

Game contingent claims

J. J. Daniel Eliasson

June 14, 2012

Abstract

Game contingent claims (GCCs), as introduced by Kifer (2000), are a generalisation of American contingent claims where the writer has the opportunity to terminate the contract, and must then pay the intrinsic option value plus a penalty. In complete markets, GCCs are priced using no-arbitrage arguments as the value of a zero-sum stochastic game of the type described in Dynkin (1969). In incomplete markets, the neutral pricing approach of Kallsen and Kühn (2004) can be used.

In Part I of this thesis, we introduce GCCs and their pricing, and also cover some basics of mathematical finance.

In Part II, we present a new algorithm for valuing game contingent claims. This algorithm generalises the least-squares Monte-Carlo method for pricing American options of Longstaff and Schwartz (2001). Convergence proofs are obtained, and the algorithm is tested against certain GCCs. A more efficient algorithm is derived from the first one using the computational complexity analysis technique of Chen and Shen (2003).

The algorithms were found to give good results with reasonable time requirements. Reference implementations of both algorithms are available for download from the author's Github page <https://github.com/del/Game-option-valuation-library>.

Keywords: Game contingent claims, game options, Israeli options, Dynkin games, zero-sum games, non-zero-sum games, Monte-Carlo simulation, pricing

Acknowledgements

I would like to thank my advisor, Harald Lang, for his support and comments during the writing of this piece. I am also indebted to Dr. Peter Ouwehand of Stellenbosch University, who received me as an affiliated student, set me on the course of game options, and advised me during my work. I also thank Henrik Hult and Ali Hamdi at KTH for reading and commenting on my penultimate draft. Their help was most useful, and any errors that remain in this work are entirely mine.

My family deserves thanks for cheering me on from across the world, and later from more close by, and for inspiring me to set off on my studies in the first place.

Finally, there is one person who has supported me through this whole journey, and who has ensured that I stayed on track to finish what I'd started: my loving girlfriend, Stephanie Strydom. To her, I dedicate this work.

Stockholm, June 14, 2012

Daniel Eliasson

Contents

1	Introduction	1
I	The theory of game contingent claims	3
2	Financial markets	5
2.1	Financial market models	5
2.1.1	Risk-free rates	5
2.1.2	Single security financial markets	6
2.1.3	Markets with multiple securities	6
2.2	Contingent claims	7
2.2.1	European options	7
2.2.2	American options	8
2.2.3	Pricing options	8
2.2.4	Complete and incomplete markets	10
3	Game contingent claims	11
3.1	Dynkin games	11
3.2	Game options	13
3.3	Pricing of GCCs	15
3.3.1	Continuous time	16
3.3.2	Discrete time	18
II	Numerical methods and applications	21
4	Algorithm 1: A least-squares Monte-Carlo method	23
4.1	Description of Algorithm 1	24
4.2	Convergence	27
4.2.1	Convergence of V_0^m to V_0	27
4.2.2	Convergence of $V_0^{m,N}$ to V_0^m	29
5	Algorithm 2: A simple Monte-Carlo method	37
5.1	Description of Algorithm 2	37

6	Possible extensions of algorithms 1 and 2	41
6.1	Multiple security markets	41
6.2	Path dependent payoffs	42
7	Examples	43
7.1	Callable puts	43
7.1.1	Results for Algorithm 1	44
7.1.2	Results for Algorithm 2	47
7.2	Convertible bonds	49
7.2.1	Results for Algorithm 1	50
7.2.2	Results for Algorithm 2	51
8	Notes on software used	57
	Bibliography	59

Chapter 1

Introduction

Financial contracts similar to options have existed since ancient times, and stock options were traded on Dutch and English markets as early as the 1600's. Options trading in the modern sense started in 1973, when the Chicago Board Options Exchange (CBOE) was established, and became the first exchange to list standardised options.

Since then, the trade in financial derivatives has grown to become a massive market, with a capitalisation several times larger than the world's gross domestic product (GDP). The outstanding value of over-the-counter financial derivatives alone exceeded \$590 billion¹ in 2008², to compare with a world GDP of \$55 billion³.

Due to their practical importance, the fair pricing of options and other derivatives has been the subject of a large body of research, and there are many practitioners, so called *quants*, in the field of quantitative finance.

While European and American-style options are commonly priced according to models that ignore counterparty risk, in reality, any financial contract has some implicit possibility of premature termination by either of the contract parties, which may then have to pay a penalty for the breach of contract. There is also the risk of one party defaulting on the contract due to insolvency. Finally, some financial instruments already exist where a buyback or transformation option is explicitly stated, as is the case with *callable puts* and *convertible bonds*.

Such financial contracts were formalised in Kifer (2000), who introduced the concept of *game contingent claims* (GCCs), also known as *game options* or *Israeli options*. In a game contingent claim, the holder has the opportunity to exercise his option at any time until a fixed maturity, whilst the writer has the opportunity to terminate the option at any time up until maturity. However, if the writer terminates the claim, he must pay to the holder the

¹American: *trillions*, i.e. 1 billion = 10^{12}

²Data from the Bank for International Settlements, <http://www.bis.org/>.

³According to the International Monetary Fund (IMF), <http://www.imf.org/>.

exercise value of the claim, plus an extra penalty.

GCCs can be seen as a generalisation of American contingent claims, and mathematically they are treated using the theory of optimal stopping which also applies to ACCs. However, due to the two-sided nature of the contracts, the optimal stopping problem is in the form of an optimal stopping game of the type described by Dynkin (1969).

Game contingent claims is a new field of study, and little work has been published on computational methods for them, even though some interesting findings have been made on the theory of such claims. This thesis focuses on the algorithmic aspect, but Part I goes over some basic theory of mathematical finance, and introduces GCCs and their pricing.

In Chapter 2, a brief refresher is given of the concepts of mathematical finance dealing with options that are of relevance here. This chapter also serves to introduce the terminology and notation used in latter chapters.

In Chapter 3, Dynkin games are discussed, and then game contingent claims are introduced, and the original pricing formula derived in Kifer (2000) is presented.

While theory is important, a derivatives trader or financial institution has as their primary concern the practical problem of pricing derivatives. The valuation of realistic game contingent claims requires the use of numerical methods.

Part II of this thesis is concerned with these numerical methods. One previously suggested method is based on simulating a judiciously chosen martingale, as described by Kühn, Kyprianou, and van Schaik (2007). Taking a different approach, in this thesis we develop a Monte-Carlo algorithm, based on an algorithm for American options that was introduced by Longstaff and Schwartz (2001), and further analysed by Clément, Lamberton, and Protter (2002).

In Chapter 4, the algorithm is described and convergence proofs are obtained. Chapter 5 investigates an improvement on the algorithm, using the techniques of Chen and Shen (2003), and in Chapter 7, the two algorithms are tested against some realistic GCCs.

The algorithm we derive has several good qualities. It is conceptually simple to understand, works for any Lévy model for underlyings, can deal with stochastic interest rates, and can be parallelised to run on distributed hardware. In Chapter 6, we explore the possibility of adapting the algorithm to options on multiple underlyings and path-dependent options.

Part I

The theory of game contingent claims

Chapter 2

Financial markets

2.1 Financial market models

2.1.1 Risk-free rates

When making financial decisions, especially risky ones such as investing in stock or options, it is useful to not only consider the possible gains to be acquired from an investment itself, but also to weigh them against those that could be made by investing in other opportunities. In economics and finance, we call this the *opportunity cost* of an investment.

One particularly relevant opportunity cost when considering a risky investment is that of instead investing the money in a safe way, such as buying government bonds. Since the government is guaranteeing the bond, and has the ability to raise taxes to collect funds if needed, these bonds are generally considered a riskless investment¹, and the interest one can earn on them is known as the *risk-free rate*.

Another characteristic of the risk-free rate is that it is known in advance. While no one can tell you what the value of a share will be in a year's time, it is possible to buy a 1-year government bond which will pay a fixed rate, known at the date of purchase.

Formally, we model the risk-free investment with a deterministic process $(B_t)_{t \in [0, T]}$, where

$$B_t = B_0 e^{rt}, \quad (2.1)$$

with r being the risk-free rate. We assume either that $r_t = r$ is constant, or that $(r_t)_{t \in [0, T]}$ is known in advance.

The fact that r_t is known is what allows us to simply consider the discounted processes described shortly, and not have to worry about the randomness of interest rates.

¹Although holders of Greek government bonds might beg to differ.

2.1.2 Single security financial markets

We consider now a market consisting of the risk-free asset (B_t), the bond, and a single process (S_t), representing the price of a stock.

Note, however, that there are many other markets that can be modelled this way, where the process (S_t) could represent the price of a commodity such as pork bellies, the price of electricity in a spot market, a weather variable, etc. In general we can call this process the *underlying*, a name that makes much more sense when contingent claims are in the picture.

The stock process can be modelled with many kinds of underlying dynamics. Common choices are the log-normal processes of the Black-Scholes model, which is described in Black and Scholes (1973), or jump-diffusion processes such as those used in the convertible bond examples of Section 7.2.

As mentioned in Section 2.1.1, we're usually interested in the excess of return over the risk-free rate, the premium we can earn from the risk we take on in an investment, and so to simplify notation and calculations, it is helpful to introduce the *discounted* stock process.

Definition 2.1. Let $(S_t)_{t \in [0, T]}$ be a stochastic process. Then the process $(\tilde{S}_t)_{t \in [0, T]}$, defined by

$$\tilde{S}_t = \frac{S_t}{B_t}, \quad (2.2)$$

is known as the *discounted form* of $(S_t)_{t \in [0, T]}$.

Remark 2.2. Note that the discounted risk-free process $(\tilde{B}_t)_{t \in [0, T]}$ has $\tilde{B}_t = 1$ for all $t \in [0, T]$.

2.1.3 Markets with multiple securities

An obvious extension of the single-security financial market model is one with multiple securities, each modelled by a stochastic process (S_t^i).

Definition 2.3. The multiple security financial market model has a risk-free process $(B_t)_{t \in [0, T]}$ and m risky securities $(S_t^1)_{t \in [0, T]}, \dots, (S_t^m)_{t \in [0, T]}$.

It is handy to name the risk-free process (S_t^0) instead, and to gather all of these processes together in a single stochastic process which takes vector values of dimension $m + 1$.

Definition 2.4. The multiple security financial market can also be modelled as a vector-valued process (S_t) defined by

$$S_t = (S_t^0, S_t^1, S_t^2, \dots, S_t^m), \quad \forall t \in [0, T], \quad (2.3)$$

where $S_t^0 = B_t^0$, the risk-free security.

The discounted process (\tilde{S}_t) corresponding to this is given by

$$\tilde{S}_t = \frac{1}{B_t} (S_t^0, \dots, S_t^m) = (1, \tilde{S}_t^1, \dots, \tilde{S}_t^m). \quad (2.4)$$

2.2 Contingent claims

Contingent claims are financial instruments where payoffs between two (or more) parties are regulated depending on some form of underlying process, hence the word *contingent* in the name. They are also known as *financial derivatives*.

Well known examples of derivatives are futures contracts on different commodities such as rice, pork bellies, orange juice and cattle; interest rate swaps and stock options.

We are here mostly interested in stock options, and more generally, option-like instruments on some form of underlying, which will be assumed to be a stock, but could in general be any time series, e.g. temperatures as used in weather derivatives.

2.2.1 European options

A *European option* is a contract between an issuer, A, and a buyer, B, that gives B the right, but not obligation, to either sell or buy shares at a pre-determined *strike price* on a given *expiration date* (also called the option's *maturity*).

If B has the right to buy the shares, the contract is a *call* option, whereas it is known as a *put* option if B has the right to sell shares.

The payoff of such a contract is easy to compute. Assume that B purchases a call option from A, allowing him to purchase a share at the strike price K upon expiration. Let the price of shares on the stock exchange on expiration be S_T . Assuming that $S_T > K$, B can make a profit by *exercising* his right to buy shares at the lower price K , and then immediately sell them on the market at the price S_T , netting a profit of $S_T - K$. If, however, $S_T < K$, then B will simply not exercise his option (since there is no use in purchasing above market price), and his payoff is 0.

Summing this up, the payoff Y_T from the call option with strike price K when the expiration date market price of the underlying is S , is given by

$$Y_T = \max(S_T - K, 0) = (S_T - K)^+. \quad (2.5)$$

For a put option, which gives B the right to sell shares at the price K , the situation is the other way around. If $S_T > K$, it would be more beneficial to sell at the market price, so B will not exercise his option, and the payoff $Y_T = 0$. If $S_T < K$, then B can buy shares at the market for S_T and immediately resell them at the price K by exercising his option, thus making a profit of $Y_T = K - S_T$. The payoff, thus, is given by

$$Y_T = \max(K - S_T, 0) = (K - S_T)^+. \quad (2.6)$$

The characteristic of a European option is that the option to exercise only exists on the expiration time T . Generalising to other forms of underlyings,

but retaining this exercise characteristics, we can talk of the larger class of *European contingent claims*, or ECCs for short.

2.2.2 American options

An *American* option is a generalisation of the European counterpart, with the difference lying in that an American option can be exercised not just on the expiration time T , but at any point up to that time, i.e. for all $t < T$. The payoff process for an American call option, then, is given by

$$Y_t = (S_t - K)^+, \quad (2.7)$$

and for an American put option, it is

$$Y_t = (K - S_t)^+. \quad (2.8)$$

Again, we can generalise to other forms of underlyings, and will then speak of *American contingent claims* or ACCs.

Bermudan options

Bermudan options are a type of option that are in between European and American, in the sense that exercise is possible on a set number of discrete time points up to maturity. When using computational methods for valuing American options, one must always consider a discretisation in time, and so, formally, in those situations the options being studied are actually Bermudan approximations to American options.

2.2.3 Pricing options

Since an option gives its holder exercisable rights, but comes with no obligations, it is clear that an option must come at a price, which the buyer pays to the issuer upon entering into the option agreement.

Determining what price should be paid for options and other derivatives is an important task for mathematical finance, and is the purpose of the algorithms developed in Part II of this work.

If it were known upon entering into the option agreement what the stock price would be on the expiration date, and each point up to it, it would be possible to say exactly what the holder of an option stands to make, and the price could reasonably be set to this. However, since this is not possible, the closest we can get is the expected value of the option's payoff process (Y_t).

Investors, however, expect to be compensated for taking on risk, and so are generally not willing to pay the full expected value of the option payoff, but some amount less than that. The specific amount would generally be expected to be a function of each investor's risk aversion, which could be

codified in a utility function. This would lead to each investor assigning his own price to an option, which is unsatisfying.

It turns out that under certain conditions, it is possible to determine a unique price for an option. The idea is to consider a situation where we can show that there is only a single price that does not lead to *arbitrage*, risk-free winnings. In such a situation, the market forces of supply and demand would push the option to the arbitrage-free price.

Consider a financial market consisting of a risk-free investment, single stock and a call option on that stock. At time $t = 0$, the stock has the price S_0 , and at the end of the time period, $t = 1$, the stock has either gone up to uS_0 , where $u > 1$, or gone down to dS_0 , where $d < 1$. During this period, the risk-free investment goes up with the interest $r > 0$. The price of a call option at time $t = 0$ is Y_0 , and the strike price is K , such that $uS_0 > K$ and $dS_0 < K$.

An investor takes on a portfolio consisting of selling 1 call option, and buying Δ shares. This portfolio costs $\Delta S_0 - Y_0$.

At time $t = 1$, the portfolio consisting of Δ shares and one sold call option can be in one of two situations:

- (i) The stock went up to uS_0 . The shares are now worth ΔuS_0 , and the holder of the option will exercise, since $uS_0 > K$. The investor must sell one share at the price K , and is left with $\Delta uS_0 - K$.
- (ii) The stock went down to dS_0 . The shares are worth ΔdS_0 , and the holder of the option will not exercise it, since $dS_0 < K$. The investor thus has ΔdS_0 in shares.

We can determine Δ in such a way that situation (i) and (ii) leave the investor in the same financial position. This happens when

$$\Delta = \frac{K}{(u - d)S_0}. \quad (2.9)$$

Now this portfolio has no uncertainty in it anymore, since we know that whichever way the stock goes, we have the same payoff. Thus, the price of this portfolio at $t = 0$ must simply be the amount of money that can be invested in the risk-free investment and which will grow until $t = 1$ to match the value of the investors portfolio at that point.

If it were anything else, it would be possible to sell the cheaper portfolio and buy the more expensive one, making a risk-free profit, called arbitrage, on the transaction. The existence of such a deal would lead to many investors wanting to purchase the cheaper portfolio, and selling the more expensive one. Through market forces, the cheaper portfolio would then increase in price, and the more expensive would decrease, until they were both priced equally and the arbitrage opportunity had ceased to be.

We are thus justified in assuming that arbitrage opportunities do not exist, for if they do, they will certainly not persist.

Thus, $(1 + r)$ times the value of the portfolio at $t = 0$ is equal to the value of the portfolio at $t = 1$, or

$$(1 + r)(\Delta S_0 - Y_0) = \Delta dS_0. \quad (2.10)$$

Inserting (2.9) and solving for Y_0 , we get

$$Y_0 = \frac{1 + r - d}{(1 + r)(u - d)}K. \quad (2.11)$$

The situation depicted above is simplistic, but by choosing u and d judiciously, and moving to a lattice of many time steps, where this pricing equation is carried out repeatedly, starting at the last time point and working back towards $t = 0$, it is possible to get good valuations for options. The resulting model is called the Cox-Ross-Rubinstein model (Cox, Ross, and Rubinstein, 1979).

It can also be shown that as the size of the time steps $\Delta t \rightarrow 0$, the Cox-Ross-Rubinstein model converges to the famous continuous-time Black-Scholes model described in Black and Scholes (1973).

2.2.4 Complete and incomplete markets

The no-arbitrage pricing strategy used in Section 2.2.3 depends on being able to create a portfolio of securities that replicate the contingent claim that is being priced. If the financial market is such that all contingent claims can be created as portfolios of securities, then the market is called *complete*.

Conversely, if there are contingent claims that do not have replicating portfolios, for instance if the market has significant transaction costs or other friction, we say that the market is *incomplete*.

Chapter 3

Game contingent claims

3.1 Dynkin games

Dynkin games is a class of zero-sum *optimal stopping games* which were introduced in Dynkin (1969).¹

Consider a game played between two players, A and B, where each day A and B must let each other know if they want to stop on that day, or continue the game. When either player chooses to stop the game, B will receive some amount of money from A. The specific amount B receives is governed by three stochastic processes: one for the amount B receives if A stops the game first, one for the amount that B receives if B stops the game first, and one for the amount that B receives if both players choose to stop on the same day.

Clearly, in such a game, B will attempt to maximise the amount he receives, while A attempts to minimise the payout.

Mathematically, we consider payoff processes which are generated from some underlying Markov process, i.e. a process that is stochastically static, and let A and B choose stopping times to decide when to end the game.

Definition 3.1. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space equipped with a filtration $(\mathcal{F}_t)_{t \in [0, T]}$. The *Dynkin game* is defined as a game played between players A and B, where A chooses a stopping time $\sigma \in \mathcal{T}_{0, T}$ and B a stopping time $\tau \in \mathcal{T}_{0, T}$. At the time $\sigma \wedge \tau$, A receives the payoff

$$X(S_\sigma)\mathbf{1}_{\{\sigma < \tau\}} + Y(S_\tau)\mathbf{1}_{\{\tau < \sigma\}} + Z(S_\sigma)\mathbf{1}_{\{\sigma = \tau\}}, \quad (3.1)$$

where the *indicator function* $\mathbf{1}_A$ satisfies

$$\mathbf{1}_A = \begin{cases} 1, & \text{if } \omega \in A, \\ 0, & \text{if } \omega \notin A. \end{cases} \quad (3.2)$$

¹Dynkin games have been generalised to a nonzero-sum version as well, but it is not of interest here.

The expected payoff to B is given by

$$M_s(\sigma, \tau) = \mathbb{E}_s [X(S_\sigma)\mathbf{1}_{\{\sigma < \tau\}} + Y(S_\tau)\mathbf{1}_{\{\tau < \sigma\}} + Z(S_\sigma)\mathbf{1}_{\{\sigma = \tau\}}], \quad (3.3)$$

where $X \geq Z \geq Y$ are Borel functions and S is a strong Markov process which begins in $S_0 = s$.

Player A strives to minimise this payoff, while B strives to maximise it.

Remark 3.2. A Dynkin game can also have an infinite horizon $T = \infty$.

Since B receives a payoff from A, it is reasonable that B must pay A some amount of money to entice him to play. Following Ekström and Peskir (2008), we can find the value of a Dynkin game using the notions of Nash and Stackelberg equilibria.

Assuming that both A and B are playing the game in the optimal way, B will be trying to find the stopping time that gives him the highest payoff, under the condition that A has found a stopping time that gives the lowest payout. Conversely, A must assume that B has found an optimal stopping time that gives the highest payoff, and A must then try to find a stopping time that minimises the payoff under those conditions. The strategies of A and B give rise to the *upper* and *lower values* of the Dynkin game.

Definition 3.3. The *upper* and *lower values* of a Dynkin game are defined respectively by

$$\begin{aligned} V^*(s) &\stackrel{\text{def}}{=} \text{ess inf}_\tau \text{ess sup}_\sigma M_s(\sigma, \tau), \\ V_*(s) &\stackrel{\text{def}}{=} \text{ess sup}_\sigma \text{ess inf}_\tau M_s(\sigma, \tau). \end{aligned} \quad (3.4)$$

Definition 3.4. If there exists optimal strategies for A and B that are optimal even when the other player is not cooperating, we have a *Nash equilibrium*. Loosely defined, the Nash equilibrium is a set of strategies such that no player can increase his payoff by changing his strategy, if all other players' strategies remain unchanged. In other words, the Nash equilibrium is a saddle point in the payoff function of each player.

Remark 3.5. Due to the fact that a Nash equilibrium is one where no player can better his payoff while every other player holds his strategy constant, Nash equilibria are also known as *non-cooperative equilibria*. It might in general be possible for players to achieve higher payoffs through cooperation.

Definition 3.6. If a Nash equilibrium does exist in the Dynkin game, then $V^*(s)$ must be equal to $V_*(s)$, due to the fact that the σ and τ that the players find are independent of each other. If this holds, then we are justified in calling it the unique value of the game, and define $V(s) \stackrel{\text{def}}{=} V^*(s) = V_*(s)$.

This value of the game is the fair price that B should pay to A in order to get him to play.

Remark 3.7. If there exist stopping times σ^* , τ^* such that

$$M_s(\sigma, \tau^*) \leq M_s(\sigma^*, \tau^*) \leq M_s(\sigma^*, \tau^*), \quad (3.5)$$

for all $\sigma, \tau \in \mathcal{T}_{[0, T]}$ and for all s , then a Nash equilibrium holds.

Remark 3.8. The special case of a Nash equilibrium from Definition 3.6 is called a *Stackelberg equilibrium*.

We can now present the theorem that states that the Stackelberg equilibrium and unique value of the game exists. In a slightly more restricted form, this was proven already in Dynkin (1969), but the version presented here is due to Ekström and Peskir (2008).

Theorem 3.9. *Consider the Dynkin game in Definition 3.1.*

(i) *If S is a càdlàg process, the Stackelberg equilibrium of Definition 3.6 holds, with $V(s) = V^*(s) = V_*(s)$ being a measurable function.*

(ii) *If S is a càdlàg and quasi-left-continuous process, the Nash equilibrium of Remark 3.7 holds, with*

$$\sigma^* = \inf \{t: S_t \in \{V = X\}\}, \quad \tau^* = \inf \{t: S_t \in \{V = Y\}\}. \quad (3.6)$$

Proof. See Ekström and Peskir (2008), Theorem 2.1. □

For more on the general theory of stochastic processes, including optimal stopping, refer to Nikeghbali (2006).

3.2 Game options

A *game contingent claim* (GCC), introduced in Kifer (2000), is a derivative contract between a seller A, and a buyer B. The claim in question is frequently an option, and we will then call the GCC a *game option*², and the seller A will be known as the *writer*, whilst the buyer B is the *holder*. Similarly to an American contingent claim (ACC), the buyer can exercise the contract at any time until a final timepoint called the *maturity*, but unlike an ACC, in the GCC case the seller too can terminate the contract, at a penalty.

More precisely, consider a financial market as defined in Section 2.1. Let $(X_t)_{t \in [0, T]}$ and $(Y_t)_{t \in [0, T]}$ be adapted càdlàg processes, with $\mathbb{E}|X_t|^2 < \infty$ and $\mathbb{E}|Y_t|^2 < \infty$ for all $t \in [0, T]$. Further, let $X_t \geq Y_t$ for all $t \in [0, T]$. These processes are considered to be the payoff processes such that if A terminates

²Sometimes also known as *Israeli options*, as suggested by Kifer.

the contract at time t , he pays to B the sum X_t , whilst if B exercises at time t , he receives from A the sum Y_t .³

If A chooses to terminate at the same time t as B chooses to exercise, B receives Y_t , i.e. this case is considered to be equal to exercise.⁴

In other words, assuming that A chooses to terminate the option at a stopping time $\sigma \in [0, T]$, and B exercises at $\tau \in [0, T]$, A pays to B an amount $R(\sigma, \tau)$ as defined below.

Definition 3.10. The *payoff* of a GCC that B receives from A is given by

$$R(\sigma, \tau) \stackrel{\text{def}}{=} X_\sigma \mathbf{1}_{\{\sigma < \tau\}} + Y_\tau \mathbf{1}_{\{\tau \leq \sigma\}}. \quad (3.7)$$

Remark 3.11. Note that this is a Dynkin game, as described in Section 3.1.

Assuming that such a price exist, we need a symbol for the *fair price* of the GCC.

Definition 3.12. The *fair price* that B must pay to A at time $t = 0$ for a GCC is called V . The price at any time t in the future is called V_t . Note that $V \stackrel{\text{def}}{=} V_0$.

Remark 3.13. Note that the fair price of Definition 3.12 hasn't been fully specified. The notion of a fair price used depends on the situation. The specific definition used depends on the financial market model, but is usually the lowest price of a hedging strategy, as in arbitrage-free pricing.

If a hedge does not exist, it is under certain circumstances still possible to find a unique fair price for the GCC. Kallsen and Kühn (2004) describe the *neutral pricing* approach. They assume a market wherein participants are expected utility maximisers in a market with balanced derivative supply and demand, and replace the equivalent martingale measure of a complete market with a *neutral pricing measure*. It is shown that under this measure, the fair price of a GCC corresponds again to the value of a Dynkin game.

In this work, we assume that the fair price is the value of the Dynkin game.

In many real situations, a party to a financial contract can get out of his contractual obligation, but will usually then need to pay a penalty for it. The difference between the payoff when B exercises the option and when A terminates it can be considered a penalty that A must pay in order to get out of a contract. This means that GCCs can be used to model contracts where A should not be able to terminate the contract, but realistically can.

³It is possible to generalise this to include a process $(W_t)_{t \in [0, T]}$, $X_t \geq W_t \geq Y_t \geq 0$, $\forall t \in [0, T]$, with $W_T = Y_T$, such that the payoff is W_t if termination and exercise coincide. This, however, does not change the price of the GCC, see Kifer (2000), Remark 2.2.

⁴The GCC could also be defined such that the payoff when exercise and termination coincide is X_t . If $X_T = Y_T$, this does not change the price of the GCC (Kifer, 2000, Remark 2.2).

Remark 3.14. Define $(\delta_t)_{t \in [0, T]}$ by $\delta_t = X_t - Y_t$. Note that $\delta_t \geq 0$ since $X_t \geq Y_t$. The process (δ_t) represents the penalty that A must pay B for terminating the contract. This way, we can write $R(\sigma, \tau) = Y_{\sigma \wedge \tau} + \delta_\sigma \mathbf{1}_{\{\sigma < \tau\}}$, where $a \wedge b = \min(a, b)$.

Game options are an extension of American options, which in turn are extensions of European options. The following remarks show how to consider ACCs and ECCs as part of the GCC framework.

Remark 3.15. If A is not allowed to terminate the claim at any time before maturity, the GCC becomes an ACC. Kifer (2000) points out that the ACC case can be studied as a GCC where it is never optimal for the writer to terminate. This can be achieved, for instance, when $\delta > \sup_{0 \leq t \leq T} \mathbb{E}[Y_t]$.

Remark 3.16. Kifer (2000) remarks that if B is not allowed to exercise until maturity, the GCC becomes a European contingent claim. This case can be considered as a GCC if $Y_t = 0$ for $t < T$ and $Y_T > 0$.

Since the writer of a game option has a possibility to terminate the option, which does not exist for the writer of American options, the value of a GCC must be lower or equal to that of an ACC. How much lower depends on the penalty the writer must pay to terminate the contract. If this penalty becomes zero, then either the holder will exercise the option (if he believes the value in the future will be less than now), or the writer will terminate it (if he believes the value in the future will be higher than now), and so the option will be stopped immediately.

Remark 3.17. If $\delta_0 = 0$, it is optimal for either writer or holder to stop immediately, and the price of the GCC must be Y_0 . As pointed out in Remark 3.15, if δ is big enough, the price of the GCC is equal to that of an ACC. Together, this means that V is an increasing function of the penalty, with $Y_0 \leq V \leq \sup_{0 \leq t \leq T} \mathbb{E}[Y_t]$.

3.3 Pricing of GCCs

A number of results have been derived regarding the pricing and hedging of game contingent claims. In this section we will only cover the pricing of GCCs in complete markets by following the arguments in Kifer (2000), where the unique price of a GCC is derived as the value of a Dynkin game.

For the interested reader, there are a number of other papers that make for a good start in reading up on GCCs.

Kunita and Seko (2004) study fixed-penalty game call and put options in a complete market, and find exercise regions for writer and holder. They show that the writer of a game call option either terminates the claim when the price is equal to the strike price, or not at all. If the underlying pays no dividend, the holder never exercises; if it does pay a dividend, the holder will exercise whenever the price hits a non-increasing exercise boundary.

For game put options, the results are similar, although the holder's exercise region is never empty.

Kallsen and Kühn (2004) describe the *neutral pricing* approach. They assume an incomplete market wherein participants are expected utility maximisers in a market with balanced derivative supply and demand, and replace the equivalent martingale measure of a complete market with a *neutral pricing measure*. It is shown that under this measure, the fair price of a GCC corresponds again to the value of a Dynkin game.

In Kallsen and Kühn (2004), the neutral pricing approach is concerned with a financial market with many speculators in game options. In such a market, similarly to the case with ACCs, it is never optimal to exercise a GCC before maturity; the holder should instead opt to sell it. If instead a market with a single writer and a single holder is studied, the opportunity to sell the claim disappears, and the game aspect of the game option surfaces. Kühn (2004) studies this case from a utility maximisation perspective, where the trading possibilities in the underlying are explicitly considered.

It is known that for both American put options and Russian options, the finite-horizon problem ($t \in [0, T]$) is harder than the infinite-horizon one, where t takes values in $[0, \infty)$. For the latter case, both types of options have closed-form solutions in a Black-Scholes framework, see for instance McKean (1965); Shepp and Shiryaev (1994). Similarly, in the case of game put options and game Russian options, there are closed-form solutions for the perpetual case, which are derived in Kyprianou (2004).

Getting back to pricing GCCs in complete markets, Kifer (2000) considers the continuous time and discrete time cases, and also the special case of the Cox-Ross-Rubinstein model.

3.3.1 Continuous time

Assume the financial market model from Section 2.1 with the underlying driven by the Black-Scholes model mentioned in Section 2.2.3. Since the market is complete, there exists a unique risk-neutral measure \mathbb{Q} , and all expectation values here are meant to be taken with respect to \mathbb{Q} .

Remark 3.18. Note that the condition $\mathbb{E} |X_t|^2 < \infty, \forall t \in [0, T]$, which is part of the definition of GCCs used in Section 3.2, is quite strong. The results of this section still hold if this condition is weakened to

$$\mathbb{E} \left[\sup_{t \in [0, T]} X_t \right] < \infty. \quad (3.8)$$

The seller A of the GCC will seek to minimise his liability to the buyer B by choosing a stopping time σ such that $\mathbb{E} [R(\sigma, \tau)]$ is minimised. Conversely, the buyer is attempting to choose a stopping time τ that maximises this value. Since the players cannot foresee the future, it must hold $\sigma, \tau \in \mathcal{T}_{0, T}$.

This situation leads to a zero-sum Dynkin game. As proven in Dynkin (1969), such a game has a unique value in the sense that

$$\operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t] = \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,T}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t] \quad \text{a.s.} \quad (3.9)$$

Kifer (2000) shows that the value of the Dynkin game is the unique no-arbitrage price of the GCC using hedging arguments.

Definition 3.19. A *hedge* against a GCC is a pair (σ, π) of a stopping time $\sigma \in \mathcal{T}_{0,T}$ and a self-financing portfolio strategy π , such that the value of the portfolio at time $\sigma \wedge t$ is higher than $R(\sigma, t)$ almost surely for each $t \in [0, T]$.

As is typical in option pricing theory, the *fair price of a GCC* is the infimum of positive prices such that there exists a hedge against the GCC with the price as initial endowment.

Definition 3.20. The *value process* of a GCC is the càdlàg process $(V_t)_{t \in [0, T]}$ such that with probability one,

$$\begin{aligned} V_t &= \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,T}} \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t] \\ &= \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,T}} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,T}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t]. \end{aligned} \quad (3.10)$$

Theorem 3.21. *The fair price of a GCC is given by $V \stackrel{\text{def}}{=} V_0$. Furthermore, for each $t \in [0, T]$, the stopping times*

$$\begin{aligned} \sigma_t^* &= \inf \{s \geq t : X_s \leq V_s\} \wedge T, \\ \tau_t^* &= \inf \{s \geq t : Y_s \geq V_s\}, \end{aligned} \quad (3.11)$$

are the unique optimal stopping strategies for the writer and holder, respectively, and it holds that

$$V_t = \mathbb{E} [R(\sigma_t^*, \tau_t^*) | \mathcal{F}_t] \quad \text{a.s.} \quad (3.12)$$

Lastly, there exists a self-financing portfolio strategy π^ such that (σ_0^*, π^*) is a hedge against the GCC with initial endowment V_0 , and this strategy is almost surely unique up to the time $\sigma_0^* \wedge \tau_0^*$.*

Proof. See Kifer (2000), Theorem 3.1. □

The results also hold for the perpetual contingent claim case.

Theorem 3.22. *Let the conditions of Theorem 3.21 be satisfied with $T = \infty$, in particular the condition in Remark 3.18.*

Then the fair price of the perpetual GCC is given by V_0 , where V_t is defined as

$$\begin{aligned} V_t &= \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,\infty}} \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,\infty}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t] \\ &= \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{t,\infty}} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{t,\infty}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_t]. \end{aligned} \quad (3.13)$$

Proof. See Kifer (2000), Proposition 3.3. □

3.3.2 Discrete time

General case

In general, it is not possible to derive closed-form solutions for the price of a GCC, and so numerical methods become necessary. The foundation of these methods is to approximate the continuous (American) exercise property with a discrete one (Bermudan), as described in Section 2.2.2.

Definition 3.23. $\mathcal{T}_{k,T}^{(n)}$ is the subset of $\mathcal{T}_{k,T}$ of stopping times taking values $jn^{-1}T$ for $j = k, k+1, \dots, n$. The value of the discrete Dynkin game when stopping is only allowed in $\mathcal{T}_{k,T}^{(n)}$ is defined by

$$\begin{aligned} V_k^{(n)} &= \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{k,T}^{(n)}} \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{k,T}^{(n)}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_{kn^{-1}T}] \\ &= \operatorname{ess\,sup}_{\tau \in \mathcal{T}_{k,T}^{(n)}} \operatorname{ess\,inf}_{\sigma \in \mathcal{T}_{k,T}^{(n)}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_{kn^{-1}T}]. \end{aligned} \quad (3.14)$$

Remark 3.24. The discrete Dynkin game value satisfies the relation

$$V_k^{(n)} = \min \left(X_{kn^{-1}T}, \max \left(Y_{kn^{-1}T}, \mathbb{E} \left[V_{k+1}^{(n)} \middle| \mathcal{F}_{kn^{-1}T} \right] \right) \right), \quad (3.15)$$

which makes it possible to calculate the fair price approximation $V_0^{(n)}$.

Of course, for this to be useful, the approximation must converge to the correct fair price, which the following theorem states.

Theorem 3.25. *The value of the discrete Dynkin game of Definition 3.23 converges to the value of the continuous Dynkin game in Definition 3.20 as n goes to infinity. In other words,*

$$V = V_0 = \lim_{n \rightarrow \infty} V_0^{(n)}. \quad (3.16)$$

Proof. Refer to Kifer (2000), Proposition 3.2. □

GCCs in Cox-Ross-Rubinstein's model

Consider the discrete financial market of Section 2.1 with the underlying driven by the Cox-Ross-Rubinstein model described in Section 2.2.3. Since the market model is complete, there exists an equivalent martingale measure $\mathbb{Q} = \{p^*, 1 - p^*\}^L$, given by

$$p^* = \frac{r - d}{u - d}. \quad (3.17)$$

The expectations below are taken with respect to this measure.

Definition 3.26. The value process of a GCC, $(V_j)_{j=0,\dots,L}$, is given by

$$\begin{aligned} V_j &= \min_{\sigma \in \mathcal{T}_{j,L}} \max_{\tau \in \mathcal{T}_{j,L}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_j] \\ &= \max_{\sigma \in \mathcal{T}_{j,L}} \min_{\tau \in \mathcal{T}_{j,L}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_j]. \end{aligned} \quad (3.18)$$

Lemma 3.27. *The value process of the GCC can also be derived recursively from the dynamic programming principle*

$$\begin{cases} V_L = Y_L, \\ V_j = \min (X_j, \max (Y_j, \mathbb{E} [V_{j+1} | \mathcal{F}_j])), \end{cases} \quad (3.19)$$

where $j = 0, \dots, L - 1$.

Theorem 3.28. *The fair price of a GCC is given by V_0 from Definition 3.26. Furthermore, for each $j = 0, \dots, L$, the stopping times*

$$\begin{aligned} \sigma_j^* &= \min\{k \geq j : X_k = V_k\} \wedge L, \\ \tau_j^* &= \min\{k \geq j : Y_k = V_k\}, \end{aligned} \quad (3.20)$$

are in $\mathcal{T}_{j,L}$, and satisfy

$$\mathbb{E} [R(\sigma_j^*, \tau) | \mathcal{F}_j] \leq V_j \leq \mathbb{E} [R(\sigma, \tau_j^*) | \mathcal{F}_j], \quad (3.21)$$

for any $\sigma, \tau \in \mathcal{T}_{j,L}$.

Finally, there exists a self-financing portfolio strategy π^* such that (σ_0^*, π^*) is a hedge against the GCC with initial capital V_0 , and this strategy is almost surely unique up to the time $\sigma_0^* \wedge \tau_0^*$.

Proof. See Kifer (2000), Theorem 2.1. □

Part II

Numerical methods and applications

Chapter 4

Algorithm 1: A least-squares Monte-Carlo method

Game contingent claims generalise American contingent claims, and numerical methods for pricing them encounter the same difficulties as methods for ACCs, plus a few complications of their own. When dealing with claims of American or game type, one must solve an optimal stopping problem. Diffusion models for optimal stopping are difficult to solve using classical PDE methods such as finite difference methods. To remedy this problem, Monte-Carlo methods can be employed. The main difficulty encountered when applying these methods to the optimal stopping problem is the evaluation of conditional expectations.

A secondary difficulty in applying Monte-Carlo methods to American option pricing is that the exercise characteristics of an American option are continuous, but a computer can handle only discrete cases. The standard way of approaching the numerical valuation of American options is to approximate the continuously exercisable option by one which is exercisable only at certain discrete times. An option with such exercise characteristics is called *Bermudan*, and the convergence of Bermudan option prices to American has been shown, for instance in Lamberton (2002).¹

For pricing American options, Longstaff and Schwartz (2001) developed a Monte-Carlo algorithm which addresses the problem of evaluating conditional expectations by regressing them on a finite number of functions of the underlying. This method has earned widespread adoption among practitioners due to being simple to implement and efficient for high-dimensional problems. It is also possible to apply parallel computing techniques to it by using a singular value decomposition method to perform the least-squares regression, as described by Choudhury, King, Kumar, and Sabharwal (2008).

Owing to being a Monte-Carlo method, the Longstaff and Schwartz algorithm can also handle some path-dependence in the payoff functions. This is

¹For the corresponding convergence result for GCCs, see Kifer (2000), Proposition 3.2.

limited, however, by the requirement that the underlying is a Markov chain. The convergence of Longstaff and Schwartz's algorithm, along with certain rate of convergence results, was proven by Clément, Lamberton, and Protter (2002).

In this chapter, we will describe an algorithm for pricing game contingent claims which is essentially an extension of the algorithm in Longstaff and Schwartz (2001), and prove the convergence using methods inspired by Clément et al. (2002). However, due to the two actors involved in the optimal stopping problem for a game option, the convergence proofs become quite a bit more cumbersome than in Clément et al. (2002).

4.1 Description of Algorithm 1

As discussed, we will study the discrete optimal stopping problem where the GCC is exercisable and cancellable at discrete times $\{0, \dots, L\}$ only.

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $(\mathcal{F}_j)_{j=0, \dots, L}$. The underlying of the GCC is an adapted Markov chain $(S_j)_{j=0, \dots, L}$ with state space (E, \mathcal{E}) . The discounted² payoff processes $(X_j)_{j=0, \dots, L}$ and $(Y_j)_{j=0, \dots, L}$ are adapted, with $\mathbb{E}|X_j|^2 < \infty$, $\mathbb{E}|Y_j|^2 < \infty$ for $j = 0, \dots, L$.

Recall that the return, representing the payoff for the buyer of the GCC, is defined as $R(\sigma, \tau) = X_\sigma \mathbf{1}_{\{\sigma < \tau\}} + Y_\tau \mathbf{1}_{\{\tau \leq \sigma\}}$, where the stopping times σ, τ are the chosen stopping strategies of the seller and buyer, respectively.

We are looking for the present value of the GCC, given by

$$V_0 = \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{0,L}} \operatorname{ess\,sup}_{\sigma \in \mathcal{T}_{0,L}} \mathbb{E} [R(\sigma, \tau)], \quad (4.1)$$

where $\mathcal{T}_{j,L}$ is the set of all stopping times with values in $\{j, \dots, L\}$.

To obtain this value, one can use the dynamic programming principle to calculate the value for each time $j = 0, \dots, L$, starting at L and working backwards. The value at time j is given by

$$V_j = \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{j,L}} \operatorname{ess\,sup}_{\sigma \in \mathcal{T}_{j,L}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_j], \quad (4.2)$$

and the dynamic programming principle is

$$\begin{cases} V_L = Y_L, \\ V_j = \min(X_j, \max(Y_j, \mathbb{E}[V_{j+1} | \mathcal{F}_j])), \quad j = 0, \dots, L-1. \end{cases} \quad (4.3)$$

Introducing the stopping times

$$\begin{aligned} \sigma_j &= \min \{k \geq j : X_k = V_k\} \wedge L, \\ \tau_j &= \min \{k \geq j : Y_k = V_k\}, \end{aligned} \quad (4.4)$$

²The proofs assume a constant interest rate for simplicity, but hold for any adapted interest rate, i.e. if $r_j \in \mathcal{F}_j$ for $j = 1, \dots, L$.

where $j = 0, \dots, L$, it follows that

$$\begin{aligned} V_0 &= \mathbb{E} [R(\sigma_0, \tau_0)], \\ V_j &= \mathbb{E} [R(\sigma_j, \tau_j) | \mathcal{F}_j], \quad j = 1, \dots, L. \end{aligned} \quad (4.5)$$

The dynamic programming principle can be rewritten in terms of these stopping strategies, as

$$\begin{cases} \sigma_L = L \\ \sigma_j = j \mathbf{1}_{\{X_j \leq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} + \sigma_{j+1} \mathbf{1}_{\{X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}}, \\ \tau_L = L \\ \tau_j = j \mathbf{1}_{\{Y_j \geq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} + \tau_{j+1} \mathbf{1}_{\{Y_j < \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}}, \end{cases} \quad (4.6)$$

where $j = 1, \dots, L - 1$.

Recall that the underlying is a Markov chain. The payoff processes depend on the underlying so that $X_j = f(j, S_j)$, $Y_j = g(j, S_j)$ for some Borel functions $f(j, \cdot)$, $g(j, \cdot)$, and thus it follows that $V_j = W(j, S_j)$ for some function $W(j, \cdot)$, and $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j] = \mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | S_j]$. Assuming that the initial state $S_0 = s$ is deterministic, then so is V_0 .

With the setup of the algorithm in place, we will now use two separate approximations to enable the problem to be tackled numerically. The first approximation consists of replacing the hard to calculate conditional expectation $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]$ with an orthogonal projection onto the space spanned by a finite number of functions of S_j .

For this, we will consider a sequence $(e_k(x))_{k \geq 1}$ of \mathcal{F}_j -measurable functions $e_k: E \rightarrow \mathbb{R}$ that satisfy

A₁: $(e_k(x))_{k \geq 1}$ is a total sequence in $L^2(\sigma(S_j))$ for $j = 1, \dots, L - 1$.

A₂: For $j = 1, \dots, L - 1$ and $m \geq 1$, if $\sum_{k=1}^m \lambda_k e_k(S_j) = 0$ a.s., then $\lambda_k = 0$, for all $k = 0, \dots, m$.

In other words, the sequence is total and linearly independent. Examples of such sequences are Hermite and Laguerre polynomials, which yield good numerical results in tests.

For $j = 1, \dots, L - 1$, we denote by P_j^m the orthogonal projection from $L^2(\Omega)$ onto the vector space generated by $\{e_1(S_j), \dots, e_m(S_j)\}$. We will write

$$P_j^m (R(\sigma_{j+1}^m, \tau_{j+1}^m)) = \alpha_j^m \cdot e^m(S_j), \quad (4.7)$$

where $u \cdot v$ is the Euclidean inner product, and $e^m(S_j)$ is defined as the vector valued function $(e_1(S_j), \dots, e_m(S_j))$. Under **A₂**, $\alpha_j^m \in \mathbb{R}^m$ can be explicitly written as

$$\alpha_j^m = (A_j^m)^{-1} \mathbb{E} [R(\sigma_{j+1}^m, \tau_{j+1}^m) e^m(S_j)], \quad j = 1, \dots, L - 1, \quad (4.8)$$

where the matrix $A_j^m \in \mathbb{R}^{m \times m}$ has coefficients given by

$$(A_j^m)_{1 \leq k, l \leq m} = \mathbb{E} [e_k(S_j) e_l(S_j)] \quad (4.9)$$

Using this notation, we introduce the approximating stopping times σ_j^m, τ_j^m with their dynamic programming principle

$$\begin{cases} \sigma_L^m = L \\ \sigma_j^m = j \mathbf{1}_{\{X_j \leq \alpha_j^m \cdot e^m(S_j)\}} + \sigma_{j+1}^m \mathbf{1}_{\{X_j > \alpha_j^m \cdot e^m(S_j)\}}, \\ \tau_L^m = L \\ \tau_j^m = j \mathbf{1}_{\{Y_j \geq \alpha_j^m \cdot e^m(S_j)\}} + \tau_{j+1}^m \mathbf{1}_{\{Y_j < \alpha_j^m \cdot e^m(S_j)\}}, \end{cases} \quad (4.10)$$

where $j = 1, \dots, L-1$.

With these stopping times, the value function can be approximated as

$$V_0^m = \min \left(X_0, \max(Y_0, \mathbb{E} [R(\sigma_1^m, \tau_1^m)]) \right) \quad (4.11)$$

Recall that $X_0 = f(0, s)$, $Y_0 = g(0, s)$ are deterministic.

Our second approximation is to use a Monte-Carlo method to numerically estimate $\mathbb{E} [R(\sigma_1^m, \tau_1^m)]$. Given the N independent simulated paths $(S_j^{(1)}, \dots, S_j^{(N)})$ of the Markov chain $(S_j)_{j=0, \dots, L}$, we denote the payoff functions, for $j = 0, \dots, L$ and $n = 1, \dots, N$, by $X_j^{(n)} = f(j, S_j^{(n)})$ and $Y_j^{(n)} = g(j, S_j^{(n)})$. We introduce for each path n stopping times $\sigma_j^{m, n, N}, \tau_j^{m, n, N}$ and their associated dynamic programming principle

$$\begin{cases} \sigma_L^{m, n, N} = L \\ \sigma_j^{m, n, N} = j \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j^{m, n, N} \cdot e^m(S_j^{(n)})\}} + \sigma_{j+1}^{m, n, N} \mathbf{1}_{\{X_j^{(n)} > \alpha_j^{m, n, N} \cdot e^m(S_j^{(n)})\}}, \\ \tau_L^{m, n, N} = L \\ \tau_j^{m, n, N} = j \mathbf{1}_{\{Y_j^{(n)} \geq \alpha_j^{m, n, N} \cdot e^m(S_j^{(n)})\}} + \tau_{j+1}^{m, n, N} \mathbf{1}_{\{Y_j^{(n)} < \alpha_j^{m, n, N} \cdot e^m(S_j^{(n)})\}}, \end{cases} \quad (4.12)$$

for $j = 1, \dots, L-1$, and where $\alpha_j^{m, n, N} \in \mathbb{R}^m$, $j = 1, \dots, L-1$ is the least square estimator (LSE)

$$\alpha_j^{m, n, N} = \arg \min_{a \in \mathbb{R}^m} \sum_{n=1}^N \left(R(\sigma_{j+1}^{m, n, N}, \tau_{j+1}^{m, n, N}) - a \cdot e^m(S_j^{(n)}) \right), \quad (4.13)$$

and it is understood that

$$\begin{aligned} R(\sigma_j^{m, n, N}, \tau_j^{m, n, N}) &= X_{\sigma_j^{m, n, N}}^{(n)} \mathbf{1}_{\{\sigma_j^{m, n, N} < \tau_j^{m, n, N}\}} \\ &\quad + Y_{\tau_j^{m, n, N}}^{(n)} \mathbf{1}_{\{\tau_j^{m, n, N} \leq \sigma_j^{m, n, N}\}}. \end{aligned} \quad (4.14)$$

From these stopping times, the value function is estimated as

$$V_0^{m,N} = \min \left(X_0, \max \left(Y_0, \frac{1}{N} \sum_{n=1}^N R \left(\sigma_1^{m,n,N}, \tau_1^{m,n,N} \right) \right) \right), \quad (4.15)$$

again recalling that X_0, Y_0 are deterministic and known.

4.2 Convergence

As described above, the algorithm can be seen as two successive approximations,

1. approximating the conditional expectation $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]$ with an orthogonal projection onto the space spanned by a finite number of functions of S_j , and,
2. approximating $\mathbb{E} [R(\sigma_1^m, \tau_1^m)]$ with a numerical estimate obtained by Monte-Carlo simulation.

We will first prove that as m goes to infinity, the orthogonal projection in approximation 1 converges in probability to the conditional expectation, and thus that the approximated value function converges in probability to the original value function. Second, we will prove that for a fixed number of functions m , the Monte-Carlo simulation's numerical estimate converges almost surely to the value function from approximation 1 as the number of simulated paths N tends to infinity.

Together, this will show that the value function estimator from the algorithm converges to the correct value function, when both the number of functions in the projection and the number of simulated paths go to infinity.

4.2.1 Convergence of V_0^m to V_0

The convergence

$$V_0^m \xrightarrow[m \rightarrow \infty]{P} V_0 \quad (4.16)$$

follows directly from the following theorem.

Theorem 4.1. *Let $\mathbb{P}(X_j = \mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]) = 0$ a.s. and $\mathbb{P}(Y_j = \mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]) = 0$ a.s. Then, for $j = 1, \dots, L$, as $m \rightarrow \infty$,*

$$(i) \quad \sigma_j^m \xrightarrow{P} \sigma_j,$$

$$(ii) \quad \tau_j^m \xrightarrow{P} \tau_j,$$

$$(iii) \quad \mathbb{E} \left[R \left(\sigma_j^m, \tau_j^m \right) \middle| \mathcal{F}_{j-1} \right] \xrightarrow{P} \mathbb{E} [R(\sigma_j, \tau_j) | \mathcal{F}_{j-1}],$$

Proof. We'll proceed by induction over j .

For $j = L$, by definition $\sigma_L^m = L = \sigma$ and $\tau_L = L = \tau$, so that (i) and (ii) hold. It follows

$$\mathbb{E} [R(\sigma_L^m, \tau_L^m) | \mathcal{F}_{L-1}] = \mathbb{E} [Y_L | \mathcal{F}_{L-1}] = \mathbb{E} [R(\sigma_L, \tau_L) | \mathcal{F}_{L-1}], \quad (4.17)$$

which means that (iii) holds. Therefore, the hypothesis holds for $j = L$.

Assuming that the hypothesis holds for $j + 1$, we shall show that it also holds for j ($1 \leq j \leq L - 1$).

Again, we will first prove (i) and (ii), whereupon (iii) will follow. Consider

$$\begin{aligned} \sigma_j^m - \sigma_j &= j \left(\mathbf{1}_{\{X_j \leq \alpha_j^m \cdot e^m(S_j)\}} - \mathbf{1}_{\{X_j \leq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \right) \\ &\quad + \sigma_{j+1}^m \left(\mathbf{1}_{\{X_j > \alpha_j^m \cdot e^m(S_j)\}} - \mathbf{1}_{\{X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \right) \\ &\quad + (\sigma_{j+1}^m - \sigma_{j+1}) \mathbf{1}_{\{X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \end{aligned} \quad (4.18)$$

$$\begin{aligned} \Rightarrow |\sigma_j^m - \sigma_j| &\leq j \left| \mathbf{1}_{\{\mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j] < X_j \leq \alpha_j^m \cdot e^m(S_j)\}} - \mathbf{1}_{\{\alpha_j^m \cdot e^m(S_j) < X_j \leq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \right| \\ &\quad + \sigma_{j+1}^m \left| \mathbf{1}_{\{\mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j] \geq X_j > \alpha_j^m \cdot e^m(S_j)\}} - \mathbf{1}_{\{\alpha_j^m \cdot e^m(S_j) \geq X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \right| \\ &\quad + |\sigma_{j+1}^m - \sigma_{j+1}| \mathbf{1}_{\{X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} \\ &\leq j \mathbf{1}_{\{|X_j - \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]| \leq |\alpha_j^m \cdot e^m(S_j) - \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]|\}} \\ &\quad + \sigma_{j+1}^m \mathbf{1}_{\{|X_j - \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]| \leq |\alpha_j^m \cdot e^m(S_j) - \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]|\}} \\ &\quad + |\sigma_{j+1}^m - \sigma_{j+1}|. \end{aligned} \quad (4.19)$$

Now, by definition,

$$\alpha_j^m \cdot e^m(S_j) = P_j^m (R(\sigma_{j+1}^m, \tau_{j+1}^m)), \quad (4.20)$$

and since P_j^m represents an orthogonal projection onto the space of \mathcal{F}_j -measurable functions,

$$P_j^m (R(\sigma_{j+1}^m, \tau_{j+1}^m)) = P_j^m (\mathbb{E} [R(\sigma_{j+1}^m, \tau_{j+1}^m) | \mathcal{F}_j]). \quad (4.21)$$

Therefore, under \mathbf{A}_1 ,

$$\alpha_j^m \cdot e^m(S_j) \xrightarrow{m \rightarrow \infty} \mathbb{E} [R(\sigma_{j+1}^m, \tau_{j+1}^m) | \mathcal{F}_j], \text{ in } L^2. \quad (4.22)$$

By the induction hypothesis, this converges to $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]$, so for each $\varepsilon > 0$:

$$\begin{aligned} \limsup |\sigma_j^m - \sigma_j| &\leq (j + \sigma_{j+1}^m) \mathbf{1}_{\{|X_j - \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]| \leq \varepsilon\}} \\ &\quad + \limsup |\sigma_{j+1}^m - \sigma_{j+1}|. \end{aligned} \quad (4.23)$$

The second term goes to 0 due to the induction hypothesis, and when ε goes to 0, the first term vanishes, since $\mathbb{P}(X_j = \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]) = 0$ a.s. This yields $\sigma_j^m \xrightarrow[m \rightarrow \infty]{P} \sigma_j$, i.e. (i) holds.

Analogously, it can be shown that $\tau_j^m \xrightarrow[m \rightarrow \infty]{P} \tau_j$, and (ii) holds.

Finally,

$$\mathbb{E}[R(\sigma_{j+1}^m, \tau_{j+1}^m) | \mathcal{F}_j] = \mathbb{E}\left[X_{\sigma_{j+1}^m} \mathbf{1}_{\{\sigma_{j+1}^m < \tau_{j+1}^m\}} + Y_{\tau_{j+1}^m} \mathbf{1}_{\{\tau_{j+1}^m \leq \sigma_{j+1}^m\}} \middle| \mathcal{F}_j\right], \quad (4.24)$$

which converges in probability, as $m \rightarrow \infty$, to

$$\mathbb{E}\left[X_{\sigma_{j+1}} \mathbf{1}_{\{\sigma_{j+1} < \tau_{j+1}\}} + Y_{\tau_{j+1}} \mathbf{1}_{\{\tau_{j+1} \leq \sigma_{j+1}\}} \middle| \mathcal{F}_j\right] = \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j], \quad (4.25)$$

and so (iii) holds as well. Thus, the hypothesis holds for $j = 1, \dots, L$. \square

4.2.2 Convergence of $V_0^{m,N}$ to V_0^m

Notation

To simplify the presentation of the next convergence proof, it is necessary to introduce some notation and make a few remarks.

Under **A₂**, the least squares estimator $\alpha_j^{m,N}$, $j = 1, \dots, L-1$ has the explicit form

$$\alpha_j^{m,N} = \left(A_j^{m,N}\right)^{-1} \frac{1}{N} \sum_{n=1}^N R\left(\sigma_{j+1}^{m,n,N}, \tau_{j+1}^{m,n,N}\right) e^m\left(S_j^{(n)}\right), \quad (4.26)$$

where $A_j^{m,N} \in \mathbb{R}^{m \times m}$ has coefficients given by

$$\left(A_j^{m,N}\right)_{1 \leq k, l \leq m} = \frac{1}{N} \sum_{n=1}^N e_k\left(S_j^{(n)}\right) e_l\left(S_j^{(n)}\right). \quad (4.27)$$

Note that due to the strong law of large numbers, $A_j^{m,N} \rightarrow A_j^m$, $N \rightarrow \infty$ a.s., and so $A_j^{m,N}$ is invertible for large enough N .

We define matrices $\alpha^m = (\alpha_1^m, \dots, \alpha_{L-1}^m)$ and $\alpha^{m,N} = (\alpha_1^{m,N}, \dots, \alpha_{L-1}^{m,N})$, and given parameters $a^m = (a_1^m, \dots, a_{L-1}^m) \in \mathbb{R}^{m^{L-1}}$ and deterministic vectors $x = (x_1, \dots, x_L) \in \mathbb{R}^L$, $y = (y_1, \dots, y_L) \in \mathbb{R}^L$ and $s = (s_1, \dots, s_L) \in E^L$, we define the vector fields $F = (F_1, \dots, F_L)$ and $G = (G_1, \dots, G_L)$ as

$$\begin{cases} F_L(a^m, x, s) = L, \\ F_j(a^m, x, s) = x_j \mathbf{1}_{\{x_j \leq \alpha_j^m \cdot e^m(s_j)\}} + F_{j+1}(a^m, x, s) \mathbf{1}_{\{x_j > \alpha_j^m \cdot e^m(s_j)\}}, \\ G_L(a^m, y, s) = L, \\ G_j(a^m, y, s) = y_j \mathbf{1}_{\{y_j \geq \alpha_j^m \cdot e^m(s_j)\}} + G_{j+1}(a^m, y, s) \mathbf{1}_{\{y_j < \alpha_j^m \cdot e^m(s_j)\}}. \end{cases} \quad (4.28)$$

Introducing, for $j = 1, \dots, L$, the sets

$$B_j = \{x_j > a_j^m \cdot e^m(s_j)\}, \quad (4.29)$$

$$C_j = \{y_j < a_j^m \cdot e^m(s_j)\}, \quad (4.30)$$

and the stopping times

$$\tilde{\sigma}_j = j \mathbf{1}_{B_j^C} + \sum_{i=j+1}^{L-1} i \mathbf{1}_{B_j \dots B_{i-1} B_i^C} + L \mathbf{1}_{B_j \dots B_{L-1}}, \quad (4.31)$$

$$\tilde{\tau}_j = j \mathbf{1}_{C_j^C} + \sum_{i=j+1}^{L-1} i \mathbf{1}_{C_j \dots C_{i-1} C_i^C} + L \mathbf{1}_{C_j \dots C_{L-1}}, \quad (4.32)$$

where $B_i B_j = B_i \cap B_j$.

Now we can write

$$F_j(a^m, x, s) = x_{\tilde{\sigma}_j}, \quad (4.33)$$

$$G_j(a^m, y, s) = y_{\tilde{\tau}_j}. \quad (4.34)$$

Note that $F_j(a^m, X, S), G_j(a^m, Y, S)$ are independent of $(a_1^m, \dots, a_{j-1}^m)$, and

$$F_j(\alpha^m, X, S) = X_{\sigma_j^m}, \quad (4.35)$$

$$F_j(\alpha^{m,N}, X^{(n)}, S^{(n)}) = X_{\sigma_j^{m,n,N}}^{(n)}, \quad (4.36)$$

$$G_j(\alpha^m, Y, S) = Y_{\tau_j^m}, \quad (4.37)$$

$$G_j(\alpha^{m,N}, Y^{(n)}, S^{(n)}) = Y_{\tau_j^{m,n,N}}^{(n)}. \quad (4.38)$$

For $j = 2, \dots, L$, let H_j denote the vector valued function

$$H_j(a^m, x, y, s) = \left(F_j(a^m, x, s) \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}} + G_j(a^m, y, s) \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}} \right) e^m(s_{j-1}), \quad (4.39)$$

and define φ, ψ, ξ as

$$\varphi_j(a^m) = \mathbb{E}[F_j(a^m, X, S)], \quad (4.40)$$

$$\psi_j(a^m) = \mathbb{E}[G_j(a^m, Y, S)], \quad (4.41)$$

$$\xi_j(a^m) = \mathbb{E}[H_j(a^m, X, Y, S)], \quad (4.42)$$

so that

$$\alpha_j^m = (A_j^m)^{-1} \xi_{j+1}(\alpha^m), \quad (4.43)$$

and, similarly, for $j = 1, \dots, L-1$,

$$\alpha_j^{m,N} = (A_j^{m,N})^{-1} \frac{1}{N} \sum_{n=1}^N H_{j+1}(\alpha^{m,N}, X^{(n)}, Y^{(n)}, S^{(n)}). \quad (4.44)$$

After all this new notation, it will please the reader to know that in the coming section, we will drop the superscript m wherever it would not confuse matters, as it is taken to be fixed.

Convergence

Theorem 4.2. Let $\mathbb{P}(X_j = \alpha_j \cdot e(S_j)) = \mathbb{P}(Y_j = \alpha_j \cdot e(S_j)) = 0$ a.s. and $\mathbb{P}\left(X_j^{(n)} = \alpha_j \cdot e\left(S_j^{(n)}\right)\right) = \mathbb{P}\left(Y_j^{(n)} = \alpha_j \cdot e\left(S_j^{(n)}\right)\right) = 0$ a.s. for all $n = 1, \dots, N$. Then, for $j = 1, \dots, L$, as $N \rightarrow \infty$,

(i) $V_0^{m,N} \rightarrow V_0^m$ a.s., and

(ii) $\frac{1}{N} \sum_{n=1}^N R\left(\sigma_j^{m,n,N}, \tau_j^{m,n,N}\right) \rightarrow \mathbb{E}\left[R\left(\sigma_j^m, \tau_j^m\right)\right]$ a.s.

With the notation just introduced, this means that for all $j = 1, \dots, L$, we must prove

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N F_j\left(\alpha^N, X^{(n)}, S^{(n)}\right) = \varphi_j(\alpha), \quad (4.45)$$

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N G_j\left(\alpha^N, Y^{(n)}, S^{(n)}\right) = \psi_j(\alpha). \quad (4.46)$$

To prove this, we shall rely on three lemmas.

Lemma 4.3. For $j = 1, \dots, L-1$,

$$(i) |F_j(a, X, S) - F_j(b, X, S)| \leq \left(\sum_{i=j}^L |X_i| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i - b_i \cdot e(S_i)| \leq |a_i - b_i| |e(S_i)|\}} \right),$$

$$(ii) |G_j(a, Y, S) - G_j(b, Y, S)| \leq \left(\sum_{i=j}^L |Y_i| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|Y_i - b_i \cdot e(S_i)| \leq |a_i - b_i| |e(S_i)|\}} \right).$$

Proof. Let $B_j = \{X_j > a_j \cdot e(S_j)\}$ and $\tilde{B}_j = \{X_j > b_j \cdot e(S_j)\}$. Then,

$$\begin{aligned} F_j(a, X, S) - F_j(b, X, S) &= X_j \left(\mathbf{1}_{B_j^C} - \mathbf{1}_{\tilde{B}_j^C} \right) \\ &+ \sum_{i=j+1}^{L-1} X_i \left(\mathbf{1}_{B_j \dots B_{i-1} B_i^C} - \mathbf{1}_{\tilde{B}_j \dots \tilde{B}_{i-1} \tilde{B}_i^C} \right) \\ &+ X_L \left(\mathbf{1}_{B_j \dots B_{L-1}} - \mathbf{1}_{\tilde{B}_j \dots \tilde{B}_{L-1}} \right). \end{aligned} \quad (4.47)$$

Looking closer at the first two terms, we have

$$\begin{aligned} \left| \mathbf{1}_{B_j^C} - \mathbf{1}_{\tilde{B}_j^C} \right| &= \left| \mathbf{1}_{\{X_j \leq a_j \cdot e(S_j)\}} - \mathbf{1}_{\{X_j \leq b_j \cdot e(S_j)\}} \right| \\ &= \mathbf{1}_{\{b_j \cdot e(S_j) < X_j \leq a_j \cdot e(S_j)\}} + \mathbf{1}_{\{a_j \cdot e(S_j) < X_j \leq b_j \cdot e(S_j)\}} \\ &\leq \mathbf{1}_{\{|X_j - b_j \cdot e(S_j)| \leq |a_j - b_j| |e(S_j)|\}}, \end{aligned} \quad (4.48)$$

and

$$\begin{aligned}
\left| \mathbf{1}_{B_j \dots B_{i-1} B_i^C} - \mathbf{1}_{\tilde{B}_j \dots \tilde{B}_{i-1} \tilde{B}_i^C} \right| &\leq \sum_{k=j}^{i-1} \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| + \left| \mathbf{1}_{B_i^C} - \mathbf{1}_{\tilde{B}_i^C} \right| \\
&= \sum_{k=j}^i \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right|.
\end{aligned} \tag{4.49}$$

Please note that $\left| \mathbf{1}_{B_k^C} - \mathbf{1}_{\tilde{B}_k^C} \right| = \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right|$. It follows that

$$\begin{aligned}
|F_j(a, X, S) - F_j(b, X, S)| &\leq |X_j| \left| \mathbf{1}_{B_j} - \mathbf{1}_{\tilde{B}_j} \right| \\
&\quad + \sum_{i=j+1}^{L-1} \left(|X_i| \sum_{k=j}^i \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \right) + |X_L| \sum_{k=j}^{L-1} \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \\
&= \sum_{i=j}^{L-1} \left(|X_i| \sum_{k=j}^i \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \right) + |X_L| \sum_{k=j}^{L-1} \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \\
&\leq \left(\sum_{i=j}^{L-1} |X_i| \right) \left(\sum_{k=j}^{L-1} \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \right) + |X_L| \sum_{k=j}^{L-1} \left| \mathbf{1}_{B_k} - \mathbf{1}_{\tilde{B}_k} \right| \\
&= \left(\sum_{i=j}^L |X_i| \right) \left(\sum_{i=j}^{L-1} \left| \mathbf{1}_{B_i} - \mathbf{1}_{\tilde{B}_i} \right| \right) \\
&\leq \left(\sum_{i=j}^L |X_i| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i - b_i \cdot e(S_i)| \leq |a_i - b_i| |e(S_i)|\}} \right).
\end{aligned} \tag{4.50}$$

This proves (i). The proof of (ii) is entirely analogous. \square

Definition 4.4. Let $\tilde{\sigma}_j^{n,N}$ and $\tilde{\tau}_j^{n,N}$ be defined for all $n = 1, \dots, N$, as

$$\begin{cases} \tilde{\sigma}_L^{n,N} = L \\ \tilde{\sigma}_j^{n,N} = j \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j^N \cdot e(S_j^{(n)})\}} + \tilde{\sigma}_{j+1}^{n,N} \mathbf{1}_{\{X_j^{(n)} > \alpha_j^N \cdot e(S_j^{(n)})\}}, \end{cases} \tag{4.51}$$

$$\begin{cases} \tilde{\tau}_L^{n,N} = L \\ \tilde{\tau}_j^{n,N} = j \mathbf{1}_{\{Y_j^{(n)} \geq \alpha_j^N \cdot e(S_j^{(n)})\}} + \tilde{\tau}_{j+1}^{n,N} \mathbf{1}_{\{Y_j^{(n)} < \alpha_j^N \cdot e(S_j^{(n)})\}}, \end{cases}$$

Lemma 4.5. Let $\mathbb{P} \left(X_j^{(n)} = \alpha_j \cdot e \left(S_j^{(n)} \right) \right) = \mathbb{P} \left(Y_j^{(n)} = \alpha_j \cdot e \left(S_j^{(n)} \right) \right) = 0$ a.s. for all $n = 1, \dots, N$, and $\alpha_j^N \rightarrow \alpha_j$ as $N \rightarrow \infty$. Then, for all $j = 1, \dots, L$, as $N \rightarrow \infty$,

(i) $\tilde{\sigma}_j^{n,N} \rightarrow \tilde{\sigma}_j$ a.s., and

(ii) $\tilde{\tau}_j^{n,N} \rightarrow \tilde{\tau}_j$ a.s.

Proof. We will proceed by induction on j . For $j = L$, by definition $\tilde{\sigma}_j^{n,N} = L = \tilde{\sigma}_j$, and $\tilde{\tau}_j^{n,N} = L = \tilde{\tau}_j$, so (i) and (ii) hold.

Assuming that (i) and (ii) hold for $j + 1, \dots, L$, it remains to show that it also holds for j . Consider

$$\begin{aligned} \tilde{\sigma}_j^{n,N} - \tilde{\sigma}_j &= j \left(\mathbf{1}_{\{X_j^{(n)} \leq \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j \cdot e(S_j^{(n)})\}} \right) \\ &\quad + \tilde{\sigma}_j^{n,N} \left(\mathbf{1}_{\{X_j^{(n)} > \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} > \alpha_j \cdot e(S_j^{(n)})\}} \right) \\ &\quad + \left| \tilde{\sigma}_j^{n,N} - \tilde{\sigma}_j \right| \mathbf{1}_{\{X_j^{(n)} > \alpha_j \cdot e(S_j^{(n)})\}}. \end{aligned} \quad (4.52)$$

Thus,

$$\begin{aligned} \left| \tilde{\sigma}_j^{n,N} - \tilde{\sigma}_j \right| &\leq j \left| \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j \cdot e(S_j^{(n)})\}} \right| \\ &\quad + \tilde{\sigma}_j^{n,N} \left| \mathbf{1}_{\{X_j^{(n)} > \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} > \alpha_j \cdot e(S_j^{(n)})\}} \right| \\ &\quad + \left| \tilde{\sigma}_{j+1}^{n,N} - \tilde{\sigma}_{j+1} \right| \mathbf{1}_{\{X_j^{(n)} > \alpha_j \cdot e(S_j^{(n)})\}}. \end{aligned} \quad (4.53)$$

Let us consider these three terms separately. For the first, it holds,

$$\begin{aligned} &j \left| \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} \leq \alpha_j \cdot e(S_j^{(n)})\}} \right| \\ &= j \left(\mathbf{1}_{\{\alpha_j \cdot e(S_j^{(n)}) < X_j^{(n)} \leq \alpha_j^N \cdot e(S_j^{(n)})\}} + \mathbf{1}_{\{\alpha_j^N \cdot e(S_j^{(n)}) < X_j^{(n)} \leq \alpha_j \cdot e(S_j^{(n)})\}} \right) \\ &\leq j \mathbf{1}_{\{|X_j^{(n)} - \alpha_j \cdot e(S_j^{(n)})| \leq |\alpha_j^N - \alpha_j| |e(S_j^{(n)})|\}}. \end{aligned} \quad (4.54)$$

Since $\alpha_j^N \rightarrow \alpha_j$ as $N \rightarrow \infty$ and $\mathbb{P} \left(X_j^{(n)} = \alpha_j \cdot e(S_j^{(n)}) \right) = 0$ a.s., this term converges almost surely to zero as $N \rightarrow \infty$.

Consider the second term,

$$\begin{aligned} &\tilde{\sigma}_j^{n,N} \left| \mathbf{1}_{\{X_j^{(n)} > \alpha_j^N \cdot e(S_j^{(n)})\}} - \mathbf{1}_{\{X_j^{(n)} > \alpha_j \cdot e(S_j^{(n)})\}} \right| \\ &= j \left(\mathbf{1}_{\{\alpha_j \cdot e(S_j^{(n)}) > X_j^{(n)} \geq \alpha_j^N \cdot e(S_j^{(n)})\}} + \mathbf{1}_{\{\alpha_j^N \cdot e(S_j^{(n)}) > X_j^{(n)} \geq \alpha_j \cdot e(S_j^{(n)})\}} \right) \\ &\leq j \mathbf{1}_{\{|X_j^{(n)} - \alpha_j \cdot e(S_j^{(n)})| \leq |\alpha_j^N - \alpha_j| |e(S_j^{(n)})|\}}. \end{aligned} \quad (4.55)$$

Just as for the first term, this converges almost surely to zero as $N \rightarrow \infty$.

Finally, the third term converges almost surely to zero as $N \rightarrow \infty$; this follows directly from the induction hypothesis that $\tilde{\sigma}_{j+1}^{n,N} \rightarrow \tilde{\sigma}_{j+1}$ a.s., as $N \rightarrow \infty$.

This shows that as $N \rightarrow \infty$, $\tilde{\sigma}_j^{n,N} \rightarrow \tilde{\sigma}_j$ a.s., which proves (i). The proof of (ii) is analogous. \square

Lemma 4.6. *Let $\mathbb{P}(X_j = \alpha_j \cdot e(S_j)) = \mathbb{P}(Y_j = \alpha_j \cdot e(S_j)) = 0$ a.s. and $\mathbb{P}(X_j^{(n)} = \alpha_j \cdot e(S_j^{(n)})) = \mathbb{P}(Y_j^{(n)} = \alpha_j \cdot e(S_j^{(n)})) = 0$ a.s. for all $n = 1, \dots, N$. Then, for $j = 1, \dots, L$, as $N \rightarrow \infty$, α_j^N converges almost surely to α_j .*

Proof. We will employ induction on j . For $j = L - 1$, we have

$$\alpha_{L-1}^N = (A_{L-1}^N)^{-1} \frac{1}{N} \sum_{n=1}^N Y_L^{(n)} e(S_L^{(n)}), \quad (4.56)$$

$$\alpha_{L-1} = (A_{L-1})^{-1} \mathbb{E}[Y_L e(S_L)]. \quad (4.57)$$

Since $Y_L^{(n)}, n = 1, \dots, N$ are i.i.d., by the strong law of large numbers, $\alpha_{L-1}^N \rightarrow \alpha_{L-1}$ a.s. as $N \rightarrow \infty$.

Now, assuming the hypothesis holds for $j, \dots, L - 1$, we shall show that it also holds for $j - 1$. Consider

$$\alpha_{j-1}^N = (A_{j-1}^N)^{-1} \frac{1}{N} \sum_{n=1}^N H_j(\alpha^N, X^{(n)}, Y^{(n)}, S^{(n)}). \quad (4.58)$$

By the strong law of large numbers, $A_{j-1}^N \rightarrow A_{j-1}$ a.s., so it must be proven that

$$\frac{1}{N} \sum_{n=1}^N H_j(\alpha^N, X^{(n)}, Y^{(n)}, S^{(n)}) \xrightarrow{N \rightarrow \infty} \xi_j(\alpha) \text{ a.s.} \quad (4.59)$$

The strong law of large numbers, the usual suspect, gives the convergence

$$\frac{1}{N} \sum_{n=1}^N H_j(\alpha, X^{(n)}, Y^{(n)}, S^{(n)}) \xrightarrow{N \rightarrow \infty} \xi_j(\alpha) \text{ a.s.}, \quad (4.60)$$

so all that remains is to prove that as $N \rightarrow \infty$,

$$\frac{1}{N} \sum_{n=1}^N \left(H_j(\alpha^N, X^{(n)}, Y^{(n)}, S^{(n)}) - H_j(\alpha, X^{(n)}, Y^{(n)}, S^{(n)}) \right) \rightarrow 0 \text{ a.s.}, \quad (4.61)$$

and we will be done. We will denote by H_N the sum in (4.61). Now consider

$$\begin{aligned}
|H_N| \leq & \frac{1}{N} \sum_{n=1}^N \left| e \left(S_{j-1}^{(n)} \right) \right| \left[\right. \\
& \left| F_j \left(\alpha^N, X^{(n)}, S^{(n)} \right) - F_j \left(\alpha, X^{(n)}, S^{(n)} \right) \right| \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}} \\
& + \left| G_j \left(\alpha^N, Y^{(n)}, S^{(n)} \right) - G_j \left(\alpha, Y^{(n)}, S^{(n)} \right) \right| \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}} \quad (4.62) \\
& + \left| F_j \left(\alpha^N, X^{(n)}, S^{(n)} \right) \right| \left| \mathbf{1}_{\{\tilde{\sigma}_j^{n,N} < \tilde{\tau}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}} \right| \\
& \left. + \left| G_j \left(\alpha^N, Y^{(n)}, S^{(n)} \right) \right| \left| \mathbf{1}_{\{\tilde{\tau}_j^{n,N} \leq \tilde{\sigma}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}} \right| \right].
\end{aligned}$$

Since $\alpha_i^N \rightarrow \alpha_i$ a.s. for $i = j, \dots, L-1$ and $N \rightarrow \infty$ under the induction hypothesis, it follows from Lemma 4.5 that $\tilde{\sigma}_j^{n,N} \rightarrow \tilde{\sigma}_j$ a.s. and $\tilde{\tau}_j^{n,N} \rightarrow \tilde{\tau}_j$ a.s., and thus also $\mathbf{1}_{\{\tilde{\sigma}_j^{n,N} < \tilde{\tau}_j^{n,N}\}} \rightarrow \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}}$ a.s., and $\mathbf{1}_{\{\tilde{\tau}_j^{n,N} \leq \tilde{\sigma}_j^{n,N}\}} \rightarrow \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}}$ a.s. This means that the last two terms of (4.62) converge almost surely to zero as $N \rightarrow \infty$.

Now, using Lemma 4.3, we get

$$\begin{aligned}
|H_N| \leq & \frac{1}{N} \sum_{n=1}^N \left| e \left(S_{j-1}^{(n)} \right) \right| \left[\right. \\
& \left(\sum_{i=j}^L \left| X_i^{(n)} \right| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i^{(n)} - \alpha_i \cdot e(S_i^{(n)})| \leq |\alpha_i^N - \alpha_i| |e(S_i^{(n)})|\}} \right) \\
& + \left(\sum_{i=j}^L \left| Y_i^{(n)} \right| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|Y_i^{(n)} - \alpha_i \cdot e(S_i^{(n)})| \leq |\alpha_i^N - \alpha_i| |e(S_i^{(n)})|\}} \right) \quad (4.63) \\
& + \left| F_j \left(\alpha^N, X^{(n)}, S^{(n)} \right) \right| \left| \mathbf{1}_{\{\tilde{\sigma}_j^{n,N} < \tilde{\tau}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}} \right| \\
& \left. + \left| G_j \left(\alpha^N, Y^{(n)}, S^{(n)} \right) \right| \left| \mathbf{1}_{\{\tilde{\tau}_j^{n,N} \leq \tilde{\sigma}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}} \right| \right].
\end{aligned}$$

Under the induction hypothesis, $\alpha_i^N \rightarrow \alpha_i$ a.s. for $i = j, \dots, L-1$ and

$N \rightarrow \infty$, so for each $\varepsilon > 0$,

$$\begin{aligned}
\limsup |H_N| &\leq \limsup \frac{1}{N} \sum_{n=1}^N \left| e(S_{j-1}^{(n)}) \right| \left[\right. \\
&\quad \left(\sum_{i=j}^L |X_i^{(n)}| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i^{(n)} - \alpha_i \cdot e(S_i^{(n)})| \leq \varepsilon |e(S_i^{(n)})|\}} \right) \\
&\quad + \left(\sum_{i=j}^L |Y_i^{(n)}| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|Y_i^{(n)} - \alpha_i \cdot e(S_i^{(n)})| \leq \varepsilon |e(S_i^{(n)})|\}} \right) \\
&\quad + \left| F_j(\alpha^N, X^{(n)}, S^{(n)}) \right| \left| \mathbf{1}_{\{\tilde{\sigma}_j^{n,N} < \tilde{\tau}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\sigma}_j < \tilde{\tau}_j\}} \right| \\
&\quad \left. + \left| G_j(\alpha^N, Y^{(n)}, S^{(n)}) \right| \left| \mathbf{1}_{\{\tilde{\tau}_j^{n,N} \leq \tilde{\sigma}_j^{n,N}\}} - \mathbf{1}_{\{\tilde{\tau}_j \leq \tilde{\sigma}_j\}} \right| \right]. \tag{4.64}
\end{aligned}$$

As earlier noted, the last two terms converge to zero, and by the strong law of large numbers, the first two terms converge to an expectation value,

$$\begin{aligned}
\limsup |H_N| &\leq \mathbb{E} \left[\left| e(S_{j-1}) \right| \left(\sum_{i=j}^L |X_i| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i - \alpha_i \cdot e(S_i)| \leq \varepsilon |e(S_i)|\}} \right) \right] \\
&\quad + \mathbb{E} \left[\left| e(S_{j-1}) \right| \left(\sum_{i=j}^L |X_i| \right) \left(\sum_{i=j}^{L-1} \mathbf{1}_{\{|X_i - \alpha_i \cdot e(S_i)| \leq \varepsilon |e(S_i)|\}} \right) \right]. \tag{4.65}
\end{aligned}$$

Since $\mathbb{P}(X_j = \alpha_j \cdot e(S_j)) = \mathbb{P}(Y_j = \alpha_j \cdot e(S_j)) = 0$ a.s., this converges to zero when $\varepsilon \rightarrow 0$. \square

The proof of Theorem 4.2 is very similar to that of Lemma 4.6, and so is omitted.

Chapter 5

Algorithm 2: A simple Monte-Carlo method

The algorithm due to Longstaff and Schwartz that underlies Algorithm 1 can be simplified by removing the least-squares approximation of the conditional expectation, and instead using the known continuation values for each path. This approach was taken by Chen and Shen (2003), who showed that it not only reduces the computational complexity, and therefore the running time of the algorithm, but also gives more accurate values than the LSE method.

Using this approach, we have developed Algorithm 2 as a simplification of Algorithm 1, where the LSE step is removed and replaced with the perfect-foresight method.

Since the least-squares approximation of the conditional expectation is replaced by the future knowledge, the convergence proofs of Section 4.2.2 no longer hold in this case. Alternative proofs have not been obtained, leaving Algorithm 2 on less solid foundation than Algorithm 1. However, the convergence is intuitively likely, and it can be seen in Chapter 7 that the algorithm does empirically converge to similar values as Algorithm 1 does, and indeed does so quicker and with less computational effort.

5.1 Description of Algorithm 2

As in Chapter 4, we will consider a discrete financial market consisting of probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $(\mathcal{F}_j)_{j=0, \dots, L}$. The underlying of the GCC is an adapted Markov chain $(S_j)_{j=0, \dots, L}$ with state space (E, \mathcal{E}) . The discounted payoff processes $(X_j)_{j=0, \dots, L}$ and $(Y_j)_{j=0, \dots, L}$ are adapted, with $\mathbb{E}|X_j|^2 < \infty$, $\mathbb{E}|Y_j|^2 < \infty$ for $j = 0, \dots, L$.

We are looking for the present value of the GCC, given by

$$V_0 = \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{0,L}} \operatorname{ess\,sup}_{\sigma \in \mathcal{T}_{0,L}} \mathbb{E} [R(\sigma, \tau)]. \quad (5.1)$$

The value of the GCC at time j can be written as

$$V_j = \operatorname{ess\,inf}_{\tau \in \mathcal{T}_{j,L}} \operatorname{ess\,sup}_{\sigma \in \mathcal{T}_{j,L}} \mathbb{E} [R(\sigma, \tau) | \mathcal{F}_j], \quad (5.2)$$

using the dynamic programming principle

$$\begin{cases} V_L = Y_L, \\ V_j = \min(X_j, \max(Y_j, \mathbb{E}[V_{j+1} | \mathcal{F}_j])), \quad j = 0, \dots, L-1. \end{cases} \quad (5.3)$$

By introducing the stopping times

$$\begin{aligned} \sigma_j &= \min \{k \geq j : X_k = V_k \text{ or } k = L\}, \\ \tau_j &= \min \{k \geq j : Y_k = V_k \text{ or } k = L\}, \end{aligned} \quad (5.4)$$

where $j = 0, \dots, L$, it follows that

$$\begin{aligned} V_0 &= \mathbb{E} [R(\sigma_0, \tau_0)], \\ V_j &= \mathbb{E} [R(\sigma_j, \tau_j) | \mathcal{F}_j], \quad j = 1, \dots, L. \end{aligned} \quad (5.5)$$

In terms of these stopping strategies, the dynamic programming principle becomes

$$\begin{cases} \sigma_L = L \\ \sigma_j = j \mathbf{1}_{\{X_j \leq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} + \sigma_{j+1} \mathbf{1}_{\{X_j > \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}}, \\ \tau_L = L \\ \tau_j = j \mathbf{1}_{\{Y_j \geq \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}} + \tau_{j+1} \mathbf{1}_{\{Y_j < \mathbb{E}[R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]\}}, \end{cases} \quad (5.6)$$

where $j = 1, \dots, L-1$.

In Chapter 4, we describe an algorithm which uses two separate approximations of the stopping strategies and value process:

- (i) an estimate of the conditional expectation $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]$ using a least-squares orthogonal projection onto the space spanned by a finite number of functions of S_j , and
- (ii) a Monte-Carlo simulation to numerically estimate $\mathbb{E} [R(\sigma_1^m, \tau_1^m)]$.

In this chapter, approximation (i) will be left out, and only approximation (ii) will be used. Instead of basing the estimation of $\mathbb{E} [R(\sigma_1^m, \tau_1^m)]$ on the least-squares approximation of $\mathbb{E} [R(\sigma_{j+1}, \tau_{j+1}) | \mathcal{F}_j]$, this conditional expectation will be replaced with the holding values in the simulated paths.

Let $(S_j^{(1)}, \dots, S_j^{(N)})$ be N independent simulated paths of the underlying S_j , $j = 0, \dots, L$, and denote by $X_j^{(n)} = f(j, S_j^{(n)})$, $Y_j^{(n)} = g(j, S_j^{(n)})$ the

payoff functions for $j = 0, \dots, L$ and $n = 1, \dots, N$. For each path n , we introduce the stopping times $\sigma_j^{n,N}$ and $\tau_j^{n,N}$ with their dynamic programming principle

$$\begin{cases} \sigma_L^{n,N} = L \\ \tau_L^{n,N} = L \\ \sigma_j^{n,N} = j \mathbf{1}_{\{X_j^{(n)} \leq R(\sigma_{j+1}^{n,N}, \tau_{j+1}^{n,N})\}} + \sigma_{j+1}^{n,N} \mathbf{1}_{\{X_j^{(n)} > R(\sigma_{j+1}^{n,N}, \tau_{j+1}^{n,N})\}}, \\ \tau_j^{n,N} = j \mathbf{1}_{\{Y_j^{(n)} \geq R(\sigma_{j+1}^{n,N}, \tau_{j+1}^{n,N})\}} + \tau_{j+1}^{n,N} \mathbf{1}_{\{Y_j^{(n)} < R(\sigma_{j+1}^{n,N}, \tau_{j+1}^{n,N})\}}, \end{cases} \quad (5.7)$$

where $R(\sigma_j^{n,N}, \tau_j^{n,N}) = X_{\sigma_j^{n,N}}^{(n)} \mathbf{1}_{\{\sigma_j^{n,N} < \tau_j^{n,N}\}} + Y_{\tau_j^{n,N}}^{(n)} \mathbf{1}_{\{\tau_j^{n,N} \leq \sigma_j^{n,N}\}}$ and the time runs through $j = 1, \dots, L - 1$.

From these stopping times, the value function is estimated as

$$V_0^N = \min \left(X_0, \max \left(Y_0, \frac{1}{N} \sum_{n=1}^N R(\sigma_1^{n,N}, \tau_1^{n,N}) \right) \right), \quad (5.8)$$

again recalling that X_0, Y_0 are deterministic.

Chapter 6

Possible extensions of algorithms 1 and 2

Algorithms 1 and 2, as defined in Chapters 4 and 5 are both defined for single-security financial markets with Markovian underlyings.

An extension to multiple securities is straightforward, and as long as each underlying is Markovian, so is the multiple-security market, and the proofs of Section 4.2 still hold.

Handling non-Markovian path-dependent payoff functions is a more subtle matter. Owing to being Monte-Carlo methods, the algorithms have potential to be useful in this case.

In the following sections, these extensions are discussed, but no proofs will be presented. It must be stressed that we have not implemented and tested any of these extensions. In essence, this chapter consists of our ideas of interesting future possibilities for the algorithms.

6.1 Multiple security markets

Let $(S_j^1)_{j=0,\dots,L}, \dots, (S_j^p)_{j=0,\dots,L}$ represent p distinct security value processes. For each $i = 1, \dots, p$, S^i is an adapted Markov chain with state space (E^i, \mathcal{E}^i) .

The payoff processes depend on the underlyings via some Borel functions f, g that map $\{0, \dots, L\} \times E^1 \times \dots \times E^p \mapsto [0, \infty)$, i.e.

$$\begin{aligned} X_j &= f(j, S_j^1, \dots, S_j^p), \\ Y_j &= g(j, S_j^1, \dots, S_j^p). \end{aligned} \tag{6.1}$$

It follows that the value process of the GCC depends on the underlyings through some function

$$V_j = W(j, S_j^1, \dots, S_j^p). \tag{6.2}$$

In Algorithm 2, this suffices to introduce multiple securities. In Algorithm 1, the situation is trickier. The LSE in (4.8) must be extended to contain functions of all the underlyings, as well as their cross products. This would suggest an exponential growth in the number of basis functions, m .

However, as noted in Longstaff and Schwartz (2001), Section 2.2, it is possible that in actuality, the growth would be only polynomial. Even so, for multiple securities, it is likely that Algorithm 2 will be a great improvement on Algorithm 1, and the use of the latter is simply not called for.

6.2 Path dependent payoffs

To allow for path dependent payoffs, the requirement that X and Y only depend on the current value of the underlying is relaxed, to allow a more general form

$$\begin{aligned} X_j &= f(\{0, \dots, j\}, (S_k^1)_{k=0, \dots, j}, \dots, (S_k^p)_{k=0, \dots, j}), \\ Y_j &= g(\{0, \dots, j\}, (S_k^1)_{k=0, \dots, j}, \dots, (S_k^p)_{k=0, \dots, j}), \end{aligned} \tag{6.3}$$

where f, g are some Borel functions.

With the payoffs thus defined, the convergence proofs in Section 4.2 do not hold anymore. The LSE (4.8) must be extended to include functions of all previous security values. In the multiple-security case, also of cross terms both between the individual securities, as well as between separate timesteps. The dimensionality of this LSE will grow fast enough that it is possible that the LSE approach of Algorithm 1 is entirely untenable under these circumstances.

Algorithm 2 fares better, due to the fact that it works through direct Monte-Carlo simulation of the paths. In this case, the dimensionality of the problem does not grow when path-dependence is allowed, and the algorithm should be able to function still.

Chapter 7

Examples

The legendary computer scientist Donald Knuth famously quipped

Beware of bugs in the above code; I have only proved it correct,
not tried it.

Alas, the rest of us shall have to resort to testing. In this chapter we put our algorithms to the test against two examples of GCCs: callable puts and convertible bonds.

We study the convergence characteristics of the algorithms when the number of simulated paths, number of time steps and number of basis functions in the LSE vary. We find that both algorithms provide good estimates of the GCC values. Algorithm 2 performs similarly to Algorithm 1, but with quicker convergence and superior computational speed.

7.1 Callable puts

Consider the financial market of Section 2.1 with an underlying following the Black-Scholes model described in Black and Scholes (1973).

We introduce the discounted payoff functions

$$Y_t = (K - S_t)^+, \quad t \in [0, T], \quad (7.1)$$

and

$$\begin{cases} X_t = (K - S_t)^+ + \delta, & t \in [0, T) \\ X_T = (K - S_T)^+, \end{cases} \quad (7.2)$$

where $\delta > 0$ is a fixed penalty. Simply put, this is a put that can be called; it will be called a callable put.¹ In Kühn and Kyprianou (2007), the callable put is characterised as a composite exotic option, and the value function is studied.

¹Sorry about that.

Remark 7.1. Note that the callable put is a regular game put option. Callable puts were in use before Kifer coined the term game option.

Of more interest here is Kühn, Kyprianou, and van Schaik (2007), in which a pathwise GCC pricing algorithm is presented, and the callable put is studied as one example. Kühn et al. use Canadisation (Carr, 1998) to obtain a true value approximation of the prices for five callable puts with parameters

$$\kappa = 0.4, r = 0.06, K = 100, T = 0.5, \delta = 5, \quad (7.3)$$

for different initial prices $S_0 = 80, 90, 100, 110, 120$ of the underlying. These values will serve as a useful benchmark against which to compare the values from Algorithms 1 and 2.

To study the convergence of the algorithms for varying input parameters, we used Algorithm 1 to value these options in a series of simulations with different number of time steps L , number of simulated paths N and number of functions m in the projection. A similar series of simulations were performed for Algorithm 2, although without m , of course. The next two sections discuss the results of the convergence study.

7.1.1 Results for Algorithm 1

To test Algorithm 1, we ran a series of simulations with increasing number of simulated price paths, N , discretisation time steps, L , and least-squares basis functions, m . For each set of N , L and m together with initial stock price $S_0 = 80, 90, 100, 110, 120$, the algorithm was run 20 times, and the results averaged to smooth the values.² The value $S_0 = 100$ was dropped from all graphs since the algorithm almost immediately converges there to the true value of 5, which makes for some rather dull graphs.

Letting $N = 3000$ and $L = 400$ be constant, we consider first the characteristics of the algorithm for varying values of m , i.e. varying numbers of basis functions in the least-squares approximation. We see in Figure 7.1 (p. 45) that the performance of the algorithm does not heavily depend on m once past $m = 50$. Consequently, we kept the constant $m = 60$ for the tests of varying number of simulated paths and time steps.

Next, we ran the simulations for a grid of increasing values for the number of sample paths, N , and time steps, L , while holding the number of basis functions m constant at $m = 60$. The resulting graphs show the convergence characteristics of Algorithm 1 as a function of N and L for different S_0 . These results are presented first as a function of N in Figure 7.2, and then as a function of L in Figure 7.3.

²The underlying assumption that the relative errors are normally distributed was verified using the quantile-quantile plots of them against the normal distribution, for each combination of N , m , L and S_0 , as well as the Shapiro-Wilk normality test.




Figure 7.1: The convergence characteristics of Algorithm 1 for callable puts of different S_0 with a varying number of least-squares basis functions m . $N = 3000$ and $L = 400$ are constant. The dashed lines represent the values where V settles for large N and L .




Figure 7.2: The convergence characteristics of Algorithm 1 for callable puts of different S_0 . Here, the option price V is considered a function of N for different levels of L .

Considering first Figure 7.2 (p. 46), we see a “small multiples” plot of the option price V as a function of N , with each panel in the graph showing this function for a given set of L and S_0 .

It is clear that the algorithm converges, as each panel settles towards a value of V when N grows. The convergence is rather quick, with $N = 3000$ being close to the final value in most cases. It can also be seen that the option value is dependent on the number of time steps L , with each refinement bringing V to settle at a lower value.

Figure 7.3 (p. 47) shows the same dataset, but now each panel shows V as a function of L for a given set of N , S_0 .

Again, we see a clear convergence pattern as L increases in each panel, and jumps between panels as larger values of N move the convergence point. It is clear that most movement between panels is for smaller values of N , with little discernable difference after $N = 3000$, which is in line with what we saw in the last graph.

The convergence in L is a bit different from that in N , still improving at $L = 800$, but with slowed pace. Since the algorithm works by finding the optimal exercise strategies in the Bermudan approximation of the option, it is important that this approximation is good.

In sum, we see that while not too many sample paths are needed to get the Monte Carlo estimation to converge, the discrete time approximation of the option converges slower to the continuous time value.

Computationally, since the algorithm’s run time increases linearly in both N and L , time is better spent by limiting the number of sample paths,




Figure 7.3: The convergence characteristics of Algorithm 1 for callable puts of different S_0 . Here, V is considered a function of L for different levels of N .

and using a finer grid in the time discretisation.

To answer the question of whether the algorithm converges to reasonable values, we compare it in Figure 7.4 (p. 48) to the values calculated in Kühn, Kyprianou, and van Schaik (2007). The graph shows the calculated values of V as a function of S_0 , with the values from Kühn et al. marked in the graph with crosses. The alignment is good, but with slightly high values when $S_0 > 100$.

7.1.2 Results for Algorithm 2

For Algorithm 2, we ran a similar set of simulations with increasing number of simulated paths and discretisation time steps. Due to the lack of a least-squares approximation, there is no parameter m for Algorithm 2, and it also runs significantly faster, by a factor of about 6.5. This allowed us to run the N and L values to higher reaches for Algorithm 2.

Figures 7.5 (p. 48) and 7.6 (p. 48) are the corresponding small multiples plots for Algorithm 2.

Considering first Figure 7.5, we see a plot of the option price V as a function of N , with each panel in the graph showing this function for a given set of L and S_0 .

Again we compare the convergence values of the algorithm with those calculated in Kühn, Kyprianou, and van Schaik (2007). Figure 7.7 shows the calculated values of V as a function of S_0 , with the values from Kühn et al. marked in the graph with crosses.

The alignment is again good, with even better fit to the values from




Figure 7.4: The relative error of callable put values for different S_0 as function of N , using Algorithm 1. $L = 400$ and $m = 60$ are kept fixed.




Figure 7.5: The convergence characteristics of Algorithm 2 for callable puts of different S_0 . Here, the option price V is considered a function of N for different levels of L .




Figure 7.6: The convergence characteristics of Algorithm 2 for callable puts of different S_0 . Here, V is considered a function of L for different levels of N .

Kühn, Kyprianou, and van Schaik (2007) than Algorithm 1, except for the case of $S_0 = 80$. Why this algorithm overestimates the value at $S_0 = 80$ is unclear, but it is also possible that the converse is true; that Algorithm 2 is actually closer to the true value than Kühn et al. (2007).

7.2 Convertible bonds

A convertible bond is a bond that can be converted by the holder to a fixed amount of stock at any time up to the maturity. The writer can at any time recall the bond at a fixed price, but the holder then has the opportunity to convert to stock.

The convertible bond can be modelled as a GCC with discounted payoff functions

$$X_t = \max(\gamma S_t, K) \text{ and } Y_t = \gamma S_t, \quad \text{for } t \in [0, T), \quad (7.4)$$

and $X_T = Y_T = \max(1, \gamma S_T)$. $K > 1$ is the recall price, while $0 < \gamma < 1$ is the number of stocks the bond can be converted into. Note that the holder is guaranteed a minimum payment of 1 at the maturity; this is the bond component of the convertible bond.

In Kühn, Kyprianou, and van Schaik (2007), the convertible bond is studied in the case when the underlying is modelled as a continuous dividend paying stock following a jump-diffusion process with non-negative, exponentially distributed jumps. Formally, we have

$$S_t = \exp(\kappa W_t + \mu t + J_t), \quad (7.5)$$




Figure 7.7: The relative error of callable put values for different S_0 as function of N , using Algorithm 2. $L = 400$ is kept fixed.

where W is a standard Brownian motion and J a compound Poisson process with jump intensity $\eta > 0$ and increments following an exponential distribution with parameter $\vartheta > 1$. Further, J is independent of W . The stock pays a continuous dividend $0 < d < r$, so with the drift $\mu = r - d - \kappa^2/2 + \eta/(1 - \vartheta)$, \mathbb{P} is an equivalent martingale measure.

Similarly to Section 7.1, we use the Canadisation true value approximations from Kühn, Kyprianou, and van Schaik (2007) as a guiding value, and run our algorithm using the parameters

$$\begin{aligned} \kappa &= 0.4, \quad r = 0.06, \quad K = 1.3, \quad T = 0.5, \\ d &= 0.02, \quad \gamma = 0.9, \quad \eta = 10, \quad \text{and } \vartheta = 7, \end{aligned} \tag{7.6}$$

while $S_0 = 0.8, 1.0, 1.2, 1.3, 1.4$, as the values of N , L and m are varied. As before, we run 20 simulations per level of N , L , m and S_0 , to smooth away the variation of the Monte-Carlo simulation.

7.2.1 Results for Algorithm 1

Figures 7.8, 7.9 and 7.10 show the convergence characteristics of Algorithm 1's convertible bond values for different values of the parameters L , m and N when $S_0 = 0.8, 1.0, 1.2, 1.3, 1.4$.




Figure 7.8: The convergence characteristics of Algorithm 1 for convertible bonds of different S_0 . Here, V is considered a function of L , with $N = 1000$ and $m = 20$ held constant.

The behaviour for increasing L is graphed in Figure 7.8 (p. 51). We see that the values fluctuate with less than 1 % from 200 to 2000 time steps, so the algorithm here converges very quickly indeed. In the case of $S_0 = 1.4$, a more stable trend towards settling is found.

The number of basis functions, m , is shown in Figure 7.9 (p. 52), and tells roughly the same story. The differences are on the order of a few percent when going from 20 to 200 basis functions, with a jump to higher precision occurring around 120.

When the number of simulated paths N grows, the algorithm quickly settles in to within a few percent of its final solutions, with relative stability from around 4000 to 8000 sample paths; see Figure 7.10 (p. 53).

7.2.2 Results for Algorithm 2

Figures 7.11 and 7.12 show the convergence characteristics of Algorithm 2 when valuing convertible bonds with varying parameters L and N .

In Figure 7.11 (p. 54), the number of time steps L is varied. The algorithm quickly moves towards its settling value, coming within 1 % before




Figure 7.9: The convergence characteristics of Algorithm 1 for convertible bonds of different S_0 . Here, V is considered a function of m , with $N = 1000$ and $L = 100$ held constant. The dashed lines represent the “true values” as estimated by Kühn et al. (2007)

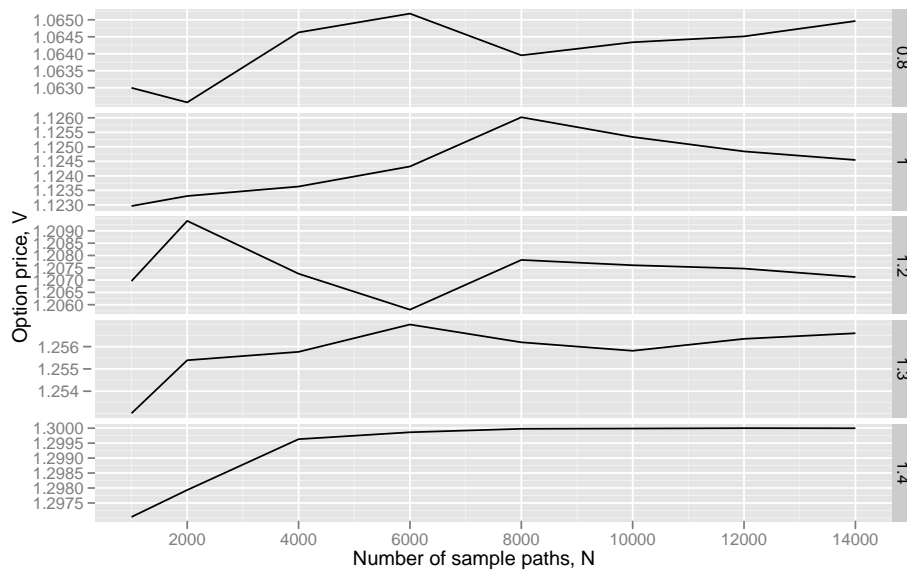


Figure 7.10: The convergence characteristics of Algorithm 1 for convertible bonds of different S_0 . Here, V is considered a function of N , with $L = 100$ and $m = 20$ held constant.

$L = 400$, and then settles slower until changes are minute when $L > 1600$.

As N increases, there is little trending behaviour at all. The algorithm settles almost immediately at a relative error below 1 % and remains there. This can be seen in Figure 7.12 (p. 55).

It is clear that Algorithm 2 has excellent convergence characteristics, settling already at low values of simulated sample paths and time steps. It can be seen, however, that L is the governing parameter in the quality of valuations from Algorithm 2, while even small values of N give good results. As was the case in Section 7.1, Algorithm 2 was also significantly faster than Algorithm 1, in this case a factor about 6.5 when $m = 20$.




Figure 7.11: The convergence characteristics of Algorithm 1 for convertible bonds of different S_0 . Here, V is considered a function of L , with $N = 1000$ held constant.




Figure 7.12: The convergence characteristics of Algorithm 2 for convertible bonds of different S_0 . Here, V is considered a function of N , with $L = 100$ held constant.

Chapter 8

Notes on software used

Both Algorithm 1 and Algorithm 2 were implemented in the *Python* programming language (van Rossum, 1991) with the *NumPy* (Ascher, Dubois, Hinsin, Hugunin, and Oliphant, 1999) and *SciPy* (Jones, Oliphant, Peterson, et al., 2001–) numerical and scientific computing packages. All the source code and documentation of the reference implementation is available for download on the author’s Github page on <https://github.com/del/Game-option-valuation-library>.

The code is designed to make it easy to implement and price general game contingent claims, and should handle gracefully the ECC and ACC cases as well. It is possible to choose between several models for the underlyings, and easy to implement new ones as necessary.

All graphs in this thesis were created using the *R* statistical computing language (R Development Core Team, 2011), together with the *ggplot2* package (Wickham, 2009). This combination was of great use in rapidly testing ideas, visualising data, and creating graphs.

This document was typeset using the X_YTEX dialect of L^AT_EX¹, academia’s favourite typesetting software.

I am deeply grateful for the hard work put into all this free software by the many contributors. Without them, this thesis would doubtlessly have suffered.

¹<http://ctan.org/>

Bibliography

- David Ascher, Paul F. Dubois, Konrad Hinsien, James Hugunin, and Travis Oliphant. *Numerical Python*. Lawrence Livermore National Laboratory, Livermore, CA, ucrl-ma-128569 edition, 1999.
- Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *The Journal of Political Economy*, 81(3):637–654, 1973.
- Peter Carr. Randomization and the american put. *The Review of Financial Studies*, 11(3):597–626, 1998.
- A. S. Chen and P. F. Shen. Computational complexity analysis of least-squares monte carlo (lsm) for pricing us derivatives. *Applied Economics Letters*, 10(4):223–229, March 2003.
- Anamitra R. Choudhury, Alan King, Sunil Kumar, and Yogish Sabharwal. Optimizations in financial engineering: The least-squares monte carlo method of longstaff and schwartz. In *IPDPS - IEEE International Parallel and Distributed Processing Symposium*, 2008.
- Emmanuelle Clément, Damien Lamberton, and Philip Protter. An analysis of a least squares regression method for american option pricing. *Finance and Stochastics*, 6:449–471, 2002.
- John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: a simplified approach. *Journal of Financial Economics*, 7:229–263, 1979.
- E. B. Dynkin. Game variant of a problem on optimal stopping. *Soviet Mathematics Doklady*, 10(2):270–274, 1969.
- Erik Ekström and Goran Peskir. Optimal stopping games for markov processes. *SIAM Journal on Control and Optimization*, 47(2):684–702, 2008.
- Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL <http://www.scipy.org/>.
- Jan Kallsen and Christoph Kühn. Pricing derivatives of american and game type in incomplete markets. *Finance and Stochastics*, 8:261–284, 2004.

- Yuri Kifer. Game options. *Finance and Stochastics*, 4:443–463, 2000.
- Christoph Kühn. Game contingent claims in complete and incomplete markets. *Journal of Mathematical Economics*, 40:889–902, 2004.
- Christoph Kühn and Andreas E. Kyprianou. Callable puts as composite exotic options. *Mathematical Finance*, 17(4):487–502, October 2007.
- Christoph Kühn, Andreas E. Kyprianou, and K. van Schaik. Pricing israeli options: a pathwise approach. *Stochastics*, 79:117–137, 2007.
- Hiroshi Kunita and Susumu Seko. Game call options and their exercise regions. Technical report, Nanzan Academic Society Mathematical Sciences and Information Engineering, November 2004.
- Andreas E. Kyprianou. Some calculations for israeli options. *Finance and Stochastics*, 8:73–86, 2004.
- Damien Lamberton. Brownian optimal stopping and random walks. *Applied Mathematics and Optimization*, 45(3):283–324, July 2002.
- 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.
- H. McKean. Appendix: A free boundary problem for the heat equation arising from a problem of mathematical economics. *Industrial Management Review*, 6:32–39, 1965.
- Ashkan Nikeghbali. An essay on the general theory of stochastic processes. *Probability Surveys*, 3:345–412, 2006.
- R Development Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2011. URL <http://www.R-project.org>. ISBN 3-900051-07-0.
- L. A. Shepp and A. N. Shiryaev. A new look at pricing of the "russian option". *Theory of Probability and its Applications*, 39(1):103–119, 1994.
- Guido van Rossum. Python programming language – official website, 1991. URL <http://www.python.org/>.
- Hadley Wickham. *ggplot2: elegant graphics for data analysis*. Springer New York, 2009. ISBN 978-0-387-98140-6. URL <http://had.co.nz/ggplot2/book>.