the exclusion restriction remains an identifying assumption rather than a directly testable fact.

In the first stage, we estimate a probit model to predict the probability of a firm being a "high organizational capital" firm (*High_OC*), defined as an organizational capital ratio above the industry-year median. As reported in Part A of Table 5.12, the coefficient on *Ind_Avg_OC* is highly significant (coefficient = -0.361, t = -6.24), confirming the strong relevance of our instrument.

In the second stage (Part B of Table 5.12), we incorporate the Inverse Mills Ratio (IMR), calculated from the first-stage model, into the baseline regressions to account for potential selection. The coefficients on the IMR are statistically insignificant for both the $NCSKEW_{t+1}$ (t = -0.37) and $DUVOL_{t+1}$ (t = -0.12) specifications. The coefficients on organizational capital remain negative and statistically significant at the 1% level for both $NCSKEW_{t+1}$ (coefficient = -0.040, t = -2.78) and $DUVOL_{t+1}$ (coefficient = -0.028, t = -2.85). These results provide additional evidence that the baseline negative association is not solely attributable to the selection mechanism captured by the Heckman specification, while not eliminating other possible sources of endogeneity.

Table 5.12: **Heckman Two-Stage Results**

	Part A: First-Stage	**Part B: Second-Stage**	
	(1)	(2)	(3)
Variables	**High_OC**	**NCSKEW$_{t+1}$**	**DUVOL$_{t+1}$**
OC_TA		-0.040***	-0.028***
		(-2.78)	(-2.85)
Ind_Avg_OC	-0.361***		
	(-6.24)		
IMR		-0.047	-0.010
		(-0.37)	(-0.12)
Controls	YES	YES	YES
Year Fixed Effects	YES	YES	YES
Industry Fixed Effects	YES	YES	YES
Cluster	Firm	Firm	Firm
Observations	18,681	13,378	13,378
Pseudo R2	0.083		
R2		0.047	0.051
Adj. R2		0.042	0.046

This table presents the estimation results of the Heckman two-stage selection model used to address potential sample self-selection bias. Part A (Column 1) reports the first-stage probit regression predicting the likelihood of a firm having high organizational capital (*High_OC*), defined as an organizational capital ratio above the industry-year median. *Ind_Avg_OC*, the industry-average organizational capital excluding the focal firm, serves as the instrumental variable (exclusion restriction). Part B (Columns 2 and 3) reports the second-stage ordinary least squares (OLS) regression results for future stock price crash risk (NCSKEW$_{t+1}$ and DUVOL$_{t+1}$), incorporating the Inverse Mills Ratio (IMR) calculated from the first stage. All continuous variables are winsorized at the 1st and 99th percentiles. Robust t-statistics clustered at the firm level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

5.4.8 Propensity Score Matching (PSM)

To address the concern of potential self-selection bias—where firms with specific observable characteristics might be more likely to accumulate high organizational capital—we employ a Propensity Score Matching (PSM) approach. This method improves the comparability of high- and low-organizational-capital firms on observed characteristics and allows us to assess whether the baseline association persists in a more balanced matched sample.

We first classify firms into a treatment group (High OC) and a control group based on whether their organizational capital (OC/TA) is above the sample median. We then estimate propensity scores using a logit model that includes all the control variables from our baseline

regression, along with industry and year fixed effects. Subsequently, we perform a one-to-one nearest neighbour matching process to construct a matched sample that balances the covariates between the two groups.

Table 5.13 reports the regression results using this matched sample of 8,177 observations. Consistent with our main findings, the coefficient on OC/TA remains negative and statistically significant at the 1% level for both crash risk measures. Specifically, the coefficient is -0.043 ($t = -2.72$) when the dependent variable is NCSKEW_{t+1}, and -0.031 ($t = -2.94$) when the dependent variable is DUVOL_{t+1}. These results suggest that the negative association between organizational capital and stock price crash risk remains robust after matching on observed firm characteristics. PSM cannot, however, eliminate bias arising from unobserved differences between the matched groups.

Table 5.13: **Robustness Check: Propensity Score Matching (PSM)**

	(1)	(2)
	NCSKEW_{t+1}	**DUVOL_{t+1}**
OC_TA	-0.043***	-0.031***
	(-2.72)	(-2.94)
Controls	YES	YES
Year Fixed Effects	YES	YES
Industry Fixed Effects	YES	YES
Cluster	Firm	Firm
Observations	8,177	8,177
Adj. R^2	0.049	0.055

This table presents the regression results using a propensity score matched sample to address potential endogeneity arising from self-selection bias. Firms are matched based on the probability of having high organizational capital using one-to-one nearest neighbor matching. The sample includes 8,177 matched firm-year observations. The dependent variables are one-year-ahead negative conditional skewness (NCSKEW_{t+1}) in Column (1) and down-to-up volatility (DUVOL_{t+1}) in Column (2). Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of variables.

5.4.9 Entropy Balancing

To further mitigate concerns regarding selection bias and model dependence that may not be fully addressed by traditional matching methods (which often result in sample size reduction), we employ the Entropy Balancing method proposed by Hainmueller (2012). Unlike standard matching techniques such as PSM that may discard unmatched observations, Entropy Balancing achieves covariate balance by reweighting the control group data to match the moments of the treatment group covariates exactly, while retaining the full information of the sample.

We define firms with organizational capital above the sample median as the treatment group and reweight the control group to match the first moment (mean) of all control variables utilized in our baseline model. We then re-estimate the baseline regression using these calculated weights.

Table 5.14 presents the results from the entropy-balanced sample. Consistent with our previous findings, the coefficient on organizational capital (OC/TA) remains negative and statistically significant at the 1% level for both crash risk measures. Specifically, in Column 1 (NCSKEW), the coefficient on OC/TA is -0.038 ($t = -2.65$). Similarly, in Column 2 (DUVOL), the coefficient is -0.025 ($t = -2.64$). These results show that the negative association between organizational capital and crash risk remains after reweighting the sample to balance observed firm characteristics. As with PSM, entropy balancing addresses imbalance in observed covariates but does not rule out unobserved confounding.

Table 5.14: **Robustness Check: Entropy Balancing**

	(1)	(2)
	NCSKEW$_{t+1}$	**DUVOL**$_{t+1}$
OC_TA	-0.038***	-0.025***
	(-2.65)	(-2.64)
Controls	YES	YES
Year Fixed Effects	YES	YES
Industry Fixed Effects	YES	YES
Cluster	Firm	Firm
Observations	13,490	13,490
Adj. R^2	0.051	0.055

This table reports the regression results using Entropy Balancing to address potential selection bias. The control group is reweighted to match the mean of covariates in the treatment group (High OC). The sample includes 13,490 firm-year observations. The dependent variables are one-year-ahead negative conditional skewness (NCSKEW$_{t+1}$) in Column (1) and down-to-up volatility (DUVOL$_{t+1}$) in Column (2). Robust t-statistics corrected for clustering at the firm level are presented in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels respectively. See the Appendix for more details about the definition of variables.

5.4.10 Evidence from the machine learning approach

In this section, we complement the econometric analysis with a machine learning framework to examine the predictive importance of organizational capital in a nonlinear setting. Traditional Ordinary Least Squares (OLS) regressions are highly interpretable but may not fully capture nonlinearities and higher-order interactions among covariates. Our theoretical framework suggests that the association between organizational capital and crash risk may depend on other governance structures, such as CEO generalist skills and compensation schemes, and may therefore be nonlinear.

To capture this complex interaction structure and examine out-of-sample predictive performance, we employ the eXtreme Gradient Boosting (XGBoost) model (Chen and Guestrin, 2016). XGBoost is a tree-based ensemble machine learning architecture. Unlike linear models, decision trees recursively partition the state space via binary splits, allowing the algorithm to naturally map nonlinearities, threshold constraints, and multi-dimensional interactions without requiring the researcher to pre-specify complex interaction terms (Jones et al., 2017).

By applying XGBoost, we evaluate the predictive importance of organizational capital for future stock price crash risk. The model's predictive structure is summarized using the Relative Variable Importance (RVI) metric. RVI measures how much a predictor contributes to

improvements in the fitted tree ensemble across sequential boosting steps. Because variable importance is a predictive rather than causal quantity, it should not be interpreted as identifying a structural effect. The metric is normalized on a scale from 0 to 100, where the most influential predictor scores 100.

As detailed in Appendix B, OC/TA ranks as the third most important predictor of future crash risk, with an RVI score of 65.2. Its relatively high ranking in a model capable of capturing nonlinear interactions indicates that organizational capital contains predictive information beyond a simple linear specification. This result complements the regression evidence by showing that OC remains relevant in the predictive model. It does not, however, identify a causal mechanism or demonstrate that the observed relationships arise from a particular structural channel.

5.5 Conclusion

This study provides comprehensive empirical evidence on the relationship between organizational capital and stock price crash risk in the U.S. market. Contrary to the concern that intangible assets exacerbate opacity, we document that higher organizational capital is associated with lower future stock price crash risk. The cross-sectional evidence is consistent with the view that operational efficiency and information transparency are potential channels through which OC may be related to crash risk.

Our analysis further shows that the negative association between Organizational Capital and crash risk is not uniform across firms. It is stronger in several subsamples characterized by Generalist CEOs, higher equity or cash compensation incentives, greater financial transparency, and stronger external monitoring. These cross-sectional patterns are consistent with the proposed governance channels, while remaining observational rather than causal evidence of those mechanisms.

These findings have important implications. For investors, our results suggest that organizational capital should be viewed as a signal of resilience rather than a source of agency risk. For corporate boards, the study highlights the importance of matching high-OC environments with capable generalist leadership and aligned incentive structures to maximize shareholder value protection. Future research may further explore how specific components of organizational capital interact with emerging risks in the digital economy.

5.6 Appendix

5.6.1 Appendix A. Variable definitions

Variable	Definition
Dependent Variable	
NCSKEW	A measure of stock price crash risk that captures negative conditional skewness of firm-specific weekly returns over the year.
DUVOL	Down-to-up volatility measure of crash risk, which is calculated as the natural logarithm of the ratio of the standard deviation of firm-specific weekly returns in the "down" weeks to the standard deviation in the "up" weeks.
Independent Variable	
OC_TA	The stock of organizational capital estimated using the Perpetual Inventory Method (Peters and Taylor, 2017), scaled by total assets.
OC_Ln	The natural logarithm of the organizational capital stock.
OC_MVE	The stock of organizational capital scaled by the market value of equity.
Control Variable	
DTURN	Change in average monthly stock turnover.
RET	The mean of firm-specific weekly returns over the fiscal year.
SIGMA	The standard deviation of firm-specific weekly returns over the fiscal year.
SIZE	The natural logarithm of the firm's total assets.
MTB	The market-to-book ratio, calculated as the market value of equity divided by the book value of equity.
LEV	Financial leverage, calculated as total debt divided by total assets.
ROA	Return on assets, calculated as income before extraordinary items divided by lagged total assets.
ACCM	The absolute value of discretionary accruals estimated from the modified Jones model (Dechow et al., 1995).
MA	Managerial Ability Score, measured using the Data Envelopment Analysis (DEA) method following Demerjian et al. (2012).
Mechanism Variables	
GAI	General Ability Index, a composite measure of a CEO's generalist experience based on the past number of positions, number of firms, number of industries, CEO experience at other firms, and conglomerate experience (Custódio et al., 2013).
EQ_INC	CEO equity incentives, measured as the natural logarithm of the CEO's portfolio delta, defined as the dollar change in the value of the CEO's stock and option portfolio for a 1% change in the stock price (Core and Guay, 2002; Kim et al., 2011b).
CASH_INC	CEO cash incentives, measured as the natural logarithm of the sum of the CEO's salary and bonus (Core et al., 2003).

Variable	Definition
IO	Institutional ownership, measured as the percentage of outstanding shares held by institutional investors.
Analyst Coverage	The natural logarithm of one plus the number of analysts issuing earnings forecasts for the firm.
Financial Opacity	The absolute value of discretionary accruals estimated from the modified Jones model, used as a proxy for financial reporting opacity (Hutton et al., 2009; Dechow et al., 1995).
Risk Perception	Managerial risk perception index constructed in this study from firm characteristics intended to capture a more risk-averse managerial posture, including lower R&D intensity, lower leverage, and higher cash holdings.
Instrumental Variables	
Ind_Avg_OC	The industry-average organizational capital excluding the focal firm in a given year, used as the instrumental variable in the Heckman first-stage model.
High_OC	An indicator variable that equals 1 if the firm's organizational capital ratio (OC_TA) is above the industry-year median, and 0 otherwise.
IMR	Inverse Mills Ratio, calculated from the first-stage probit regression of the Heckman selection model to control for self-selection bias.

5.6.2 Appendix B: Machine Learning XGBoost

This supplement details the outcomes from utilizing the extreme gradient boosting machine learning framework. To align with the analytical approach of our prior sections, we integrate 10 predictive variables. Research has shown that sophisticated machine learning techniques, notably ensemble methods such as boosting, can perform robustly in complex prediction settings (Barboza et al., 2017; Jones et al., 2017).

Developed by Chen and Guestrin (2016), the extreme gradient boosting model represents an advanced, efficient version of gradient boosting, a method initially introduced by Friedman (2001). This technique retains the core methodology of traditional gradient boosting but introduces several optimizations and enhancements for improved performance and scalability (Chen and Guestrin, 2016). As described by Friedman (2001), the gradient boosting model performs the following sequences:

Gradient Boosting Algorithm. 1. Initialize $f_0(x) = \text{argmin}_\gamma \sum_{i=1}^{N} L(y_i, \gamma)$. 2. For $m = 1$ to M: (a) For $i = 1, 2, \ldots, N$ compute $r_{im} = -\left[\frac{\partial L(y_i, f(x_i))}{\partial f(x_i)}\right]_{f=f_{m-1}}$. (b) Fit a regression tree to the targets r_{im} giving terminal regions $R_{jm}, j = 1, 2, \ldots, J_m$. (c) For $j = 1, 2, \ldots, J_m$ compute

$$\gamma_{jm} = \text{argmin}_\gamma \sum_{x_i \in R_{jm}} L(y_i, f_{m-1}(x_i) + \gamma)$$

. (d) Update $f_m(x) = f_{m-1}(x) + \sum_{j=1}^{J_m} \gamma_{jm} I(x \subset R_{jm})$. 3. Output $f(x) = f_M(x)$.

where L is the loss function, y_i is the outcome variable, x is a set of explanatory variables, $m - 1$ is the current model, M represents the number of iterations, and $J_m, m = 1, 2, \ldots, M$ is the size of each of the constituent trees. Intuitively, the model starts from a weak learner's perspective, which improves through an iterative function of gradient descent algorithms.

In Appendix E.1 (below), we report the relative variable importance (RVI) rankings, which indicate the overall predictive power of individual variables on the future stock price crash risk. The ranks are relative to each predictor's contribution over the tree. The variable with the highest sum of improvements scores 100, while the remaining variables have lower scores descending toward 0. Values closer to 0 indicate less importance in the estimation compared with the top-performing variable (see Jones 2017).

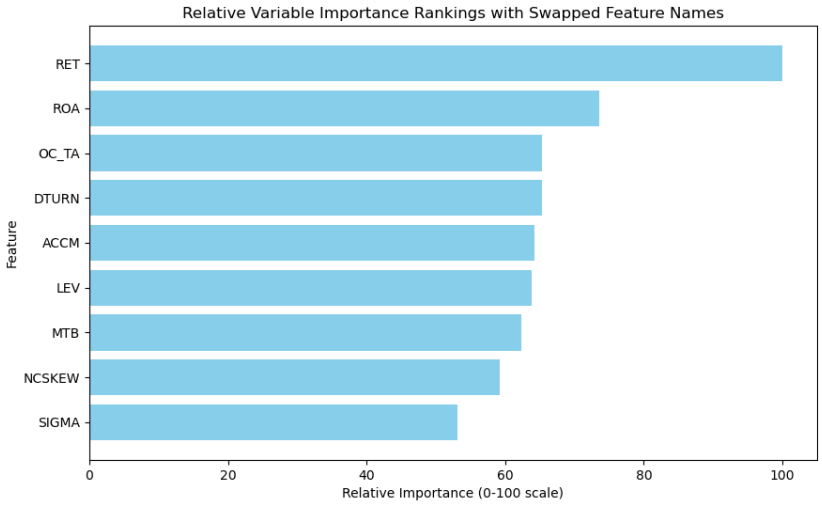


Figure 5.1: Relative Variable Importance Rankings

Chapter 6

Conclusion

This doctoral dissertation has investigated the pricing, management, and mitigation of risk in modern financial markets, traversing the domains of natural resource economics, energy transition finance, and corporate governance. As the global economy enters an era defined by "polycrisis"—characterized by accelerating climate change, structural policy shifts, and complex agency conflicts—traditional valuation frameworks predicated on assumptions of linearity and low-dimensional state spaces have proven increasingly insufficient. Through four self-contained but thematically integrated essays, this research has developed a unified computational and empirical perspective showing how domain-specific physical and organizational information can be integrated with modern numerical and machine-learning methods to study valuation and risk-management problems.

A primary contribution of this thesis lies in addressing the computational challenges associated with high-dimensional and nonlinear valuation problems. In the domain of natural resource management, the first essay addressed the "curse of dimensionality" inherent in aquaculture harvesting decisions. By applying Deep Q-Learning (DQL), the study demonstrated that Deep Reinforcement Learning (DRL) offers a scalable, nonparametric alternative to traditional regression-based methods like Least-Squares Monte Carlo (LSM). The results show that DQL can learn an economically coherent harvesting policy with average stopping times broadly comparable to the LSMC benchmark, although its option-value estimates are lower in the reported calibration. The principal contribution is therefore methodological scalability: the neural-network framework provides a flexible route for extending the optimal-stopping problem to richer state spaces without manually expanding a polynomial basis.

Building on the need for robust valuation under structural uncertainty, the third essay extended these computational innovations to the energy sector. By comparing several function-approximation techniques within a physics-informed real-options framework, the analysis showed that numerical methods can differ in their treatment of the non-smooth climate threshold. Tree-based ensemble methods, specifically Random Forest and XGBoost, produced distinct valuation and timing implications near the non-smooth climate threshold relative to

global polynomial approximators. The comparison highlights that methods based on recursive partitioning are naturally suited to representing hard regime boundaries and that numerical-method choice can matter when the value surface is non-smooth.

Beyond methodological advancements, this dissertation underscores the critical importance of integrating physical and biological realities into financial modeling. The second essay bridged biological modeling and financial engineering to value biomass futures contracts, revealing that the economic benefit of hedging is highly state-dependent. The findings showed that hedging instruments are most valuable under scenarios of "extreme" environmental volatility, fundamentally altering optimal harvesting strategies to mitigate biological downside risk. Similarly, the third essay interpreted the additional value associated with anticipation of a high-carbon-price crisis regime as a “fear premium”, highlighting how the possibility of a physical tipping point can affect the option value of carbon-free baseload generation before the threshold is crossed. Together, these chapters demonstrate that financial models must explicitly account for state-dependent physical regimes to accurately price transition risks.

Complementing the analysis of external physical risks, the fourth essay examined Organizational Capital (OC) as a governance-related intangible asset. Using a comprehensive panel of U.S. firms, the study documented a significant negative association between OC and future stock price crash risk. The relation is conditional rather than uniform across firms and is stronger in several settings characterized by external monitoring, generalist CEOs, and stronger compensation incentives. These patterns challenge the view of intangibles as solely a source of opacity and are consistent with OC playing a governance-related role when complemented by capable leadership and monitoring, while the empirical evidence remains observational rather than causal.

The implications of this research extend to practitioners, policymakers, and future academic inquiry. The findings suggest that financial institutions and energy firms must accelerate the adoption of "physics-informed" valuation models that look beyond historical financial data to incorporate biological and climatic state variables. Furthermore, the demonstrated superiority of ensemble methods in handling regime-switching dynamics suggests a need to revisit standard real options valuation methodologies in the presence of structural breaks. Future research should further bridge the gap between physical science and finance, potentially by integrating richer physical-state dynamics directly into reinforcement-learning frameworks and thereby capturing feedback between economic decisions and the climate system with greater fidelity. Ultimately, this dissertation concludes that in a nonlinear world, the successful pricing of risk requires a synthesis of advanced computation, physical science, and organizational theory.

Bibliography

K.K. Aase. A pricing model for quantity contracts. *Journal of Risk and Insurance*, 71(4): 617–642, 2004.

D. Aboody and B. Lev. Information asymmetry, r&d, and insider gains. *The Journal of Finance*, 55(6):2747–2766, 2000.

A. Ajdari, P. Noursalehi, and A. Sobhani. To wait or not to wait: Optimal stopping for radiotherapy. *Working Paper*, 2019.

Leif Andersen and Mark Broadie. Primal-dual simulation algorithm for pricing multidimensional american options. *Management Science*, 50(9):1222–1234, 2004.

R. C. Anderson, S. A. Mansi, and D. M. Reeb. Board structure, accounting profitability, and patent applications. *Journal of Law and Economics*, 52(4):619–645, 2009.

P. C. Andreou, C. Antoniou, J. Horton, and C. Louca. Corporate governance and firm-specific stock price crashes. *European Financial Management*, 22(5):916–956, 2016.

Isaiah Andrews, James H. Stock, and Liyang Sun. Identification-robust confidence sets for moments of policy effects. *Econometrica*, 86(6):1925–1963, 2018.

R. Arnason. Optimal feeding schedules and harvesting time in aquaculture. *Marine Resource Economics*, 7(1):15–35, 1992.

F. Asche and T. Bjorndal. *The economics of salmon aquaculture*. John Wiley & Sons, 2011.

Andrew Atkeson and Patrick J Kehoe. Modeling and measuring organization capital. *Journal of Political Economy*, 113(5):1026–1053, 2005.

N. Attig and S. Cleary. Organizational capital and investment-cash flow sensitivity: the effect of management quality practices. *Financial Management*, 43(3):473–504, 2014.

N. Attig and S. El Ghoul. Organization capital and the cost of equity financing in medium-sized manufacturing firms. *Contemporary Accounting Research*, 35(3):1616–1644, 2018.

R. Ball. Market and political/regulatory perspectives on the recent accounting scandals. *Journal of Accounting Research*, 47(2):277–323, 2009.

Flavio Barboza, Herbert Kimura, and Edward Altman. Machine learning models and bankruptcy prediction. *Expert Systems with Applications*, 83:405–417, 2017.

J. Barney. Firm resources and sustained competitive advantage. *Journal of Management*, 17 (1):99–120, 1991.

Robert J Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.

Mary E Barth, Ron Kasznik, and Maureen F McNichols. Intangible assets, it and insider trading. *Journal of Accounting and Economics*, 31(1):1–26, 2001.

A. Bauer, B. Henderson, and M. Welker. Organization capital and the market for corporate control. *Working Paper*, 2025.

S. Becker, P. Cheridito, and A. Jentzen. Deep optimal stopping. *Journal of Machine Learning Research*, 20(74):1–25, 2019.

Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Pricing and hedging american-style options with deep learning. *Journal of Risk and Financial Management*, 13(7):158, 2020.

E. Benmelech, E. Kandel, and P. Veronesi. Stock-based compensation and ceo (dis) incentives. *The Quarterly Journal of Economics*, 125(4):1769–1820, 2010.

D. Bergstresser and T. Philippon. Ceo incentives and earnings management. *Journal of Financial Economics*, 80(3):511–529, 2006.

Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, Belmont, MA, 3rd edition, 2005.

Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.

A. Betzer, H. S. G. Lee, P. Limbach, and J. M. Salas. Are generalists beneficial to corporate shareholders? evidence from exogenous executive turnovers. *Journal of Financial and Quantitative Analysis*, 55(2):581–619, 2020.

Petter Bjerksund. Contingent claims evaluation when the convenience yield is stochastic: Analytical results. Working paper, Norwegian School of Economics and Business Administration, 1991.

T. Bjorndal. Optimal harvesting of farmed fish. *Marine Resource Economics*, 5:139–159, 1988.

S. E. Black and L. M. Lynch. Measuring organizational capital in the new economy. In *Measuring capital in the new economy*, pages 205–236. University of Chicago Press, 2005.

A. Bleck and X. Liu. Market transparency and the accounting regime. *Journal of Accounting Research*, 45(2):229–256, 2007.

N. Bloom, R. Sadun, and J. Van Reenen. The organization of firms across countries. *The Quarterly Journal of Economics*, 127(4):1663–1705, 2012.

Oliver Boguth, David Newton, and Mikhail Simutin. The fragility of organization capital. *Journal of Financial and Quantitative Analysis*, 57(3):857–887, 2022. doi: 10.1017/S0022109021000144.

Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001.

J.R. Brett. Temperature and fish. *Chesapeake Science*, 10:275–276, 1969.

J.R. Brett. Environmental factors and growth. In W.S. Hoar, D.J. Randall, and J.R. Brett, editors, *Fish Physiology*, volume 8, pages 599–675. Academic Press, 1979.

N. Burns and S. Kedia. The impact of performance-based compensation on misreporting. *Journal of Financial Economics*, 79(1):35–67, 2006.

J. L. Callen and X. Fang. Institutional investor stability and crash risk: Monitoring versus short-termism? *Journal of Banking & Finance*, 37(8):3047–3063, 2013.

J. L. Callen and X. Fang. Religion and stock price crash risk. *Journal of Financial and Quantitative Analysis*, 50(1-2):169–195, 2015.

B. I. Carlin, B. Chowdhry, and M. J. Garmaise. Investment in organization capital. *Journal of Financial Intermediation*, 21(2):268–286, 2012.

E. Caroli and J. Van Reenen. Skill-biased organizational change? evidence from a panel of british and french establishments. *The Quarterly Journal of Economics*, 116(4):1449–1492, 2001.

Konan Chan, Mei-Xuan Li, Chu-Bin Lin, and Yanzhi Wang. Organization capital effect in stock returns—the role of r&d. *Journal of Business Finance & Accounting*, 49(7-8):1237–1263, 2022.

X. Chang, Y. Chen, and L. Zolotoy. Stock liquidity and stock price crash risk. *Journal of Financial and Quantitative Analysis*, 52(4):1605–1637, 2017.

Carl R. Chen, Hai Hong, and Jeremy C. Stein. Forecasting crashes: Trading volume and predictability in stock returns. *Journal of Financial Economics*, 61(1):317–332, 2001.

J. Chen, K. C. Chan, W. Dong, and F. Zhang. Internal control and stock price crash risk: Evidence from china. *European Accounting Review*, 26(1):125–152, 2017.

L. Chen, A. C. H. Lei, and C. Song. Organization capital and stock price crash risk. *Pacific-Basin Finance Journal*, 88:102587, 2024.

S. Chen, M. Insley, and T. Wirjanto. The impact of stochastic convenience yield on long-term forestry investment decisions. *SSRN Electronic Journal*, 2012.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794. ACM, 2016.

Y. Chen, J. Wan, and Y. Ge. Q-learning for pricing american options. *Journal of Derivatives*, 27(1):1–15, 2019.

Shijun Cheng. R&d expenditures and ceo compensation. *The Accounting Review*, 79(2): 305–328, 2004.

CME. CME Hedging Crop Insurance. https://www.cmegroup.com/articles/2023/hedging-crop-insurance-with-new-crop-weekly-options.html, 2025.

J. Core and W. Guay. Estimating the value of employee stock option portfolios and their sensitivities to price and volatility. *Journal of Accounting Research*, 40(3):613–630, 2002.

J. E. Core, W. R. Guay, and R. E. Verrecchia. Price versus non-price performance measures in optimal ceo compensation contracts. *The Accounting Review*, 78(4):957–981, 2003.

N. Crouzet, J. C. Eberly, A. L. Eisfeldt, and D. Papanikolaou. The economics of intangible capital. *Journal of Economic Perspectives*, 36(3):29–52, 2022.

C. Custódio, M. A. Ferreira, and P. Matos. Generalists versus specialists: Lifetime work experience and chief executive officer pay. *Journal of Financial Economics*, 108(2):471–492, 2013.

Z. Dai, H. Yu, B.K.H. Low, and P. Ho. Bayesian optimization meets bayesian optimal stopping. In *International Conference on Machine Learning (ICML)*, 2019.

P. M. Dechow, R. G. Sloan, and A. P. Sweeney. Detecting earnings management. *The Accounting Review*, 70(2):193–225, 1995.

P. Demerjian, B. Lev, and S. McVay. Quantifying managerial ability: A new measure and validity tests. *Management Science*, 58(7):1229–1248, 2012.

Elroy Dimson. Risk measurement when shares are subject to infrequent trading. *Journal of Financial Economics*, 7(2):197–226, 1979.

A. K. Dixit and R. S. Pindyck. *Investment under Uncertainty*. Princeton University Press, Princeton, 1994.

A. L. Eisfeldt and D. Papanikolaou. The value and ownership of intangible capital. *American Economic Review*, 104(5):189–194, 2014.

Andrea L Eisfeldt and Dimitris Papanikolaou. Organization capital and the cross-section of expected returns. *The Journal of Finance*, 68(4):1365–1406, 2013.

J. Ery and L. Michel. Solving optimal stopping problems with deep q-learning. *SSRN Electronic Journal*, 2021.

Christian-Oliver Ewald, Ruolan Ouyang, and Tak Kuen Siu. On the market-consistent valuation of fish farms: Using the real option approach and salmon futures. *American Journal of Agricultural Economics*, 99(1):207–224, 2017.

X. Fang, J. Pittman, and Y. Zhao. The role of auditor personal characteristics in audit quality: Evidence from china. *Working Paper*, 2017.

X. Fang, C. Girardone, Y. Li, and Y. Zeng. Generalist ceos and stock price crash risk. *Journal of Business Finance & Accounting*, 50(1-2):182–221, 2023.

A. Fathan and E. Delage. Deep reinforcement learning for optimal stopping with application in financial engineering. *arXiv preprint arXiv:2105.08877*, 2021.

Gero W. Fischer and W.E. Grant. Use of a native predator to control overcrowding in warm-water polyculture ponds. *Ecological Modelling*, 72(3):205–227, 1994.

O. I. Forsberg and A. G. Guttormsen. The value of information in salmon farming. *Aquaculture Economics & Management*, 10(3):183–200, 2006.

W. B. Francis, I. Hasan, H. Zhang, and X. Zhang. The impact of organization capital on firm innovation. *Journal of Financial Stability*, 53:100829, 2021.

Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.

Xi Fu and Zhifang Zhang. Cfo cultural background and stock price crash risk. *Journal of International Financial Markets, Institutions and Money*, 62:74–93, 2019. doi: 10.1016/j.intfin.2019.05.001.

Sofia Gamito. Growth models and their use in ecological modelling: an application to a fish population. *Ecological Modelling*, 113(1):83–94, 1998.

I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2017.

D. Gounopoulos, G. Loukopoulos, and P. Loukopoulos. Ceo education and the ability to raise capital. *Corporate Governance: An International Review*, 29(1):67–99, 2021.

Robert M. Grant. Toward a knowledge-based theory of the firm. *Strategic Management Journal*, 17(Winter):109–122, 1996.

A. Guttormsen. Faustmann in the sea: Optimal rotation in aquaculture. *Marine Resource Economics*, 23(4):401–410, 2008.

J. Hainmueller. Entropy balancing for causal effects: A multivariate reweighting method. *Political Analysis*, 20(1):25–46, 2012.

J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

Lars Peter Hansen and Thomas J Sargent. Robust control and model uncertainty. *The American Economic Review*, 91(2):60–66, 2001.

M. M. Hasan and A. Cheung. Organization capital and firm life cycle. *Journal of Corporate Finance*, 48:556–578, 2018a.

M. M. Hasan and M. R. Uddin. Do intangibles matter for corporate policies? evidence from organization capital and corporate payout choices. *Journal of Banking & Finance*, 135: 106395, 2020.

M. M. Hasan, G. Taylor, and G. Richardson. Brand capital and stock price crash risk. *Management Science*, 68(10):7221–7247, 2022.

Mostafa Monzur Hasan and Adrian Cheung. Organization capital and firm risks. *China Accounting and Finance Review*, 2018b.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition, 2009.

James J Heckman. Sample selection bias as a specification error. *Econometrica*, pages 153–161, 1979.

J. M. Hernández, E. Gasca-Leyva, C. J. León, and J. M. Vergara. A growth model for gilthead seabream. *Ecological Modelling*, 165(2-3):265–283, 2003.

Bengt Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1): 74–91, 1979.

Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

R. Hu. Deep learning for ranking response surfaces with applications to optimal stopping problems. *arXiv preprint arXiv:1812.04077*, 2019.

A. P. Hutton, A. J. Marcus, and H. Tehranian. Opaque financial reports, r^2, and crash risk. *Journal of Financial Economics*, 94(1):67–86, 2009.

F. Jamshidian and M. Fein. Closed-form solution for american options. *Working Paper*, 1990.

K. Jebran, S. Chen, and R. Zhang. Board social capital and stock price crash risk. *Review of Quantitative Finance and Accounting*, 59:499–540, 2022.

Michael C Jensen and William H Meckling. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4):305–360, 1976.

N. Jia. Tournament incentives and stock price crash risk. *Accounting Horizons*, 32(3):101–121, 2018.

J. Jiang, K. R. Petroni, and I. Y. Wang. Cfos and ceos: Who have the most influence on earnings management? *Journal of Financial Economics*, 96(3):513–526, 2010.

L. Jin and S. C. Myers. r^2 around the world: New theory and new tests. *Journal of Financial Economics*, 79(2):257–292, 2006.

Stewart Jones, David Johnstone, and Roy Wilson. Predicting corporate bankruptcy: An evaluation of alternative statistical frameworks. *Journal of Business Finance & Accounting*, 44(1–2):3–34, 2017. doi: 10.1111/jbfa.12218.

J. B. Kim, Y. Li, and L. Zhang. Corporate tax avoidance and stock price crash risk: Firm-level analysis. *Journal of Financial Economics*, 100(3):639–662, 2011a.

J. B. Kim, Z. Wang, and L. Zhang. Ceo overconfidence and stock price crash risk. *Contemporary Accounting Research*, 33(4):1720–1749, 2016.

J. B. Kim, L. Y. Lu, and Y. Yu. Analyst coverage and expected crash risk: Evidence from exogenous changes in analyst coverage. *The Accounting Review*, 94(4):345–364, 2019.

Jeong-Bon Kim, Yinghua Li, and Liandong Zhang. Ceos versus cfos: Equity incentives and crashes. *Journal of Financial Economics*, 101(3):713–730, 2011b.

Y. Kim, H. Li, and S. Li. Corporate social responsibility and stock price crash risk. *Journal of Banking & Finance*, 43:1–13, 2014.

M. Kohler, A. Krzyzak, and N. Todorovic. Pricing of american options by neural networks. *Mathematical Finance*, 20(3):383–410, 2010.

Yasemin Y. Kor and Andrea Mesko. Dynamic managerial capabilities: Configuration and orchestration of top executives' capabilities and the firm's dominant logic. *Strategic Management Journal*, 34(2):233–244, 2013.

Kin Wai Lee, Baruch Lev, and Gillian Yeo. Organizational structure and earnings management. *Journal of Accounting, Auditing & Finance*, 22(2):293–331, 2007. doi: 10.1177/0148558X0702200213.

Derek Lemoine and Christian P Traeger. Watch your step: optimal policy in a tipping climate. *American Economic Journal: Economic Policy*, 6(1):137–166, 2014.

PingSun Leung, Eithan Hochman, Lawrence W. Rowland, and James A. Wyban. Modeling shrimp production and harvesting schedules. *Agricultural Systems*, 32(3):233–249, 1990.

B. Lev and S. Radhakrishnan. The valuation of organization capital. In *Measuring capital in the new economy*, pages 73–110. University of Chicago Press, 2005.

B. Lev, S. Radhakrishnan, and W. Zhang. Organization capital. *Abacus*, 45(3):275–298, 2009.

A. Lewbel. Using heteroscedasticity to identify and estimate mismeasured and endogenous regressor models. *Journal of Business & Economic Statistics*, 30(1):67–80, 2012.

M. León-Santana and J. M. Hernández. The effects of water temperature in aquaculture management. *Applied Economics*, 38(18):2159–2168, 2006.

K. Li, B. Qiu, and R. Shen. Organization capital and mergers and acquisitions. *Journal of Financial and Quantitative Analysis*, 53(4):1871–1909, 2018.

X. Li, S. S. Wang, and X. Wang. Trust and stock price crash risk: Evidence from china. *Journal of Banking & Finance*, 76:74–91, 2017.

Y. Li, C. Szepesvari, and D. Schuurmans. Learning to exercise american options. *University of Alberta Working Paper*, 2009.

L. Liyanage, D.S. Zois, and C. Chelmis. Optimal stopping for civic engagement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019.

Samuel H. Logan and Warren E. Johnston. Broodstock replacement policies for rainbow trout under photoperiod control and uncertainty. *Aquaculture*, 106(2):99–115, 1992.

Francis A. Longstaff and Eduardo S. Schwartz. Valuing american options by simulation: A simple least-squares approach. *The Review of Financial Studies*, 14(1):113–147, 2001.

X. Luo and C.B. Bhattacharya. The debate over doing good: Corporate social performance, strategic marketing levers, and firm-idiosyncratic risk. *Journal of Marketing*, 73:198–213, 2009.

Yunbiao Ma, Xiaotong Yang, Yuan Zhang, and Bing Zhu. Product market competition and stock price crash risk: Evidence from china. *The International Journal of Accounting*, 59 (03):1–37, 2024.

A. Marwick, M. M. Hasan, and T. Luo. Organization capital and corporate cash holdings. *International Review of Financial Analysis*, 68:101458, 2020.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

Arnaud Muller-Feuga. Modélisation de la croissance des poissons en élevage. Technical Report 21, IFREMER, 1990. 58 pp.

OECD Nuclear Energy Agency. Carbon pricing, power markets and the competitiveness of nuclear power. Technical report, OECD Publishing, 2011.

Bernt Øksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, 6th edition, 2003.

Lin Peng and Richard Roll. Investor diversification and the equity premium: major findings and policy implications. *Financial Analysts Journal*, 64(2):68–73, 2008.

Goran Peskir and Albert Shiryaev. *Optimal Stopping and Free-Boundary Problems*. Birkhäuser, 2006.

R. H. Peters and L. A. Taylor. Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272, 2017.

R. S. Pindyck. Irreversibilities and the timing of environmental policy. *Resource and Energy Economics*, 22(3):233–259, 2000.

J. Pizarro and E. Schwartz. Fisheries optimal harvest under price and biomass uncertainty. *Environmental and Resource Economics*, 78:147–175, 2021.

Reuters. Boeing faces delivery halts and share price decline due to 737 max production failures, 2023. URL https://www.reuters.com. Accessed: 2025-01-20.

Reuters. Intel loses $7 billion in foundry operations, erasing $32 billion in market value, 2025. URL https://www.reuters.com. Accessed: 2025-01-20.

W.E. Ricker. Growth rates and models. In *Fish Physiology, III, Bioenergetics and Growth*, pages 677–743. Academic Press, 1979.

Thomas A Rietz. The equity risk premium: a solution. *Journal of Monetary Economics*, 22 (1):117–131, 1988.

Ashok J. Robin and Hao Zhang. Do industry-specialist auditors influence stock price crash risk? *Auditing: A Journal of Practice & Theory*, 34(3):47–79, 2015.

Leonard C. G. Rogers. Monte carlo valuation of american options. *Mathematical Finance*, 12(3):271–286, 2002.

G. Rothwell. A real options approach to evaluating new nuclear power plants. *The Energy Journal*, 27(Special Issue):37–53, 2006.

J.-D. M. Saphores. Environmental uncertainty and the timing of environmental policy. *Natural Resource Modeling*, 17(2):163–190, 2004.

J.-D. M. Saphores and P. Carr. Real options and the timing of implementation of emission limits under ecological uncertainty. Available at SSRN, 2019.

Eduardo S Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997.

W. Shi and C. Xia. Human capital and stock price crash risk. *Working Paper*, 2023.

A. M. Spence. Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics*, 8(2):534–544, 1977.

James H. Stock and Motohiro Yogo. Testing for weak instruments in linear iv regression. In Donald W. K. Andrews and James H. Stock, editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press, 2005. doi: 10.1017/CBO9780511614491.006.

Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.

Z. Tan and Y. Wu. On regime switching models. *Mathematics*, 13(7):1128, 2025.

P. J. Thomas and N. Chrysanthou. Using real options to compare the economics of nuclear power and wind power with electricity from natural gas. *Proceedings of the Institution of Mechanical Engineers*, 226(A3):398–412, 2012.

H. Tong. On a threshold model. In *Pattern Recognition and Signal Processing*, pages 575–586. Sijthoff & Noordhoff, 1978.

J. Tvedt. Optimal entry and exit decisions under uncertainty and the impact of mean reversion. *Oper. Res. Forum*, 3:54, 2022.

T. Vukina, D.-F. Li, and D.M. Holthausen. Hedging with crop yield futures: A mean-variance analysis. *American Journal of Agricultural Economics*, 78(4):1015–1025, 1996.

Martin L Weitzman. On modeling and interpreting the economics of catastrophic climate change. *The Review of Economics and Statistics*, 91(1):1–19, 2009.

K. Wu and S. Lai. Intangible intensity and stock price crash risk. *Journal of Corporate Finance*, 73:101682, 2022.

X. Xu. Generalist ceos, management risk and internal control weaknesses. *Journal of Business Finance & Accounting*, 51(1-2):209–239, 2023.

F. Yu. Analyst coverage and earnings management. *Journal of Financial Economics*, 88(2): 245–271, 2008.

R. Yu and P. Leung. Optimal partial harvesting schedule for aquaculture operations. *Marine Resource Economics*, 21(3):301–315, 2006.

L. Zhu. A simulation based real options approach for the investment evaluation of nuclear power. *Computers & Industrial Engineering*, 63(3):585–593, 2012.