



جامعة خليفة
Khalifa University

Deep Real Options: Valuation of Real Options on Green Energy using Deep Learning Methods

Ahmed Alqubaisi

MSc. Thesis

July 2023

A thesis submitted to Khalifa University of Science and Technology in accordance with the requirements of the degree of MSc. in Computational Data Science in the Department of Electrical Engineering and Computer Science.



جامعة خليفة
Khalifa University

Deep Real Options: Valuation of Real Options on Green Energy using Deep Learning Methods

by

Ahmed Alqubaisi

A thesis submitted in partial fulfillment of the
requirements for the degree of

MSc. in Computational Data Science

at

Khalifa University

Thesis Committee

Dr. Yerkin Kitapbayev (Main Advisor),
Khalifa University

Dr. Jorge Passamani Zubelli (Co-
Advisor)
Khalifa University

Dr. Giorgio Consigli (RSC Member 1),
Khalifa University

Dr. Sajid Javed (RSC Member 2)
Khalifa University

July 2023

Abstract

Ahmed Alqubaisi, “**Deep Real Options: Valuation of Real Options on Green Energy using Deep Learning Methods**”, M.Sc. in Computational Data Science, Department of Electrical Engineering and Computer Science, Khalifa University of Science and Technology, United Arab Emirates, July 2023.

This paper investigates using deep neural networks to value real options in green energy context. We look at testing the capability of neural networks to value American options first and then apply them to real options. Two projects are compared, where one incorporates stochastic variable cost, a gas powered plant, and the other with fixed costs, a wind powered plant. The Least Squared Monte Carlo valuation method for American options was used as a benchmark. A simple fully connected network was tested on benchmark simulated data and showed promising results. In essence, the model learns and then predicts continuation values in discrete time intervals and compares them to the option payoff. The model however did not extend training accuracy when dealing with testing data of real options. We conclude that neural networks have potential in valuing real options though may suffer from over- fitting to its trained data.

Indexing Terms: Real Options, Deep Learning, Optimal Stopping, Option Valuation, Monte Carlo, Green Energy

Acknowledgement

I would like to express my deepest gratitude to my family and my sponsor ADIA for supporting me through my research and thesis work.

I am also thankful for the support and guidance provided by my advisors, committee members, and friends throughout the journey.

Declaration and Copyright

Declaration

I declare that the work in this thesis was carried out in accordance with the regulations of Khalifa University of Science and Technology. The work is entirely my own except where indicated by special reference in the text. Any views expressed in the thesis are those of the author and in no way represent those of Khalifa University of Science and Technology. No part of the thesis has been presented to any other university for any degree.

Author Name: Ahmed Alqubaisi

Author Signature:

A handwritten signature in black ink, appearing to read 'ahmed B', written over a horizontal line.

Date: 24/7/2023

Copyright ©

No part of this thesis may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without prior written permission of the author. The thesis may be made available for consultation in Khalifa University of Science and Technology Library and for inter-library lending for use in another library and may be copied in full or in part for any bona fide library or research worker, on the understanding that users are made aware of their obligations under copyright, i.e. that no quotation and no information derived from it may be published without the author's prior consent.

Contents

Abstract	I
Acknowledgement	II
Declaration and Copyright	III
Contents	IV
List of Figures	V
List of Tables	VI
CHAPTER 1. Introduction	1
1.1 Background and Motivation	1
1.2 Objective of the Study	2
1.3 Scope and Limitations	3
1.4 Organization of the Thesis	3
CHAPTER 2. Literature Review	4
2.1 Real Options Theory	4
2.2 Real Options in Energy Investments	6
2.3 Optimal Stopping and Neural Networks	7
CHAPTER 3. Model and Methodology	8
3.1 Data Curation	8
3.2 Real Options Valuation Framework	8
3.2.1 Valuation Method Selection	10
3.3 Deep Learning Models	10
3.3.1 Model Architecture	10
3.3.2 Model Training and Validation	13
3.3.3 Hyperparameter Optimization	13
3.4 Model Design and Evaluation	14
CHAPTER 4. Results	19
4.1 Real Options Valuation Results	20
CHAPTER 5. Conclusion	22
5.1 Summary of Findings	22
5.2 Contribution to the Field	22
5.3 Practical Implications	22
5.4 Recommendations for Future Research	23
References	24

List of Figures

1	NPV vs Real Option Over Time	1
2	Early exercise boundary for a hypothetical call real option on oil. The shaded region illustrates the prices and times where it is more beneficial to exercise the option (drill for oil) rather than hold the option (wait).	4
3	Comparison between the true function and the neural network approximation.	11
4	Simulated Asset Prices	14
5	Neural Network Architecture	16
6	Backward Induction of Model Weights	17
7	Monte Carlo Path 1	17
8	Monte Carlo Path 2	17
9	Steps to Finding Option Value	18
10	Model 2 Loss on American Options	20
11	Model 2 Loss on Real Options	20

List of Tables

1	Types of real options in investment projects [3]	5
2	Comparison of Option Values between Glasserman's textbook and our model	14
3	American Option Values with Different Simulations	19
4	Real Option Values with Different Simulations	20
5	Comparison of LSM and Neural Network with 100 Simulations	21

Chapter 1. Introduction

1.1 Background and Motivation

Projects carried out by companies are evaluated by their economic and physical feasibility. Economic feasibility is determined by the expected returns of the project. One can model the returns through measuring a project's net present value (NPV) or internal rate of return (IRR). These measures rely on estimating cash flows of an investment over a time horizon. An issue with this approach is that investments hold uncertainty with regards to how successful they may be and estimating cash flows are not guarantees. In addition, the option value of waiting to start an investment is not captured by these traditional methods. Thus, it is important to look at tools that offer flexibility in valuing these investments. Here, viewing real-world investments as options and to follow theory of valuation of financial option contracts become useful.

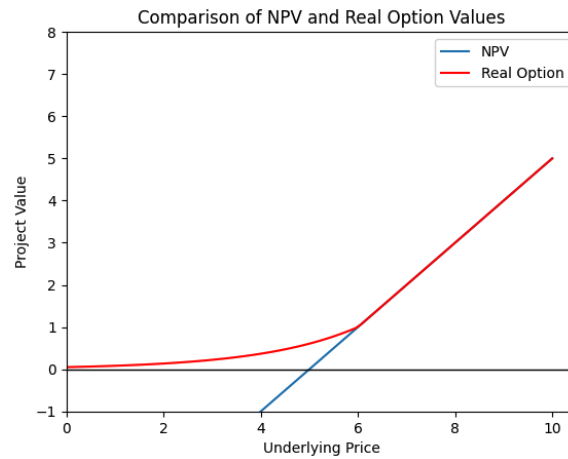


Figure 1: NPV vs Real Option Over Time

The figure above is a hypothetical case where the real option shows positive project value while NPV of the project is negative. It is only until the underlying price goes beyond 4.5 does the project break even. This is due to a real option capturing the value of waiting and gives a positive value, even when underlying prices are low.

1.2 Objectives of the Study

In light of increased complexity of investments due to interlinkage of factors, managerial assumptions on cashflow amounts, timing of cashflow, and probability of success and failure could be far from reality. The dynamic nature of elements that affect an investment cannot be fully implemented in traditional methods of valuation. The financial options framework allows for including the value of starting a project earlier or later than initially planned. However, within the options framework, specifically American options, therein lies a problem of trying to properly price the option for given stochastic model of underlying price. American options can be exercised at any time before maturity T and hence in general are more expensive than European options. This early exercise feature makes the problem of pricing and hedging more complicated. For example, application of Monte-Carlo method is not obvious in this case. It can be shown that the fair price of the American option is given by

$$V = \max_{\tau \in [0, T]} \mathbb{E}[e^{-r\tau} h(X_\tau)] \quad (1)$$

where τ is the stopping time with respect to the underlying price process X and h is the payoff function.

This is called an optimal stopping problem. Stopping time by definition is the random time with values between 0 and T that is adapted to the information generated by trajectory of X . We cannot use the future information to determine the value of stopping time τ . In our context, stopping times represent possible investment timing strategy given the current public information on the market. Real options in its simplest form is an American option with a different underlying. An optimal stopping problem is presented with valuing real options and we wish to explore this problem by applying neural networks. In this thesis, we study the problem of a firm seeking to build an energy plant. The company can choose between two alternative technologies, build a wind plant (*WP*) or a gas-fired plant (*GFP*). The model extends on the paper by Detemple and Kitapbayev [1].

In their paper, there are two state variables, the electricity price X and the gas price Y , which evolve according to

$$dX_t = (r - \delta_X)X_t dt + \sigma_X X_t dW_t \quad (2)$$

$$dY_t = (r - \delta_Y)Y_t dt + \sigma_Y Y_t dB_t \quad (3)$$

under the risk-neutral measure Q . The interest rate r , the volatility coefficients σ_X, σ_Y and the convenience yields δ_X, δ_Y are constant. Uncertainty is modeled by a pair of correlated Brownian motions W, B with the correlation parameter $\rho \in (-1, 1)$.

The WP sells electricity at price X and collects a subsidy s per MWh of electricity sold. It has a constant running cost k_w . The value of the WP once it is built can be calculated as

$$\begin{aligned} W(x; s) &= \mathbb{E} \left[\int_0^\infty e^{-rt} (X_t + s - k_w) dt \right] \\ &= \frac{x}{\delta_X} + \frac{s - k_w}{r} \end{aligned} \quad (4)$$

for $x > 0$, where we assumed that the plant is operated indefinitely.

The GFP buys gas at the price Y and sells electricity at the price X . It has a constant

running cost k_g . The value of GFP, once it is built, is given by

$$\begin{aligned} G(x; y) &= \mathbb{E} \left[\int_0^\infty e^{-rt} (X_t - Y_t - k_g) dt \right] \\ &= \left(\frac{x}{\delta_X} - \frac{y}{\delta_Y} \right) - \frac{k_g}{r} \end{aligned} \quad (5)$$

for $x, y > 0$, also assuming indefinite operation.

The investor has the option to choose the technology for power production at the time of investment. Clearly, she will select the one that maximizes the net present value of its investment, i.e., the best of the two technologies. The investor also chooses the optimal time to invest. The investor solves the following problem

$$V_{(x,y)} = \sup_{\tau \geq 0} \mathbb{E} \left[e^{-r\tau} \max(G(X_\tau, Y_\tau) - K_g, W(X_\tau) - K_w) \right] \quad (6)$$

where K_w and K_g are the investment costs. This problem has been solved in Detemple and Kitapbayev [1] using probabilistic methods. In the current thesis, we apply Monte-Carlo approach along with regression and neural network based methods.

1.3 Scope and Limitations

The target area of this research deals with real options on green energy plants that have discrete exercise dates with fixed system parameters such as risk free rate, convenience yield, volatility, and correlation. The price process of electricity and gas are assumed to be geometric Brownian motion. Training and testing neural networks get computationally intensive quickly, and such, network architecture for this research was kept simple. Another limitation of this research is that we used a model based approach. The risk of misspecification and approximation is inherited in this approach. The nature of investment opportunities is complex and assuming a single framework is suitable for every project is not possible. Sensitivity to input changes is also a limitation of this research as most input parameters are fixed.

1.4 Organization of the Thesis

The thesis is divided into four main parts, literature review explains real options and the context of green energy as well as deep learning techniques, model and methodology goes through data curation, valuation methods, and neural network architecture. Results section discusses different outputs of the model and the performance. The thesis then ends with a conclusion of our finding and recommendations for future research.

Chapter 2. Literature Review

2.1 Real Options Theory

Real options, like financial options, give the holder the right, but not the obligation to carry out a certain decision on the future. Real options hold high value to those in charge of project management. While the underlying in financial options are stocks, in real options the underlying can be the value of new plants or new products. As such, value of real options depend on the value of stated underlying assets. An analogy of strike price for the real options is seen as initial cost to develop or start the project, like drilling for an oil well that may be used in the future. In the textbook by Dixit and Pindyck [2], real options are studied in depth and it provides the logic and attributes of real options. The concept of real options is based on the irreversibility of the investment entailing sunk costs and other implications, an aspect that is not fully captured by NPV. Another aspect not captured by NPV is the value of waiting and delaying the investment. Traditional theory says that when the NPV of a project is positive, it should be started. With real option theory, the option to delay is of value. Intuitively, it's shown that the value of the option has a higher value when uncertainty is high. When firms are faced with a volatile environment, they would highly prefer an option to wait and see what will happen in the future. These uncertainties vary from project to project, but they usually fall under the realm of exogenous and endogenous prices, regulation, and technological uncertainties.

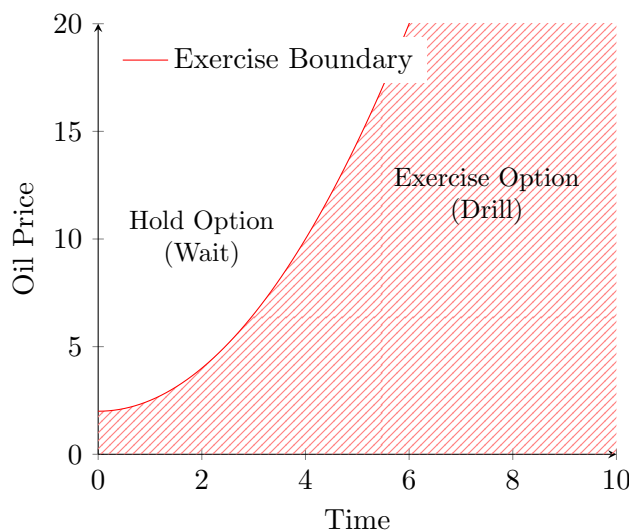


Figure 2: Early exercise boundary for a hypothetical call real option on oil. The shaded region illustrates the prices and times where it is more beneficial to exercise the option (drill for oil) rather than hold the option (wait).

The theory of real options presents similar mathematical background to that of American options. In both cases, uncertainty is usually generated by a stochastic process. They both find solutions to their differential equations through Ito's Lemma. This benefits the research of real options as breakthroughs in American option pricing transfers over to real options theory and vice versa.

Within the literature, an important chapter of Investment Under Uncertainty [2] lies within Chapter 5 as it explains the rationale of investment decisions faced by firms. Dixit

and Pindyck argue that firms must consider the costs and benefits of investing in uncertain environments. The chapter introduces the concept of a “trigger level” or a threshold, which is the level of profitability that must be reached before the firm decides to invest. This threshold is influenced by factors such as the level of uncertainty and the cost of waiting. The chapter also highlights the role of flexibility in investment decisions. Flexibility allows firms to adjust their investment decisions in response to new information or changes in the environment. Dixit and Pindyck notes that flexibility can be achieved through real options such as the option to delay, expand or contract the investment. These options can be viewed as an additional source of value for the firm, known as the option value of the investment. Dixit and Pindyck also discusses the importance of sunk costs, which are costs that cannot be recovered once they have been incurred. Sunk costs can influence investment decisions by creating a bias towards continuing to invest in a project even if it is no longer profitable.

Table 1: Types of real options in investment projects [3]

Type	Description
Option to delay	Postpone investment decisions to wait for new information. By delaying the investment, managers can learn more about the project and reduce uncertainty, allowing them to make better-informed decisions.
Option to expand	Increase the size of an investment or project if it becomes profitable. By having the option to expand, managers can scale up the project based on positive outcomes, increasing the potential return on investment. This option has value because it provides upside potential without committing to a larger investment upfront.
Option to reduce	Reduce the size of an investment or project if it becomes unprofitable. By having the option to contract, managers can limit their losses and reduce exposure to negative outcomes. This option has value because it provides downside protection without committing to a smaller investment upfront.
Option to abandon	Abandon an investment or project if it becomes unprofitable or if better opportunities arise. By having the option to abandon, managers can limit their losses and free up resources for other opportunities. This option has value because it provides managers with an exit strategy, reducing the downside risk associated with investing in a project.
Option to switch	Switch between different investment opportunities based on new information or changes in the environment. By having the option to switch, managers can adjust their investment strategy to take advantage of changing market conditions. This option has value because it provides managers with flexibility to adapt to changing circumstances.

The focus of this thesis involves the option of choosing between fixed costs and variable cost project. Abandoning the project or reducing the output is not feasible as energy needs are constant and required to meet an average quota. A paper by Decamps et al. [4] show the dichotomous nature of investing in alternative projects as well as showing an indifference region where its optimal to wait even when initial condition surpass optimal thresholds. This idea is carried on to this research as we the same problem formulation. Large amount of literature focus on single projects, such as [5], [6], [7], and [8]. Bakke [9] investigates changing different parameters of real option framework and illustrates how volatility does not necessarily increase the value of an option. Our research looks into two projects simultaneously and requires foundations made in papers such as Decamps. Real option theory is intertwined with American options research. Exotic option structures were discussed in papers such as [10]. These options are more closely related to real options than basic or vanilla American options. In the section for problem formulation, the value of the real option is that of an American option, restating the overlap.

2.2 Real Options in Energy Investments

There are several applications of real options in energy investments. Venetsanos [11] showed that an option valuation to a wind energy farm had a positive option value, but a negative net present value. As discussed in earlier sections, this difference is apparent as traditional methods do not account for the value of waiting. The traditional method used in the paper was discounted cashflows and evaluated the option's value with the BSM (Black-Scholes Merton) model. Real options do not only apply to deciding on one project, but also can be used to understand the behavior of an industry. Kjarland [12] discussed optimal decisions in investment of hydropower plants as a whole within the context of his home country of Norway. The necessity to apply such framework appeared as when electricity became a competitive market in the country, firms were making decisions in uncertain environment of electricity prices. The paper aimed to understand the behavior of the hydropower industry as a whole by using real option valuation. Real option framework was used to tackle policy decisions as well. Ritzenhofen [13] discussed optimal design of feed-in-tariffs to stimulate renewable energy investments under regulatory uncertainty. The paper concluded with statements on how regulations affect impact of tariffs on renewable energy investments and tariffs create a binary timing outcome for renewable investments. Detemple [1] show distinction from fossil fuel energy is relevant as they inflect different cost dependencies. Fossil fuel plants costs are dependent on the cost of fuel used for making energy as well as having fixed costs with regards to their plant operation and maintenance. GE plants on the other hand have free input and large fixed costs. Indeed, valuing GE investment options is a crucial tool for green energy producers as well as regulators. The importance of this valuation is key in decision making as an optimal route is studied and implemented. Though GE technology is improving, it is yet to compete with fossil fuels. Consequently, optimal investment decision bring benefit to pushing for the use of GE.

The International Energy Agency publishes a yearly outlook to world energy that addresses market developments in the energy sector. As energy producers face a volatile industry, it is a changing landscape. Political, geopolitical, and technological changes have vast impact on the market. Looking at the most recent paper of 2022, [14] the agency addressed Russia's invasion of Ukraine effect on the energy market. As there are network impacts with energy markets, regulation was hastened to adopt clean energy and energy independence. Transformation to clean electricity is also challenged by supply chain and

affordable energy issues. This is year specific and energy impacts can leak from year to year. This illustrates the fact that energy plants supply and demand is guided by market uncertainty. This again emphasizes the importance of a model to take into account such uncertainties and protect firms from taking decisions without a full picture.

2.3 Optimal stopping and Neural Networks

As previously stated, real option valuation is typically an optimal stopping time problem. In the context of real options, the optimal stopping time occurs when the payoff of exercising the option now is greater than the continuation value. The literature presents several ways of finding the optimal time to exercise an option. [15] illustrates the main techniques of integral equations for optimal exercise boundaries, parametric option pricing models, and jump diffusions. Real options usually take the route of simulation to solve the optimal exercise problem. Simulations conveniently tackle the issue of path dependencies. This line of research was based on the LSM approach (Least Squared Monte Carlo) [16]. The value of continuation in this approach is found by fitting a linear combination of basis functions. The value of the option is found on all time steps over all simulations by backward induction from its maturity. As such, the model can be fed a set of cashflows of different simulations, and the output would be a value of an option. Several numerical methods exist currently, though this research aimed to include recent computational advancements in machine learning.

Deep Learning has emerged as a machine learning tool used for many applications. These include computer vision, natural language processing, and speech recognition. The goal of deep learning algorithms is to extract features from the input. This is done by setting weights to values of the inputs, moving these weights through layers of processing the input, then trying to optimize the weight by comparing it to a 'true' value. The network of layers process the data of the input extracting features layer to layer. Then, the model produces an output of classification or prediction by applying the learned weights to a new input. For the purpose of this thesis, a basic neural network is the starting point. Starting with a basic neural network will serve as a benchmark for adopting other complex architectures. The direction of supervised vs unsupervised for the purpose of the thesis is clear; a supervised learning technique is appropriate as the problem of valuation is a regression under our model.

Previous work on deep learning and financial options include [17]. The authors use the LSM from the Longstaff Schwartz paper mentioned earlier to find optimal stopping times through feeding it to a neural network. The reasoning behind using deep neural networks is the ability for the model to learn stopping times rather than trying to solve the boundary problem explicitly. Although this paper involves hedging the value of the option, an aspect outside the scope of the research, the approach to building the neural network is valuable. The paper finds the option price by first computing an optimal stopping policy that also outputs a lower bound. An upper bound is also computed with a confidence interval.

Chapter 3. Model and Methodology

3.1 Data Curation

This research simulated data of stock prices and project values by Monte Carlo simulation. Monte Carlo simulation holds an advantage in the ability to offer large amounts of data, an aspect crucial for deep learning. In addition, as our problem is that of financial forecasting, and simulations provide an efficient way to tackle this problem. [18] Linear regression estimation of American option pricing from [19] was used as a benchmark. The results from the textbook simulation was compared to the model created for this thesis to make sure the simulation and output were correct. In addition, the data simulated were direct inputs to models and no preprocessing was needed.

1. Define geometric Brownian Motion:

$$dS_t = \mu S_t dt + \sigma S_t dW_t$$

where S_t is the price of the financial instrument at time t , μ is the expected return, σ is the volatility, and dW_t is a Wiener process.

2. Simulate Paths for the underlying

$$S_{t+\Delta t} = S_t + \mu S_t \Delta t + \sigma S_t \sqrt{\Delta t} Z$$

where Z is a random variable following a standard normal distribution. You repeat this step N times to generate N different paths.

3.2 Real Options Valuation Framework

Let us consider the two risky projects with the investment payoffs $X - Y - K_1$ and $X - K_2$, respectively, with $K_1 < K_2$. The processes are X and Y given by

$$\begin{aligned} dX_t &= (r - \delta_X) X_t dt + \sigma_X X_t dW_t \\ dY_t &= (r - \delta_Y) Y_t dt + \sigma_Y Y_t dB_t \end{aligned}$$

where $r = 0.04$, $\delta_X = \delta_Y = 0.03$, $\sigma_X = \sigma_Y = 0.25$, W and B are standard Brownian motions with $\rho = 0.3$.

Let us consider now the real options problem with the (undiscounted) payoff upon exercise at time τ

$$\tilde{h}(X_\tau, Y_\tau) = \max(X_\tau - Y_\tau - K_1, X_\tau - K_2, 0).$$

We take $K_1 = 0.5$, $K_2 = 1$, $X_0 = 1$, $Y_0 = 0.5$. The option expires in $T = 3$ years and can be exercised at twelve equally spaced dates $t_i = i/4$, $i = 1, \dots, 12$.

Here we have a payoff function dependent on the values of X and Y . The two projects present are indicative of a variable cost Y_τ and a fixed cost of K_2 . By definition, K_2 must be greater than K_1 or else the maximum payoff would lie in the fixed cost project. K_1 are inclusive of all costs associated with variable cost projects and K_2 that of a fixed cost project. Other research delve into dissecting the cost, but this research focuses on first solving the investment problem. From a financial option perspective, this is a dual strike max option

where the variable cost project depends on two state variables. From the company's point of view, it wants to exercise the option if the current payoff is greater than the expectation of continuation. Put formally:

$$V = \sup_{\tau \in \mathcal{T}} \mathbb{E} G_{\tau},$$

where V is the value of the option at time 0 and G_{τ} is the discounted payoff. In our example, \mathcal{T} is the set of exercise times.

Using the exercise dates t_i , it is optimal to exercise the option at t_n for $V_t = G_n$. As such, the stopping times are:

$$t_i := \begin{cases} t_n & \text{if } G_n \geq \mathbb{E} [G_{t_{n+1}} | X_n] \\ t_{n+1} & \text{if } G_n < \mathbb{E} [G_{t_{n+1}} | X_n] \end{cases}$$

The main problem with option valuation has to do with approximating $\mathbb{E} [G_{\tau_{n+1}} | X_n]$. In other words, the continuation value mentioned in the introduction is core of the option valuation problem. Another way to approach the expectation of $G_{t_{n+1}}$ given X_n is as a function $g(X_n)$. The property of this function is to minimize $\mathbb{E} [\{G_{t_{n+1}} - c(X_n)\}^2]$. We now can introduce the backbone of this research valuation method, the LSM approach. The LSM approach mentioned in the introduction approximates the expectation by fitting known option values on a set of basis functions. The optional values are known as underlying prices are simulated and payoffs were calculated at each time-step. These basis functions can be seen as a subset of $g(X_n)$. In the LSM approach, the continuation value is calculated by regressing $\mathbb{E} [V_{i+1}(X_{i+1}) | X_i = x] = \sum_{r=1}^M \beta_{ir} \psi_r(x)$ where β_{ir} are a constant coefficients.

Regression-Based Pricing Algorithm (*see*[19])

- (i) Simulate b independent paths $\{X_{1j}, \dots, X_{mj}\}, j = 1, \dots, b$, where m are the different timesteps.
- (ii) At terminal nodes, set $\hat{V}_{mj} = h_m(X_{mj}), j = 1, \dots, b$ where h_m is the payoff function
- (iii) Apply backward induction: for $i = m - 1, \dots, 1$, \circ given estimated values $\hat{V}_{i+1,j}, j = 1, \dots, b$, use regression as above to calculate $\hat{\beta}_i = \hat{B}_{\psi}^{-1} \hat{B}_{\psi V}$; \circ set

$$\hat{V}_{ij} = \max \{h_i(X_{ij}), CV_i(X_{ij})\}, \quad j = 1, \dots, b,$$

- (iv) Set $\hat{V}_0 = (\hat{V}_{11} + \dots + \hat{V}_{1b}) / b$. [19]

This outlines the crucial elements of this research. A summary in simple terms can be written as follows: We start out with problem that involves finding an optimal stopping decision. This is due to the fact that the process of the underlying in an American options go through unknown future trajectory. A method to find the value of the option is to find the maximum expectation of the discounted payoff of the option. Since the payoff depends an underlying with unknown future trajectory, we simulate geometric brownian motion of the underlying asset. Then timesteps are created by segmenting simulations so the problem becomes a discrete time exercise problem. The algorithm works backwards to find the payoff at each timestep. Then it performs a regression using a set of basis functions to approximate the continuation value of the option given the input of an underlying. Then it compares the continuation value with the present payoff and decides whether to exercise or not. This is done on every simulation, over all time steps. A time-zero option value is then calculated by taking the average of all option values at maturity over all simulations.

3.2.1 Valuation Method Selection

We acknowledge several valuation methods for the valuation of American options. This research however focuses on comparing results from the LSM approach to that of a deep neural network. Regression and low estimators as mentioned in Glasserman's textbook [19] are not examined. However, a low estimator is approximated from the simulations done for the LSM approach. The binomial model and Black-Scholes model are methods that are outside the scope of this research, but future research could include comparing these methods with results highlighted here.

3.3 Deep Learning Models

Deep Learning models are used for several applications due to their ability to capture non-linearity of the problems at hand. The deep learning model used for this research is a supervised model with basic network architecture. The structure or the architecture of a deep learning model involves a set of layers, each with a number of nodes, where each node performs a non-linear transformation (activation functions), such as rectified linear unit, sigmoid, step function, of the input. The aim of using a deep learning model for this specific case is that $g(X_n)$ mentioned in the LSM approach is a complex function. It is unlikely to get a good approximation by linear regression, even over a set of basis functions as described above. With a neural network, outputs are modeled by a set of weights associated with the inputs, which are non-linear representations. The model tries to represent outputs over the layers and through these representations, the model can map an input to a specific output. Instead of a one-step calculation done for regression coefficients to predict outputs, in neural networks outputs are predicted after the model has transformed inputs layer after layer in increasing complexity and learned how an input corresponds to a certain output. The tool becomes valuable in terms of option valuation as an option can be valued from a complex depiction of its underlying asset instead of an elementary regression. In our case the inputs are simulations of underlying asset prices, and as we know that option value is a function of price so we approximate V by $g(X_n)$ produced by the neural network.

3.3.1 Model Architecture

Before getting into the architecture of the model, a brief overview of a neural network is as follows. A neural network is composed of units of computation named neurons. These

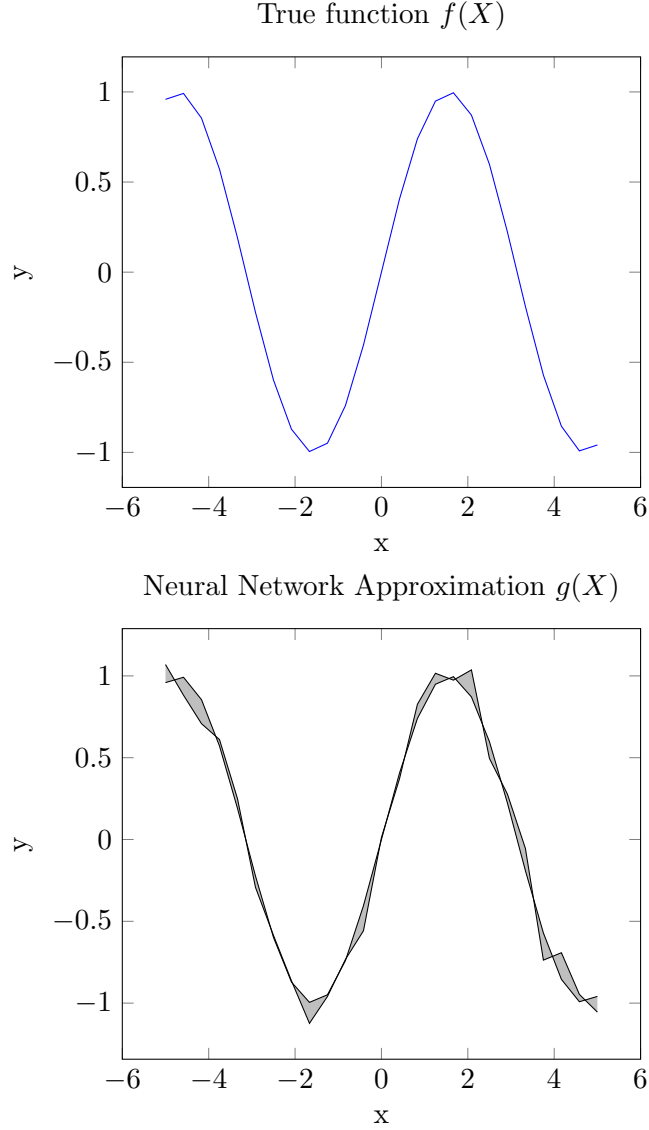


Figure 3: Comparison between the true function and the neural network approximation.

neurons take weighted sum of inputs and applies a non-linear function to it and produces and output for the next neuron. The formula is:

$$y = f\left(\sum_{i=1}^n w_i x_i + b\right)$$

where y is the output, f is the activation function, w_i are the weights assigned to the inputs x_i , b is the bias term, and n is the number of inputs. As mentioned, the output of a neuron is an input to another neuron in the next layer. The layer is just a collection of neurons represented by:

$$\hat{y} = f_L(\mathbf{w}_L f_{L-1}(\mathbf{w}_{L-1} \dots f_2(\mathbf{w}_2 f_1(\mathbf{w}_1 \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2) \dots) + \mathbf{b}_L)$$

where \hat{y} is the predicted output, \mathbf{x} is the input vector, f_l is the activation function for layer l , \mathbf{w}_l is the weight matrix for layer l , and \mathbf{b}_l is the bias vector for layer l .

The activation function can be the following functions:

Sigmoid:

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

TanhH:

$$\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$$

ReLU:

$$\text{ReLU}(x) = \max(0, x)$$

Moving to Model Architecture, the specific model for this research ran with different hyper-parameters. These hyper-parameters include the number of layers, the number of neurons in each layer, the batch size, and the number of epochs.

Layers are the basic building blocks of neural networks, and the number of layers determines the depth of the network. In our study, we are using multiple layers to capture complex patterns in the data. We expect that a shallow network will not approximate well to the true value of the option. We are also varying the number of neurons in each layer, as this can affect the network's capacity and ability to learn complex relationships in the data.

During training, the input data is divided into batches, and the network's weights are updated based on the average error computed over each batch. We are using different batch sizes to explore their effect on the network's convergence and training time.

We are also training the network for multiple epochs, which could help improve its performance. We are measuring the loss and mean absolute error (MAE) for each network during training and evaluation to assess their performance.

Loss is a commonly used evaluation metric that measures the difference between the predicted and actual values for a given input sample. We are using the mean squared error (MSE) loss function, which penalizes large errors more strongly than small errors. During training, the network aims to minimize the MSE loss function.

In addition to loss, we are also using MAE to evaluate the network's performance on a held-out validation dataset. MAE measures the average absolute difference between the predicted and actual values across all input samples. By comparing the loss and MAE for each network, we can select the best model for our task. Adjusting model architecture is normal practice in neural network research and we present these adjustment here as well.

As for overall architecture, a regression based model was chosen. A fully connected network was used as it has the basis expand upon and may serve as a benchmark for future research. The basis of our problem is minimizing loss as an objective, thus networks that are optimized architecturally for detection, generation, classification, and segmentation were not explored from the initiation of the experiment. Engaging into these different architectures introduces complications that would not add meaningful performance to the model. Neural networks are sensitive to data quality, even with abundance. Convolution Neural Networks (CNNs) and auto-encoder type of architecture perform very well for image recognition and text data respectively. Applying these models for our data may not be interpretable. CNNs work well with images as they create feature maps and use those for classifying images. [20] Auto-encoders interpret text data by dimensionality reduction and reconstruction. [21] These techniques would not help with valuation as we do not require new data generation and feature maps are not inherit to the input. The nature of the problem is investigating how a neural network is able to approximate the option value function.

This also aligns with the direction of the research. The purpose is to create a starting point for real options valuation framework starting from a basic neural network and then building upon it based on applications. Nonetheless, a 1-D convolutional neural network was tested on this problem. One dimensional convolutional networks are applied in signal processing and other various applications. [22]The network works by applying convolutional filters, scalar functions, over the input to extract features of the input data. Over its layers, the network learns which features are most important via a loss function and thus a trained network over that type of data is created. For our application, initial results on a 1-D CNN were not very encouraging and it was decided that a CNN would not be the way forward.

However, one interesting architecture that should be considered is transformers. As shown in a recent paper by Zeng et al, [23], it showed transformers can be effective in time series forecasting, so instead of generating simulations, one would calculate continuation values by using outputs from forecasting a single time-series.

3.3.2 Model Training and Validation

The model was trained using Keras API in python. Keras works on top of Tensorflow and specifically for this research we used Tensflow 2.0. The model training required input and output data. Our data is simulated asset prices and our outputs were option values. The option values were calculated explicitly at each time-step. For example, the value of the option at the third timestep is the payoff function of the asset prices at the fourth time-step. Once training input and output were created, the training input dataset was split between validation and training. The validation allows us to check the performance of the network on unseen data. This is a beneficial tool as working solely with training data could cause the model to overfit to the training data as it that would lead to the optimized loss function. The aim of model from initiation was to generalize it for many assets, and introducing a validation split is important. Another aspect that was varied was the training split. This tackles the same issue of overfitting. Generally, if most of the training data is used, say larger than 80% of the split, then overfitting is likely. General practice in deep learning aim for a split between 60% training to 80% training. For the specific case of this research, the number of parameters are high and such its beneficial to keep as much of the training data.

3.3.3 Hyperparameter Optimization

As discussed in the beginning of this section, several hyperparameters were adjusted. Care must be taken when adjusting these to not fall under the bias of data mining. The code incorporates running the neural network with varying layers, neurons, epochs, and batchsize. The range of layers were 2 to 6 layers. The approach is to see if the hypothesis that neural networks depth can properly fit into the non-linearity of option valuation. Neuron count has to be trialed to check if increasing the count too much will cause overfitting and too little will not capture the complexities of the problem. Epochs share a similar issue with neurons as its specific to the data at hand. Batch size tuning comes at the cost of converging to suboptimal solutions or taking too much time to run. So the model has to balance between all parameters to run in a timely manner without under or over fitting.

Table 2: Comparison of Option Values between Glasserman’s textbook and our model

Basis Functions	Regression	Low	LSM	Our Model
$1, S_i, S_i^2, S_i^3$	15.74	13.62	13.67	15.63
$1, S_i, S_i^2, S_i^3, S_1 S_2$	15.24	13.65	13.68	15.12
$1, S_i, S_i^2, S_i^3, S_1 S_2, \max(S_1, S_2)$	15.23	13.64	13.63	15.12
$1, S_i, S_i^2, S_i^3, S_1 S_2, S_1^2 S_2, S_1 S_2^2$	15.07	13.71	13.67	13.98
$1, S_i, S_i^2, S_i^3, S_1 S_2, S_1^2 S_2, S_1 S_2^2, \tilde{h}(S_1, S_2)$	14.06	13.77	13.79	13.94
$1, S_i, S_i^2, S_1 S_2, \tilde{h}(S_1, S_2)$	14.08	13.78	13.78	13.99

3.4 Model Design and Evaluation

The first step of creating the model was starting out with recreating the LSM approach in Glasserman’s textbook to have a starting point for benchmarking and recreate it for our real assets. We first simulated underlying prices.

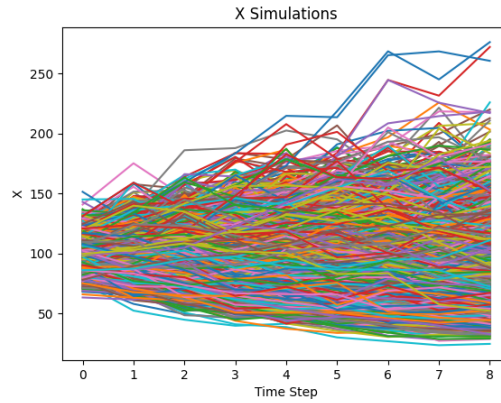



Figure 4: Simulated Asset Prices


Then we follow steps in subsection 3.2 where we conduct a regression on a set of basis functions. Table 2 shows the results

The difference between our model and the text stems from the nature of simulations. After recreating the textbook results, we implemented the neural network. The input for the neural networks would be underlying prices of timestep i and output would be the option value at timestep $i + 1$

3D Data Input Illustration



3D Data Output Illustration



These input and output are fed into a neural network and compute option values using backward induction. The neural network was made of an input layer of 1000 neurons, followed by three hidden layers of size 512,128, and 64 respectively. The output layer is a linear layer with a single output. The activation function for all hidden layers and input layer was ReLU. This activation function was chosen as it has strictly positive values, does not have a vanishing gradient problem, and is inexpensive computationally. As for loss, both mean squared error and mean absolute error were examined and recorded across epochs. The model architecture and backward induction illustration is shown below. Note that neuron count has been reduced in scale to fit the page.

The above graph illustrates how the back propagation works in the model. We estimate the option value based on the weights initially calculated by the model when we trained the last time step underlying price with the terminal value payoff. The weights as mentioned in the deep learning section are multiplied by the underlying prices at a prior timestep to get the continuation value. This continuation value is compared to the payoff and the larger value is stored in the option value matrix. The new option value matrix is then fed into the next model in the backward loop and the cycle continues til we reach initiation. We then stored each model separately to test it on a separate simulation. The testing is done in a forward manner, meaning for each simulation over every timestep, the first model predicts the continuation value based on the first timestep underlying prices. We then compared payoff and continuation value to decide whether to exercise or not for this particular simulation. If the specific simulation has a higher continuation value, we move on to the next timestep and repeat. If the simulation has a higher payoff, then we exercise for that simulation and ignore future timesteps.




Figure 5: Neural Network Architecture




Figure 6: Backward Induction of Model Weights




Figure 7: Monte Carlo Path 1




Figure 8: Monte Carlo Path 2

The graphs above illustrate how the model decides the value of the option and what values fill the array of option values for each simulation. We can see that the second generated path has an optimal stopping point labelled the exercise point where the immediate payoff is larger than what our model predicts. On the third path however, there is no stopping point and it is best to keep holding the option to maturity. To calculate the final option value, we took the average of every path and discounted it corresponding to what time it was or was not exercised.

The figure illustrates how we extracted the simulation option values and then how it became the final option value. The apparent difficulty with this model comes with evaluating




Figure 9: Steps to Finding Option Value

every simulation, given the fact that we hope to model as many simulations as we can to converge to the true option value.

Chapter 4. Results

We expected results to be lower than the theoretical true value of the option as we fit an expected value that will always be, by definition, lower than the maximum value of the option. We can see this formally below:

$$\begin{aligned} V_0 &= \max \mathbb{E}(\exp(-r\tau)G(X_\tau, Y_\tau)) \\ &= \mathbb{E}(\exp(-r\tau_*)G(X_{\tau_*}, Y_{\tau_*})) \geq \mathbb{E}[\exp(-r\tau_\theta)G(X_{\tau_\theta}, Y_{\tau_\theta})] \\ &\approx \frac{1}{N_2} \sum_{n=1}^{N_2} \exp(-r\tau_\theta)G(X_{\tau_\theta}, Y_{\tau_\theta}) \end{aligned}$$

where τ_θ is the stopping time proposed by the neural network model, τ_* is the true optimal stopping time, and N_2 is the number of simulations in the Monte Carlo. After running the model on test simulated data, the final option value is a lower bound of the true option value. After several runs on different number of simulations, the range of values are from 15.0 to 13.6. The deviation occur due to differences between simulations.

Table 3: American Option Values with Different Simulations

Simulations	Option Value
100	15.0
500	14.5
1000	13.6
2000	13.8

Comparing it with LSM, it seemed the neural network did somewhat capture the value of the options properly. Keeping in mind that the true value through a binomial model in the LSM paper was 13.90. Looking at the validation and training loss, we can see a reason why this occurred. We see in the loss, the mean absolute error decreases with epochs, but it is still high in all models. However, the pattern of validation and training loss did not show signs of overfitting. This assured that the model may be used for generalizing on other data, but it seemed this generalization is not accurate enough. For example, the second model which was used to find the third continuation value, showed the following loss graph:

This loss and variance between simulations encouraged a change within network architecture. After several hyper-parameter tuning, there was not much change with error or option value. We attempted to overfit the model to get low loss, but it seemed the neural network reached a point where loss could not be lowered further, even with high simulation, high epochs, and more hidden layers. We assume solely having price as input and training on option value predictions was enough to have the model learn option valuation. Another thing to note is that validation loss is lower than training loss. A couple of reasons for this irregular behavior. One is that training loss is measured during the epoch and validation loss is calculated at the end of the epoch where the model is slightly trained. Another reason is that the validation set is not sampled properly and this is a sacrifice we made to deal with our resources not able to compute very high numbers of training simulation data. However, we gained confidence when we trained on a certain number of simulations satisfactory for a value close to the true value, then increase simulation vastly when testing the model so there is better convergence to the true value and this strategy works.




Figure 10: Model 2 Loss on American Options

4.1 Real Options Valuation Results

When the model was used with real option valuation, results were similar to an extent. The same neural network model, model 2, loss is presented for comparison reasons with the American option valuation. Other neural network models of valuing other timesteps exhibited similar loss behavior for both real options and American options. Both model loss graphs were on the same neural network architecture and same number of simulations of 100.

Table 4: Real Option Values with Different Simulations

Simulations	Option Value
100	0.15
500	0.13
1000	0.13
2000	0.12




Figure 11: Model 2 Loss on Real Options

One reason this discrepancy exists is because the model is inherit to a single feature data set, which is just the closing price of the stock. Another reason there is a high loss may be due to using an untrained network. Although the network estimated an option value price on the same order of magnitude, the range of error is apparent when applying it to a green energy plant where costs are in the millions of dollars. A pre-trained network may learn the upper and lower bounds better and this expertise is transferred to new inputs. Comparing this model with how the LSM model performed, there is a difference of almost 0.1 in absolute terms between predicted real option values.

Table 5: Comparison of LSM and Neural Network with 100 Simulations

Basis Functions	LSM	Neural Network
$1, S_i, S_i^2, S_i^3$	0.23	0.15
$1, S_i, S_i^2, S_i^3, S_1 S_2$	0.23	0.15
$1, S_i, S_i^2, S_i^3, S_1 S_2, \max(S_1, S_2)$	0.23	0.15
$1, S_i, S_i^2, S_i^3, S_1 S_2, S_1^2 S_2, S_1 S_2^2$	0.22	0.15
$1, S_i, S_i^2, S_i^3, S_1 S_2, S_1^2 S_2, S_1 S_2^2, \tilde{h}(S_1, S_2)$	0.22	0.15
$1, S_i, S_i^2, S_1 S_2, \tilde{h}(S_1, S_2)$	0.24	0.15

Chapter 5. Conclusion

5.1 Summary of Findings

This research has found that untrained neural networks are not efficient but may generate close approximations to valuing real options under specific neural architecture when given stochastic underlying prices. The network generated a sub-optimal exercise decision that had large margin of error when compared to the Least Squared Monte Carlo approach. The network also had high running time and is computationally expensive compared to LSM. Run times for just 2000 simulations took approximately an hour. Given that this is a basic neural network with discrete time-steps, it would be unreasonable to test and customize network parameters if each run was this long. It was also utilizing high amount of RAM at 21 GB with kernel crashes not being uncommon during the experimentation phase.

5.2 Contributions to the Field

Little exploration has been done in the intersect of deep learning and real options, and hopefully this research paves a way for future developments. This research further contributes to deep learning model architecture. We explored potential outcomes and downsides to dividing models in discrete time-steps and training them in that fashion. Our research also implies an efficiency with simulating project trajectories and evaluating the real option that way. This is opposed to estimating fixed underlying project values on discrete times and then finding a real option value. The disadvantage to putting single estimates or even probability with outcomes is small deviations in inputs may vastly change the real option value at the end of calculations. It is thus recommended to follow a simulations approach, then update the model dynamically to incorporate any real time structural change in project value. This could include price shocks to the stochastic variable which is not rare in financial systems. It is also shown that using a simple untrained network will require relatively large computing power compared to LSM and does not improve valuation in our case. We hope that through this general model of valuation, other professionals are able to extract benefit from customizations that are appropriate to their respective uses.

5.3 Practical Implications

As mentioned in the introduction of this paper, real options are a practical tool for practitioners in many fields. This may be applied in regulatory areas to explore the general value creation from certain legislation as seen in Ritzenhofen [13] paper. One may argue that with neural networks, the valuation model is more robust. This entices confidence in current decision makers that rely on traditional methods. Other industries that can manipulate this model to their needs include real estate development, research and development teams, and entrepreneurs. These cases are similar to green energy projects where there is some sort of initial investment that is not much compared to the whole capital needed, then a decision is present to continue or to abandon. Land development, building a drug lab, creating prototypes are the 'premium' these option holders 'pay' to have the real option respectively. This model is not without its downsides. In its general case, it may be unsuitable for applications that have complex cost dependencies. We split projects to variable and fixed costs and assume costs are all inclusive within those variables. When payoff structure increases in complexity, the result of this model is unexplored. Another

implication deals with the nature of the underlying process at hand. Some underlying may not be represented by brownian motion accurately and thus may need another stochastic process to describe it properly. In addition, there must be sufficient computational resources to carry out this valuation process.

5.4 Recommendations for Future Research

One aspect to consider for future research is to test methods suggested by Antonov and Piterbarg in [24]. The authors state that the developed methods in the paper outperform standard deep neural networks in financial applications. It would be interesting to compare the case of real options to the methods they propose. Another aspect is to explore different payoff functions. Exploring different neural network architecture would be helpful to determine optimal starting points for different applications. This research focuses on option valuation for general usage. Future research could incorporate specific use of the model to regulatory or project management specific scenarios. Future research could include studying different underlying processes such as jump diffusion. Future improvements can be made to the network by training it on the one dimensional partial differentiate equation of option valuation and then testing it on green energy power plants.

References

- [1] J. Detemple and Y. Kitapbayev, “The value of green energy: Optimal investment in mutually exclusive projects and operating leverage,” *Review of Financial Studies*, vol. 33, 2020.
- [2] A. K. Dixit and R. S. Pindyck, *Investment under Uncertainty*. Princeton University Press, 1994.
- [3] L. Trigeorgis, *Real options : managerial flexibility and strategy in resource allocation*. MIT Press Cambridge, Mass, 1996.
- [4] J.-P. Décamps, T. Mariotti, and S. Villeneuve, “Irreversible investment in alternative projects,” *Economic Theory*, vol. 28, no. 2, pp. 425–448, 2006.
- [5] K. M. Maribu, A. Galli, and M. Armstrong, “Valuation of spark-spread options with mean reversion and stochastic volatility,” *HAL*, 2007.
- [6] S.-E. Fleten and E. Näsäkkälä, “Gas-fired power plants: Investment timing, operating flexibility and co2 capture,” *Energy Economics*, 2010.
- [7] S.-E. Fleten, K. Maribu, and I. Wangensteen, “Optimal investment strategies in decentralized renewable power generation under uncertainty,” *Energy*, 2007.
- [8] T. K. Boomsma, N. Meade, and S.-E. Fleten, “Renewable energy investments under different support schemes: A real options approach,” *European Journal of Operational Research*, 2012.
- [9] I. Bakke, S.-E. Fleten, L. I. Hagfors, V. Hagspiel, and B. Norheim, “Investment in mutually exclusive transmission projects under policy uncertainty,” *Journal of Commodity Markets*, 2016.
- [10] M. Broadie and J. Detemple, “The valuation of american options on multiple assets,” *Mathematical Finance*, 1997.
- [11] K. Venetsanos, P. Angelopoulou, and T. Tsoutsos, “Renewable energy sources project appraisal under uncertainty: the case of wind energy exploitation within a changing energy market environment,” *Energy Policy*, 2002.
- [12] F. Kjærland, “A real option analysis of investments in hydropower—the case of norway,” *Energy Policy*, vol. 35, no. 11, pp. 5901–5908, 2007.
- [13] R. Ingmar, “Robustness of renewable energy support schemes facing uncertainty and regulatory ambiguity,” *SSRN Electronic Journal*, 2014.
- [14] IEA, “World energy outlook 2022.” IEA, Paris, 2022.
- [15] J. Detemple, “Optimal exercise for derivative securities,” *Annual Review of Financial Economics*, vol. 6, pp. 459–487, 12 2014.
- [16] F. A. Longstaff and E. S. Schwartz, “Valuing american options by simulation: A simple least-squares approach,” *Review of Financial Studies*, vol. 14, 2001.

- [17] S. Becker, P. Cheridito, and A. Jentzen, “Pricing and hedging american-style options with deep learning,” *Journal of Risk and Financial Management*, vol. 13, 2020.
- [18] G. Consigli and M. M. Hosseinzadeh, “Simulation of long-term returns with stochastic correlations,” *International Journal of Finance Managerial Accounting*, 2017.
- [19] P. Glasserman, *Monte Carlo methods in financial engineering*. Springer, 2004.
- [20] J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, and T. Chen, “Recent advances in convolutional neural networks,” *Pattern Recognition*, vol. 77, pp. 354–377, 2018.
- [21] C.-Y. Liou, W.-C. Cheng, J.-W. Liou, and D.-R. Liou, “Autoencoder for words,” *Neurocomputing*, vol. 139, pp. 84–96, 2014.
- [22] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman, “1d convolutional neural networks and applications: A survey,” *Mechanical Systems and Signal Processing*, vol. 151, p. 107398, 2021.
- [23] A. Zeng, M. Chen, L. Zhang, and Q. Xu, “Are transformers effective for time series forecasting?,” *The Thirty-Seventh AAAI Conference on Artificial Intelligence*, 2022.
- [24] V. P. Alexandre Antonov, “Alternatives to deep neural networks in finance,” *SSRN Electronic Journal*, 2021.