Investment Decisions Under Uncertainties:

A Case of Nuclear Power Plants.

Shashi Jain

Investment Decisions Under Uncertainties: A Case of Nuclear Power Plants.

PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. ir. K.C.A.M. Luyben,
voorzitter van het College voor Promoties, in het openbaar te verdedigen op
maandag 10 februari 2014 om 10:00 uur

door

Shashi Jain
MTech. Mechanical Engineering,
Indian Institute of Technology – Madras, India

geboren te Ranchi, India

Dit proefschrift is goedgekeurd door de promotor:
Prof. dr. ir. C. W. Oosterlee

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. C. W. Oosterlee	Technische Universiteit Delft, promotor
Prof. dr. C. Vázquez Cendón,	Univeridade da Coruña
Prof. dr. M. E. Ricotti,	Politecnico di Milano
Prof. dr. ir. A.W. Heemink,	Technische Universiteit Delft
Prof. dr. ir. M. H. Vellekoop,	Universiteit van Amsterdam
Prof. dr. ir. H. X. Lin	Technische Universiteit Delft, reservelid
Dr. ir. J. L. Kloosterman,	Technische Universiteit Delft
Ir. F. Roelofs,	Nuclear Research and Consultancy Group (NRG)

Ir. F. Roelofs heeft als begeleider in belangrijke mate aan de totstandkoming van het proefschrift bijgedragen.

Investment under uncertainties: A case of nuclear power plants. Dissertation at Technische Universiteit Delft.

The work described in this thesis was financially supported by the Nuclear Research Consultancy Group.

ISBN my_ISBN_number

To my parents

Acknowledgements

There are several people who have directly or indirectly contributed to this thesis, and here I would like to acknowledge them.

Foremost, I would like to thank Prof. Kees Oosterlee for providing me this opportunity. His guidance has been, at every stage of the research, a constant source of motivation and knowledge. I benefited immensely from his experience on not just topics related to the research, but also on broader aspects of life. I could not have asked for a better mentor and guide.

I would like to thank Ferry Roelofs for taking over as my co-supervisor. His clear and objective explanations on various aspects of the nuclear industry helped me in my research. He played a key role in identifying important research questions that this thesis addresses. I would also like to thank Aliki van Heek, who was instrumental in formulating the project and advised me during the initial phase of the project. I would like to thank Luc van den Durpel, who introduced me to interesting practical problems that are relevant to the nuclear industry.

I would like to thank Marjon Ruijter, Lech Grzelak, Patrik Karlsson, Yanbin Shen, Qian Feng, Fei Cong, Alvaro Leitao, Suzanne de Jong with whom I had several useful discussions that influenced this research.

It was fun working at CWI mostly because of all the wonderful colleagues I had. I would specially like to thank Janis Bajars, Christoph Koehn, Aram Markosyan, Yunus Hassan, Benjamin Sanderse, Bram van Es, Willem Haverkort, Maria Suarez, Marta Pou, Bin Chen, Bowen Zhang, Anton van der Stoep, Wagner Fortes, Folkert Bleichrodt, Daniel Pelt, Linda Plantagie, Jeroen Witteveen, Joost Batenburg, Margreet Nool, Daan Crommelin, Jason Frank, Tristan van Leeuwen, Debarati Bhaumik, Nick Verheul, Jesse Dorrestijn and Paul De Zeeuw. Special thanks to Keith Myerscough and Wander Wadman for the beautiful Dutch translation of the summary. The kind help of CWI staff, especially Nada Mitrovic, Dubravka Tepsic, Henk Roose and Minnie Middelberg, is much appreciated. My sincere gratitude to Prof. Barry Koren whose advice I really value.

I would like to thank colleagues at NRG — Ed Komen, Vinay Ramohalli, Pratap

Sathiah, Sathosh Jayaraju, Afaque Shams, Steven van Haren, Aditya Thallamthattai and Varun Jain — with whom I had several friendly discussions. My sincere thanks to colleagues at TU Delft. I would like to acknowledge Cindy Bosman and Deborah Dongor for all their help.

I would like to thank my parents, sister and brother for their constant support and encouragement. Most of all I would like to thank my wife Archana, whose constant patience, love, affection and motivation were of immense help in the successful completion of this thesis. Words aren't enough to express my sincere appreciation for her.

Summary

Investment Decisions Under Uncertainties: A Case of Nuclear Power Plants.

Shashi Jain

This thesis discusses the role of *flexibility of decisions* when investing in projects that are affected by economic uncertainties. The theory of real options is extensively applied in this thesis to value such investment decisions. Many investment decisions can be considered as real options, as the investment opportunity gives the right — but not the obligation — to undertake certain business initiatives, such as deferring, abandoning or expanding a capital investment project. Real option theory provides a framework to value the flexibility of decisions which cannot easily be evaluated using the traditional discounted cash flow (DCF) analysis. The underlying numerical techniques used can be challenging as they involve finding an optimal strategy, amongst several possibilities, for making decisions. In a liberalized market, the input costs, the output costs, time duration of projects as well as the discount rates involved can be stochastic which gives rise to multiple sources of uncertainty when investing in real assets. Therefore, the associated numerical problems to be solved often suffer from the so-called *curse-of-dimensionality*.

Valuing real options has similarity with pricing of financial options, especially American options. This thesis develops efficient pricing methods for American options to value real options. Although the pricing methods developed in the thesis are quite general and can be applied to a wide range of investment problems, the focus of the thesis is on evaluating investment decisions related to nuclear power plants.

Although the prospects for nuclear power after the Fukushima accident are weaker in some regions, globally, the nuclear power capacity is projected to rise in the *New Policies Scenario* from 393 GW in 2009 to 630 GW in 2035 [39]. Most of these new reactors are planned in non-OECD countries, for example India has seven reactors under construction and has twenty new units planned; China aims to at least

quadruple its nuclear capacity by 2020; and Russia has about 14 reactors planned, some to replace existing plants, and by 2017 ten new reactors totalling at least 9.2 GWe should be operational [88].

The economics of nuclear power plants suffers from several sources of uncertainty which makes their valuation using traditional methods difficult. On one hand, the revenues from nuclear power plants can be uncertain due to fluctuating electricity prices, while on the other hand in a deregulated market prudently incurred unforeseen capital costs can no longer be passed on to the end users. Generation IV nuclear power plants are considered promising because their conceptual designs include improved safety features and results in better fuel management. Some of the Generation IV reactors can additionally benefit as they are small and medium sized reactors (SMRs). SMRs provide greater flexibility in reactor order decisions as they can be constructed in a sequential modular fashion. Sequential construction allows for the possibility to delay or abandon the planned construction of future modules if the economic conditions do not evolve favourably. The conventional Gen III and Gen III+ reactors are designed to be large units, which enables them to benefit from the economy of scale. An important question then is, how does the economy of scale for the conventional large reactors compare with the flexibility offered by SMRs. There is a need for advanced valuation methods for nuclear power plants that do not only take into account uncertainties in their costs and revenues stream, but also — values the flexibility — and provides an optimal strategy for investment decisions.

The first half of the thesis is concerned with the development of efficient Monte Carlo based numerical methods for pricing American options. As the underlying problems for pricing real option problems and financial American options are similar, the method developed here can be applied to both problems. Specifically, a new pricing method called the stochastic grid method (SGM) is discussed in Chapter 2. Although SGM suffers from drawbacks when pricing high-dimensional options, it serves as an important first step towards the development of a more robust and efficient pricing method called the stochastic grid bundling method (SGBM). Chapter 3 describes in detail SGBM and provides proofs related to the convergence and bias of the results. The chapter also illustrates the method and demonstrates its efficiency by considering different option types, such as basket options with upto fifteen assets.

The second half of the thesis exploits the pricing methods developed in the first half to evaluate investment decisions related to nuclear power plants. Chapter 4 develops the real option model for valuing modular constructions in finite decision time horizon. Valuing modular investments for nuclear power plants is a pertinent research question, owing to an increased focus of the industry and policy makers towards the benefits of small and medium sized reactors (SMRs) when compared to large units. The findings of this chapter are further employed in Chapter 5 to evaluate more detailed and realistic investment decisions, such as benefits of constructing twin reactors, or effects of learning and uncertain lifetime of operation on investment decisions.

Uncertainty in capital costs of nuclear power plants is one of the hurdles for the success of the nuclear industry and calls for an improved decision making process that can take into account these uncertainties. Chapter 6 deals with mitigating the risks involved in the costs and revenues from nuclear power plants using diversification. The outcome of the work described in this chapter is a decision making tool that helps in determining a portfolio of nuclear reactors which has the minimum risk (in terms of the variance of returns) for a given level of expected returns. The decision tool takes into account uncertainties in construction costs, construction duration, operation and maintenance costs, as well as uncertainties in the price of electricity. The decision tool also provides, under the model assumptions, optimal economic conditions to start the construction of different types of reactors, as well as the economic conditions under which the construction of a nuclear reactor should ideally be abandoned.

The stochastic simulation based method developed in the thesis is shown to efficiently price high-dimensional options, and under certain conditions also compute their sensitivities. The thesis also demonstrates how the developed pricing method can be used when valuing real options — particularly the option to delay, the option to abandon and sequential modular options. Finally, the thesis proposes a decision tool that provides, under model assumptions, an optimal policy to invest in different types of nuclear reactors. The decision tool also helps in deciding a reactor order fraction for nuclear power plants that minimizes, based on mean-variance portfolio optimization techniques, the risk that arises from their uncertain construction costs and revenues.

Samenvatting

Investeringsbeslissingen onder onzekere omstandigheden: een casus over kerncentrales.

Shashi Jain

Dit proefschrift behandelt de rol van *beslissingsflexibiliteit* bij het investeren in projecten die beïnvloed worden door economische onzekerheden. De theorie van reële opties wordt uitvoerig toegepast om de waarde van zulke investeringsbeslissingen te bepalen. Veel investeringen kunnen worden beschouwd als reële opties, omdat de investering het recht — maar niet de verplichting — geeft om bepaalde handelsinitiatieven te ondernemen, zoals het starten, stoppen, of uitbreiden van een kapitaalinvestering. De theorie van reële opties creëert een kader voor de waardebepaling van beslissingsflexibiliteit die niet gemakkelijk kan worden geëvalueerd met traditionele *verdisconteerde geldstromenanalyse*. De gebruikte onderliggende numerieke technieken zijn niet-triviaal, omdat zij uit verschillende mogelijkheden een optimale strategie voor het nemen van beslissingen moeten selecteren. In een geliberaliseerde markt kunnen de opstartkosten, de opbrengsten, de tijdsduur van het project en discontovoet stochastisch zijn, wat leidt tot meerdere bronnen van onzekerheid bij het investeren. Daarom lijden de bijbehorende numerieke problemen onder de zogeheten *vloek der dimensionaliteit*.

Waardebepaling van reële opties vertoont een gelijkenis met het prijzen van financiële opties, in het bijzonder van Amerikaanse opties. In dit proefschrift worden efficiënte methoden ontwikkeld voor het prijzen van Amerikaanse opties. De waarderingstechnieken die worden ontwikkeld zijn vrij algemeen en zijn toepasbaar op een breed scala aan investeringsproblemen. Desalniettemin ligt de nadruk hier op het evalueren van investeringsbeslissingen die gerelateerd zijn aan kerncentrales. Ondanks het feit dat na het Fukushima-ongeluk de vooruitzichten voor kernenergie in sommige delen van de wereld zijn verzwakt is de projectie in *New Policies Scenario* dat de capaciteit van kernenergie wereldwijd zal stijgen van 393 GW in 2009 tot 630 GW in 2035 [39]. De meeste van deze nieuwe reactoren staan gepland

in niet-OECD landen. India heeft zeven reactoren in aanbouw en twintig meer gepland; China streeft ernaar haar capaciteit aan kerncapaciteit minimaal te verviervoudigen vóór 2020; en Rusland heeft ongeveer veertien reactoren gepland: enkele om huidige installaties te vervangen en vóór 2017 moeten tien nieuwe reactoren met tezamen 9.2 GWe operationeel zijn [88].

De economische waarde van kerncentrales is onderhevig aan een aantal onzekere factoren, die de waardebepaling met traditionele methoden bemoeilijkt. Enerzijds zal de opbrengst van de centrale onzeker zijn door fluctuerende elektriciteitsprijzen; anderzijds kunnen de gemaakte kosten in een gereguleerde markt niet langer worden doorberekend aan de eindgebruiker. Generatie IV kerncentrales zijn veelbelovend omdat hun conceptuele ontwerp betere veiligheidsmaatregelen bevat en ze resulteren in beter brandstofbeheer. Sommige Generatie IV reactoren zijn *kleine en middelgrote reactoren* (SMRs) en hebben daarom nog meer voordelen. SMRs bieden namelijk grotere flexibiliteit in reactorbestellingsbeslissingen omdat zij op een sequentiële en modulaire wijze gebouwd kunnen worden. Sequentiële constructie geeft de mogelijkheid tot uitstel of afstel van de geplande constructie van verdere modules als de economische randvoorwaarden zich onfortuinlijk ontwikkelen. Anderzijds profiteren de Gen III en Gen III+ reactoren vanwege hun formaat van schaalvergroting. Een belangrijke vraag is dan hoe de schaalvergroting voor conventionele reactoren opweegt tegen de flexibiliteit van SMRs. Er is behoefte aan geavanceerde waarderingsmethoden voor kerncentrales die niet alleen de onzekerheden in de ontwikkeling van kosten en baten in ogenschouw nemen, maar ook de flexibiliteit waarderen en een optimale strategie voor investeringsbeslissingen geven.

De eerste helft van dit proefschrift behandelt de ontwikkeling van efficiënte numerieke technieken voor het prijzen van reële opties op basis van Monte-Carlosimulaties. Omdat de onderliggende problemen voor het prijzen van reële opties en financiële Amerikaans opties sterk overeenkomen, is de hier ontwikkelde techniek toepasbaar op beide problemen. Specifiek wordt een nieuwe waarderingstechniek, de Stochastic Grid Method (SGM), voorgesteld in Hoofdstuk 2. Ondanks het feit dat SGM lijdt onder tekortkomingen bij het prijzen van hoog-dimensionele opties, dient ze als een belangrijke eerste stap richting de ontwikkeling van een robuustere en efficiëntere waarderingstechniek, de Stochastic Grid Bundeling Method (SGBM) genaamd. Hoofdstuk 3 beschrijft SGBM in detail en bevat bewijzen omtrent de convergentie en zuiverheid van de resultaten. Het hoofdstuk illustreert ook de efficiëntie van de techniek door het toe te passen op verschillende optietypes, zoals opties op een mandje met tot vijftien aandelen.

Het tweede deel van dit proefschrift gebruikt de waarderingstechnieken van het eerste deel voor het evalueren van de investeringsbeslissingen omtrent kerncentrales. Hoofdstuk 4 ontwikkelt een model met reële opties voor de waardering van modulaire constructies met een eindige tijdshorizon voor de beslissing. De waardering van modulaire constructies is, als gevolg van een toegenomen focus van de industrie en beleidsmakers op de voordelen van SMRs in vergelijking met grotere centrales, een prangende onderzoeksvraag. De bevindingen van dit hoofdstuk wor-

den verder gebruikt in Hoofdstuk 5 voor het evalueren van realistischere investeringsbeslissingen. Voorbeelden hiervan zijn de voordelen van de bouw van dubbele reactoren of de effecten van voortschrijdend inzicht en onzekere levensduur op investeringsbeslissingen.

Onzekerheid in de kapitaalkosten van kerncentrales is een groot struikelblok voor de nucleaire industrie en vraagt om een verbeterd beslissingsproces dat deze onzekerheid in beschouwing neemt. Hoofdstuk 6 behandelt het verdelen van risicos in de kosten en baten van kerncentrales door gebruik te maken van diversificatie. Het resultaat van dit werk is een techniek die helpt bij het bepalen van een portfolio van kernreactoren met een minimaal risico (in termen van de variantie in de opbrengst) bij een gegeven winstverwachting. De techniek neemt onzekerheden omtrent constructiekosten, constructieduur, operationele kosten en onderhoudskosten in acht, alsmede onzekerheden in de elektriciteitsprijs. De techniek levert ook, onder modelaannames, de optimale economische randvoorwaarden voor de start van de bouw van verschillende reactortypes en ook de economische randvoorwaarden waaronder de bouw idealiter wordt stopgezet.

De op stochastische simulatie gebaseerde rekentechniek in dit proefschrift is efficiënt voor het prijzen van hoogdimensionale opties en onder bepaalde voorwaarden ook voor het berekenen van gevoeligheden van de berekende waarden. Het proefschrift demonstreert hoe de ontwikkelde waarderingstechniek gebruikt kan worden bij de waardering van reële opties — zoals de optie wanneer sequentiële modulaire centrales te bouwen. Tot slot stelt het proefschrift een techniek voor die, onder modelaannames, een optimaal voorschrift voor het investeren in verschillende types kernreactoren levert. De techniek ondersteunt ook beslissingen die het risico als gevolg van onzekere bouwkosten en opbrengsten minimaliseert op basis van variantiereductie optimalizatietechnieken.

Contents

CHAPTER

1 Introduction

The global electricity demand is projected to more than double to over 30,000 TWh annually by the year 2030. More than 70% of the increased energy demand will come from developing countries, led by China and India [49]. Providing sufficient energy to meet the needs of a growing world population without substantially impacting the environment will be a daunting task.

There is currently 370 GWe of nuclear power capacity in operation around the world, producing 3000 TWh each year — 15% of the world's electricity — the largest share provided by any non-greenhouse gas emitting source [88]. This reduces significantly the environmental impact of today's electricity generation and affords a greater diversity of electricity generation that enhances energy security. By the late 2000s, nuclear power was under serious consideration in over 45 countries which did not yet have it. Nuclear power was being reconsidered in the developed nations like US, UK , France, Finland. Energy from nuclear power plants was a central component of the national energy policy in the fast growing developing economies like India and China. China and India have a total of 34 new reactors under construction and many more units planned in order to meet their future energy demands [88]. This period saw, as it was termed in the press, a *nuclear renaissance*.

According to Kessides (2012) [50], the so called nuclear renaissance could be attributed to:

- An extremely strong record of global nuclear operations, with no high-profile incidents, for over two decades helped shift the perceptions about the environment and health risks of the nuclear energy.
- There was a fading memory of the Three Mile Island and Chernobyl accidents.
- High volatility in the fossil fuel prices called for an increased diversity in electricity generation, and

- Increased public concern over the greenhouse gas emissions, meant that nuclear energy was one of the leading candidates to shoulder the increased future energy demands.

The events in Fukushima however derailed the onset of a nuclear renaissance with the focus back on the safety of nuclear power plants. These events are likely to cause major regulatory changes thus further increasing the uncertainties in already uncertain economics of the nuclear industry. It can be argued that the nuclear renaissance began faltering even before the unfolding of events at Fukushima. This was due to the concerns over the large risks and uncertainties underlying the cost elements of nuclear power. These risks and uncertainties were reflected in the wide range of cost estimates for nuclear power plants. The cost overruns and schedule delays of Finland's new Olkiluoto plant and France's Flamenville plant are rekindling old fears about nuclear power being far too complex and costly. This raises new questions about the viability of new nuclear plants, especially in deregulated electricity markets.

Negative wholesale prices have become more common as European countries turn to renewables, particularly Germany with its forced march away from nuclear power, known as the *Energiewende*. Neighbours such as Poland and the Czech Republic complain that power surges from Germany are playing havoc with their grids [79]. Across Europe a strange consequence of subsidised renewables is that some governments now want to pay power companies to maintain the capacity to produce electricity from fossil fuels to ensure that backup power is available. More perversely, Europe is burning more heavily polluting coal at the expense of cleaner and more flexible gas. This is because coal is cheap, the gas market is far from liquid and the carbon-emissions system is broken [79]. Therefore, in the longer term, increasing concerns about the CO_2 emissions added to the need for electricity in bulk without intermittency, may imply stronger prospects for nuclear power.

The future of nuclear power depends on resolving the issues of safety of operations, safe management of radioactive wastes and measures to prevent proliferation (MIT, 2003)[26]. However, in a deregulated electricity market, the economics of NPPs will also be an important determinant of nuclear energy's role in the future global energy mix. The focus of the thesis lies in the development of models that can assist policy makers in taking decisions related to the economics of nuclear power plants, while accounting for the uncertainties that affect the costs and benefits of nuclear power plants. The underlying mathematical models used are generic which allows the techniques used here to be applicable to a wide range of topics which require stochastic optimization of decisions.

1.1 The Uncertain Economics Of Nuclear Power Plants

The costs of nuclear power can be categorized as:

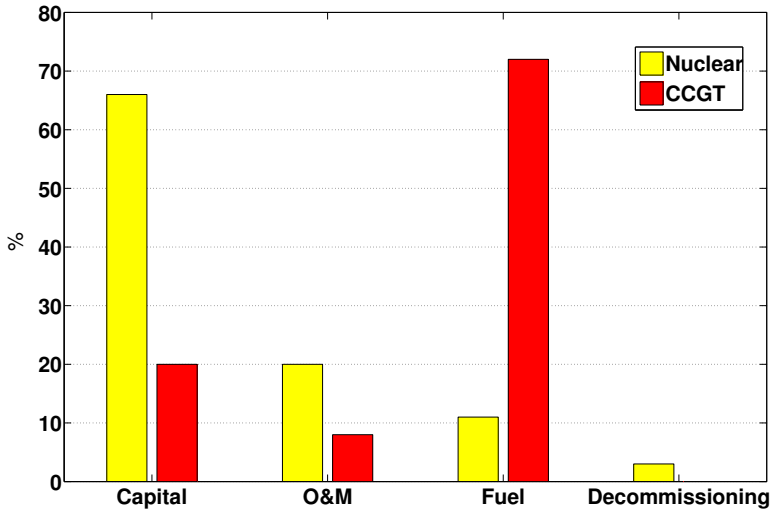

Figure 1.1: Cost profile of nuclear and combined cycled gas turbines (CCGT) based power plants. Source DTI (2007) [25]

- *Capital or construction costs* — are incurred during the planning, licensing and construction of a new nuclear power plant.
- *Operations and maintenance costs* — relate to the operations, management and maintenance of a power plant and includes the planned maintenance, labour, security, insurance, etc.
- *Fuel costs* — costs related to the back and front end fuel cycles.
- *Decommissioning costs* — relates to the decommissioning of the plant at the end of its operating life and long-term disposal and management of radioactive waste.

Figure 1.1 compares the cost breakup for a nuclear and a gas fired power plant. Much of the uncertainty in the economics of nuclear power plants relates to the construction or capital costs, which also is the most important component of the total costs. On the other hand, nuclear power plants would be fairly insensitive to the cost of fuel as it's a minor component of the total generating costs.

Kessides (2012) [50] identified the major reasons for the past escalation in construction costs of nuclear power plants as:

- Incorrect understanding of the economy of scale argument — the early cost projections while taking into account the economy of scale, usually ignored

the added costs that will be incurred due to increased complexity of larger nuclear power plants.

- Construction before design completion often necessitated costly redesign and significant construction delays which, with high interest rates, substantially increased the cost of building.
- Unwieldy licensing process and increased regulatory requirements often changing in mid-course leading to construction delays.
- Non-uniform designs inhibited the exploitation of the economy of volume principle and further compounded the complexity of the licensing process.
- Hesitant implementation of the remedial measures for emerging problems and identified risks and constraints.

While the uncertainties in the costs can to some extent be dealt with by better management and planning of the construction of the nuclear power plants, it is still important to take these uncertainties into account while doing an initial cost-benefit analysis of the plant. Real options may be used to estimate the value of keeping the nuclear option open for a firm confronted with not only uncertain capital costs but also uncertain fossil fuel, electricity, and carbon prices.

1.2 Reactor Generations

The nuclear reactors have been classified into different generations based on the reactor design and also partly on the era in which they were constructed [33]. The different generations of nuclear reactors, as shown in Figure 1.2 are:

- Generation I reactors: Generation I reactors were the initial designs built in 1950-1960s and were mostly early prototype of several designs [33].
- Generation II reactors: Generation II reactors are the commercial designs built between the 1960s and 1990s [33]. Most of the reactors operational right now belong to the Generation II type.
- Generation III reactors: Generation III and III+ reactors are evolutionary developed reactor designs largely based on the Generation II technologies but incorporate several enhancements; such as improved fuel technology and greater thermal efficiencies. Additionally, they have higher availability and are designed for longer operating life (>60 years) which vastly improves their economics. Generation III reactors have improved safety features which results in reduced probability of core meltdown and have a greater resistance to structural damage, for e.g. from impact of an aircraft. Advances to Generation III are underway, resulting in several (so called Generation III+) near-term deployable plants that are actively under development and are being considered for deployment in several countries [33].

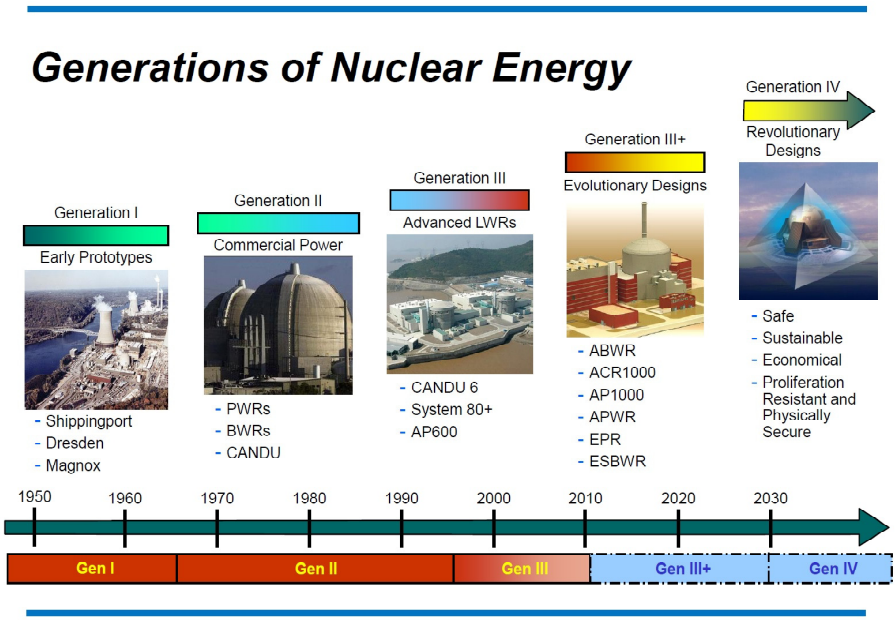

Figure 1.2: Reactor Generations [source: Generation IV roadmap ,2002 [33]]

Examples of Gen III reactors include the Advanced Boiling Water Reactor (ABWR), the Advanced Pressurized Water Reactor (APWR), AP- 600 and the enhanced CANDU 6. Examples of Gen III+ reactor designs include the Economic Simplified Boiling Water Reactor (ESBWR), AP-1000, the European Pressurized Reactor (EPR), VVER-1200, APR-1400 (see Kessides, 2012 [50]).

A common feature for Gen III and Gen III+ reactors is that they are designed to be large reactors so that they can benefit from the principle of the *economy of scale*[1].

- Generation IV reactors: Generation IV reactors represent a set of conceptual nuclear reactor designs currently being researched, with the year 2030 being considered as their earliest possible deployment date. The Generation IV reactor concepts are considered as revolutionary developed reactor designs compared to Generations II and III reactors. They were chosen amongst several innovative next generation designs with following goals in mind [33]:

  – Sustainability :

    Generate energy sustainably, and promote long-term availability of nuclear fuel;

[1] In microeconomics, economy of scale stands for the cost advantages that a business obtains due to expansion. There are factors that cause a producer's average cost per unit to fall as the scale of output is increased, mostly as fixed costs are spread out over more units of output. Often operational efficiency is also greater with increasing scale, leading to lower variable cost as well.

Minimize nuclear waste and reduce the long term stewardship burden.

- Safety & Reliability :

  Excel in safety and reliability;

  Have a very low likelihood and degree of reactor core damage;

  Eliminate the need for off-site emergency response.

- Economics :

  Have a life cycle cost advantage over other energy sources;

  Have a level of financial risk comparable to other energy projects.

- Proliferation Resistance & Physical Protection :

  Be a very unattractive route for diversion or theft of weapons-usable materials, and provide increased physical protection against acts of terrorism.

With these goals in mind the following six concepts have been identified as most promising: the very high temperature reactor (VHTR), the sodium cooled fast reactor (SFR); the supercritical water reactor (SCWR), the gas cooled fast reactor (GFR), the lead cooled fast reactor (LFR) and the molten salt reactor (MSR) [33]. Recently Locatelli et al. (2013) [52] presented a comprehensive overview of the Generation IV reactors, their main R&D areas, and their economic perspectives. A few of the Generation IV reactor concepts, like the VHTR are small and modular in nature. Small (< 300 MWe) and medium (< 700 MWe) reactors (SMRs) are considered an attractive option as they can benefit amongst others from the following facts :

- Modular construction : SMRs can be manufactured largely in a factory and delivered and installed module by module on site, bringing down the construction costs by learning effects and reduced construction time. Reduced construction times also bring down the financing costs and investment risks [62].
- Better siting options: Smaller reactors can be installed in remote locations that have little or no access to the grid, where large scale plants cannot be accommodated [62].
- Investment flexibility: Modular construction of SMRs provides greater flexibility of investment decisions, wherein, if the economic scenario, such as the cost of electricity and its demand doesn't turn out as anticipated; or the cost of the modules is higher than expected, then investment in future modules can either be delayed or abandoned completely, without affecting the modules already ordered.

One of the recurring topics discussed in later chapters involves comparing the economic benefits of flexibility for SMRs with the benefits of economy of scale for large units.

1.3 Real Options

In deregulated global electricity markets, the economics of nuclear power plants will play a key role in the future decisions to build new reactors. The various cost components of a nuclear power plant have several sources of uncertainties which makes it difficult to garner a consensus on what will be the cost of a new nuclear generating plant. Uncertain costs in a deregulated market makes the decision making process even more difficult as it is no longer possible to pass on the unexpected costs to the end users. When the costs and benefits of a project are uncertain, in addition to the conventional net present value (NPV) approach, the use of *real options* can add significant information for better valuation of the project. Projects with economic uncertainties should take into account the value of flexibility which arises from the possibility of delaying the investment or in certain cases abandoning the project at a future stage if processes underlying economic circumstances turn unfavourable.

The real options approach for making investment decisions in projects with uncertainties, pioneered by Arrow and Fischer (1974) [3], Henry(1974) [37], Brennan and Schwartz (1985) [14] and McDonald and Siegel (1986) [57] became accepted in the past decade. Dixit and Pindyck (1991) [27] and Trigeorgis (1996) [82] comprehensively describe the real options approach for investment in projects with uncertain future cash flows.

Real option analysis (ROA) has been applied to value real assets like mines (Brennan and Schwartz (1985)), oil leases (Paddock, et. al (1988)), patents and R&D (Schwartz (2002)). Pindyck (1992) [67] used real options to analyse the decisions to start, continue or abandon the construction of nuclear power plants. Rothwell (2004) [73] used ROA to compute the critical electricity price at which a new advanced boiling water reactor should be ordered in Texas.

Until recently, the valuation of investment projects was done exclusively using the discounted cashflow method (DCF), which computes the net present value (NPV) of a project, given a deterministic net cashflow structure, C_{t_m}, and a known discount factor r as:

$$\text{NPV} = \sum_{m=0}^{M} \frac{C_{t_m}}{(1+r)^{t_m}}.$$

If the net cashflow, C_{t_m}, at time t_m is positive, it indicates a cash inflow, while a negative value indicates a cash outflow. It is assumed in the model that the net cashflow at each of the M discrete time steps is known exactly. More often than not, investment projects are affected by multiple uncertain factors, that make it difficult to predict accurately the cashflows at future time steps. Additionally, DCF excludes the role of future management decisions — that can be made in the lieu of emerging information during the lifetime of the project — on the future cashflow structure.

Real option analysis can serve as an indispensable tool when the project valuation involves:

- *Option to delay :* Irreversible investments decisions that are affected by uncertain market conditions can benefit from the inherent value of the option to delay the investment. As more information is revealed with time, the management can choose to invest, if market scenarios turn positive, or not undertake the project if conditions turn otherwise. The value of this option is derived from the flexibility of the firm to optimally choose the time point when it makes the investment decision.
- *Option to abandon :* Investors, especially in projects which involve significant R&D, can benefit from the flexibility that arises from the option to abandon the project in the future.
- *Growth option :* Growth options have a strategic value and are particularly relevant for projects that are not profitable in themselves but might open up more lucrative possibilities in the future. Investment decisions that are made sequentially are often examples of growth options. For example, investing in new oil production capacity first involves an investment for exploration followed by the development of the wells. The first sequence of investment derives its value from the fact that it can lead to the possibility of the second phase of more profitable investment.
- *Option to switch :* The option derives its value from the flexibility of switching the inputs and outputs depending on the market conditions.

A factor common to the above list is the *flexibility of decisions* which can alter the future cashflow structure. Real options allow us to compute the value that arises from flexible decisions if they are made optimally in the future. To briefly introduce the real option valuation technique we take a simple toy example discussed in Dixit and Pindyck (1994) [27].

The objective of the example is to evaluate an irreversible investment decision into a widget factory. It is assumed that the factory can be set up immediately and will cost a fixed amount K = \$1600. Once set up, the factory produces a widget a year for eternity. The present cost of the widget is \$200, but may go up to \$300 and down to \$100 with probability, $q = 0.5$ and $(1-q)$, respectively. For this example it is assumed that the price will remain constant after the first year (see Figure 1.3). Also for simplicity it is assumed that risks involved are diversifiable, which allows the future cashflows to be discounted using the risk-free rate of interest, which is taken as 10 percent for this example.

A DCF analysis of the investment, which cannot take into account the value of flexibility of delaying the investment decision, will value the project as:

$$\text{NPV} = -1600 + 200 + 0.5\left(\sum_{t_m=1}^{\infty} \frac{300}{(1.1)^{t_m}}\right) + 0.5\left(\sum_{t_m=1}^{\infty} \frac{100}{(1.1)^{t_m}}\right), \tag{1.1}$$

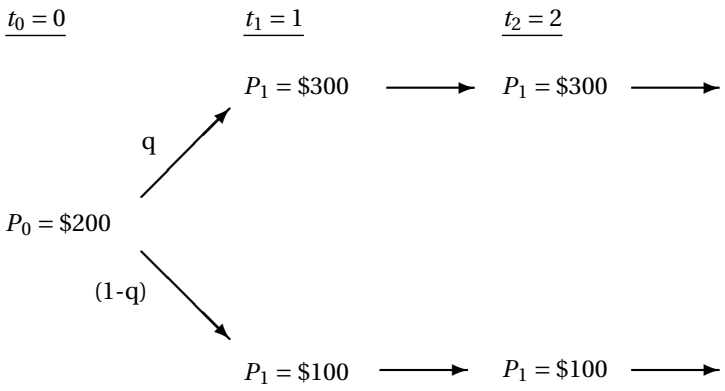

Figure 1.3: Price of widgets for two periods problem.

which is equal to \$400. As the NPV of the project is positive, according to a DCF analysis the investment should be made immediately. However, if we include the flexibility of delaying the investment, the value of the project undertaken at the next time step if the widget prices go up would be,

$$\text{NPV} = -1600 + \left(\sum_{t_m=0}^{\infty} \frac{300}{(1.1)^{t_m}} \right) = \$1700,$$

while if the widget prices go down, the value of the project will be

$$\text{NPV} = -1600 + \left(\sum_{t_m=0}^{\infty} \frac{100}{(1.1)^{t_m}} \right) = -\$500.$$

If the decision to construct the factory can be delayed up to the next time step according to an optimal decision the project will be undertaken only if the price of widgets goes up. To decide whether the project should be undertaken now or at the next time step, we compare the cashflow obtained when the decision to invest is delayed until the next time step and the optimal policy is followed in the future; with the one obtained if the factory is set up immediately. Therefore, if the investment decision is delayed until the next time step, the net present value of the project will be:

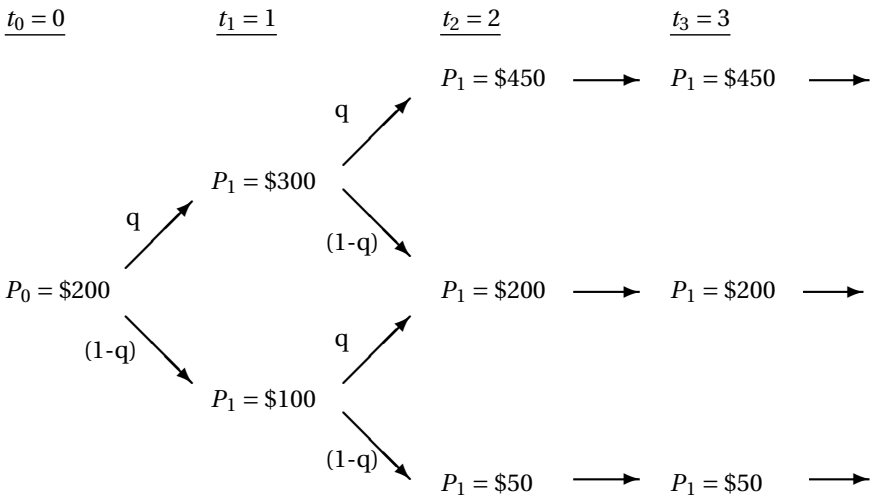

Figure 1.4: Price of widgets for three periods problem.

$$\text{NPV} = 0.5\left[\frac{-1600}{1.1} + \left(\sum_{t_m=1}^{\infty}\frac{300}{(1.1)^{t_m}}\right)\right] = \$773,$$

while if the factory is set up immediately the NPV of the project would be given by Equation (1.1), i.e. \$400. As the NPV for the factory, if the decision to set it up is made at the next time step, is greater than if it is set up immediately; for the present example it is more profitable to delay the investment decision. The inherent value of flexibility to delay the investment decision for this case is equal to \$373.

The real option pricing problem can however be difficult due to the following reasons:

- The above problem can become more challenging if the uncertainty in the costs of widgets is not limited to the first time step. Figure 1.4 depicts the above problem extended to three time steps, with the price of the widget being uncertain for the first two time steps. The problem then is to determine when it may be optimal to invest:

  1. invest immediately;
  2. wait a year and then invest if the price has gone up; but *never* invest if the price has gone down;
  3. wait a year and invest if the price has gone up, but if it went down wait another year and invest if it then goes up;
  4. wait two years and only invest if the price has gone up both times;

5. never invest.

An optimal strategy to invest here can be determined following Bellman's dynamic programming principle which is described for a more general problem in Section 1.4. In general there could be multiple time steps and the optimal investment strategy would then have to be chosen from several realized possibilities.

- In a more practical application the widget price wouldn't just go up and down to some discrete states but can with a positive probability attain any (generally positive) value.
- Another practical challenge involved in the real option valuation of projects is that projects are usually affected by *multiple sources of uncertainty* and that makes the problem a high-dimensional one. In the above example, not only the price of the widgets could be uncertain, but also the cost of setting up the factory, the time it takes to set up the factory, the number of years the factory will operate, the discount factor etc. Use of lattice, binomial tree- or finite-difference based methods may therefore be unsuitable for real option valuation, when there are more than three sources of uncertainties, as these methods suffer from the *curse of dimensionality*.

As most real options have a so-called early-exercise feature, the pricing problem involved is similar to that of their financial counterparts, namely the American options. To develop a suitable pricing model for real options it is then natural to study the pricing methods used for financial options, especially the American options and their discrete time versions — the Bermudan options.

1.4 Bermudan Options

An option is a financial contract which enables its buyer to buy an asset at a future time for a predetermined price. A Bermudan option is an option where the buyer has the right to exercise at a set (discretely spaced) of times. This is intermediate between a European option which allows exercise at a single time, namely expiry and an American option, which allows exercise at any time. With an increasing number of exercise opportunities Bermudan option values approach the value of an American option.

Bermudan options can broadly be categorized into call- and put- options. A Bermudan call option gives the buyer the right, but not the obligation to buy an asset from the seller of the option at a certain pre-specified dates (between the issue and expiry date) for a certain price (the strike price). The seller has an obligation to sell the asset to the buyer if the buyer so decides. The buyer pays a fee for this right, which is the value of holding the option. A put option on the other hand gives the buyer the right, but not the obligation, to sell an asset at certain pre-specified dates

(between the issue and maturity) for a certain price (the strike price) to the seller of the option.

The real option example discussed above can be seen as as a Bermudan call option, with the strike price equal to the cost of setting up the plant and the cashflow from selling up the widgets as the asset price. If on the other hand the cost of setting up the plant was uncertain while the price of widgets were known and constant then the problem would be similar to a Bermudan put option.

Pricing of Bermudan options especially for multi-dimensional processes is a challenging problem owing to its path-dependent settings. As discussed above, the traditional valuation methods, such as lattice and tree-based techniques are often impractical in such cases due to the curse of dimensionality and hence are used only in the low-dimensional cases. In recent years many simulation-based algorithms have been proposed for pricing Bermudan options, most of which use a combination of Monte Carlo simulations and dynamic programming to estimate the option price.

Dynamic Programming Formulation

In order to compute the value of a Bermudan option we are required to find an optimal policy to make early exercise decisions. Such a policy can be obtained by following *Bellman's optimality principle*, wherein an optimal solution is found starting from the final time step and then recursively moving backwards until we reach the initial time step, where at each time step the optimal decision is determined.

The dynamic programming formulation for the Bermudan option pricing problem is defined here. A complete probability space $(\Omega, \mathscr{F}, \mathbb{P})$ and finite time horizon $[0, T]$ are assumed. Ω is the set of all possible realizations of the stochastic economy between 0 and T. The information structure in this economy is represented by an augmented filtration $\mathscr{F}_t : t \in [0, T]$, with $\mathscr{F}_t$ the sigma field of distinguishable events at time t, and $\mathbb{P}$ is the risk-neutral probability measure on elements of $\mathscr{F}$. It is assumed that $\mathscr{F}_t$ is generated by W_t, a d-dimensional standard Brownian motion, and the state of economy is represented by an $\mathscr{F}_t$-adapted Markovian process, $\mathbf{S}_t = (S_t^1, \ldots, S_t^d) \in \mathbb{R}^d$, where $t \in [t_0 = 0, \ldots, t_m, \ldots, t_M = T]$. Let $h_t := h(\mathbf{S}_t)$ be an adapted process representing the intrinsic value of the option, i.e. the holder of the option receives $\max(h_t, 0)$, if the option is exercised at time t. With the risk-free savings account process $B_t = \exp(\int_0^t r_s \, ds)$, where r_t denotes the instantaneous risk-free rate of return, we define

$$D_{t_m} = \frac{B_{t_m}}{B_{t_{m+1}}}.$$

The special case where r_t is constant is considered. The problem is then to compute

$$V_{t_0}(\mathbf{S}_{t_0}) = \max_{\tau} \mathbb{E}\left[\frac{h(\mathbf{S}_\tau)}{B_\tau}\right],$$

where τ is a stopping time, taking values in the finite set $\{0, t_1, \ldots, T\}$. The value of the option at the terminal time T is equal to the product's pay-off,

$$V_T(\mathbf{S}_T) = \max(h(\mathbf{S}_T), 0). \tag{1.2}$$

The conditional continuation value Q_{t_m}, i.e. the expected pay-off at time t_{m+1}, is given by:

$$Q_{t_m}(\mathbf{S}_{t_m}) = D_{t_m}\mathbb{E}\left[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m}\right]. \tag{1.3}$$

The Bermudan option value at time t_m and state $\mathbf{S}_{t_m}$ is given by

$$V_{t_m}(\mathbf{S}_{t_m}) = \max(h(\mathbf{S}_{t_m}), Q_{t_m}(\mathbf{S}_{t_m})). \tag{1.4}$$

We are interested in finding the value of the option at the initial state $\mathbf{S}_{t_0}$, i.e. $V_{t_0}(\mathbf{S}_{t_0})$.

The Least Squares Method

One of the most widely used simulation based methods for pricing Bermudan options was proposed by Longstaff and Schwartz in 2001 called the least squares method (LSM) [53]. As LSM serves as an important reference model for comparing results throughout the thesis, the method is summarily described here.

Success of the LSM can be attributed to the reformulation of the dynamic programming principle, described above, in terms of the optimal stopping time rather than in terms of the value process. The LSM approximates the continuation value at each time step, given by Equation (1.3), using a least-squares regression jointly with the cross-sectional information provided by Monte Carlo simulation. By comparing the estimated continuation values with the immediate exercise values, the optimal exercise decision is identified. This procedure is repeated recursively going backwards in time. After discounting the obtained cashflows to time zero, the price of the Bermudan option is computed. Therefore, the dynamic programming formulation given by Equations (1.2) to (1.4) is rewritten in terms of optimal stopping times, denoted by τ, as:

$$\begin{cases} \tau(t_M) &= T, \\ \tau(t_m) &= t_m \mathbf{1}_{\{Q_{t_m}(\mathbf{S}_{t_m}) \leq h(\mathbf{S}_{t_m})\}} + \tau(t_{m+1})\mathbf{1}_{\{Q_{t_m}(\mathbf{S}_{t_m}) > h(\mathbf{S}_{t_m})\}}, \quad m < M-1. \end{cases} \tag{1.5}$$

The problem is then to compute the continuation value by,

$$Q_{t_m}(\mathbf{S}_{t_m}) = \mathbb{E}\left[\frac{B_{t_m}}{B_{\tau(t_{m+1})}} h(\mathbf{S}_{\tau(t_{m+1})})|\mathbf{S}_{t_m}\right].$$

In LSM this computation is done by simulating independent copies of sample paths, $\{\mathbf{S}_{t_0}(n), \ldots, \mathbf{S}_{t_M}(n)\}$, $n = 1, \ldots, N$, of the underlying process $\mathbf{S}_t$, all starting from the same initial state $\mathbf{S}_{t_0}$. The continuation value is then approximated moving backwards in time as,

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}(n)) = \sum_{k=0}^{K} \alpha_{t_m}(k)\phi_k(\mathbf{S}_{t_m}(n)),$$

which satisfies,

$$\underset{\alpha_{t_m}}{\operatorname{argmin}} \frac{1}{N} \sum_{n=1}^{N} \Big| \frac{B_{t_m}}{B_{\tau(t_{m+1})}} h(\mathbf{S}_{\tau(t_{m+1})}(n)) - \widehat{Q}_{t_m}(\mathbf{S}_{t_m}(n)) \Big|^2. \tag{1.6}$$

Here $\phi = (\phi_0, \ldots, \phi_K)'$ forms an orthogonal basis. The basis functions are usually polynomials of the state variables. The option price is then given by:

$$\tilde{V}_{t_0}(\mathbf{S}_{t_0}) = \mathbb{E}\left[\frac{h(\mathbf{S}_{\tau(t_1)})}{B_{\tau(t_1)}} \right].$$

Longstaff and Schwartz (2001) [53] suggest that using only *in-the-money* paths for the regression step, given by Equation 1.6) helps improving the result. Although, the LSM is simple to implement and computationally efficient, the approximated continuation values at the intermediate time steps are generally noisy and result in a sub-optimal early exercise policy. The improvements suggested in the literature to the original LSM, to obtain a better early exercise policy, such as boundary-distance-grouping and use of nested-simulations, (see Broadie and Cao (2009) [15]), make the method computationally expensive.

A contribution of this thesis is the introduction of the *Stochastic Grid Bundling Method* (SGBM), which overcomes some of the drawbacks of LSM, while still being computationally competent.

1.5 Reading Guide

The thesis is divided into two parts, essentially. Part 1 deals with the development of a fast and accurate pricing method for options with early exercise features, and consists of the following chapters.

Chapter 2 introduces the *Stochastic Grid Method* (SGM) for pricing Bermudan options. Although this method has favourable characteristics, such as the estimated continuation value in the intermediate time steps being less noisy when compared to LSM, it is unsuitable for high-dimensional problems. However, it still plays an important role in this thesis as it serves as a precursor to the stochastic grid bundling method (SGBM), which employs and improves upon some of the techniques used in SGM.

Chapter 3 introduces the *Stochastic Grid Bundling Method* for pricing Bermudan options and fast approximation of their sensitivities to the underlying assets. The method benefits from the two most popular approaches for Bermudan option pricing methods, namely *regression based methods* and *state-space partitioning based methods*. It uses regression for reducing the high-dimensional state space to a low-dimensional space, which makes SGBM computationally efficient, while state-space partitioning helps in better sampling of the conditional distribution, which

makes the method more accurate. The efficiency of the method is demonstrated through a sequence of numerical examples with increasing complexity.

Part 2 of the thesis is focussed on applications that involve optimal economic decisions for nuclear power plants, with the optimal decisions made using the methods developed in Part 1. It consists of the following chapters.

Chapter 4 discusses the value of flexibility that arises from the modular construction of small and medium sized reactors, SMRs, (according to IAEA, 'small' refers to reactors with power less than 350 MWe, and 'medium' with power less than 700 MWe). SMRs benefit from flexibility of investment, reduced upfront expenditure, enhanced safety, and easy integration with small sized grids. Large reactors on the other hand have been an attractive option due to the economy of scale. The focus of this chapter is to analyse the economic impact of flexibility due to modular construction of SMRs, under different considerations of decision time, uncertainty in electricity prices and constraints on the construction of units. For the real option valuation SGM is used in this chapter.

Chapter 5 extends the findings of Chapter 4 to more realistic cases. Real option analysis is used to compare different construction strategies, such as; constructing twin units vs constructing two large reactors of equivalent size independently, or modular SMRs which benefit from the learning effect and flexibility vs large reactors that benefit from the economy of scale. This chapter also analyses the effect of uncertain lifetime of operation for nuclear reactors, which could arise due to possible lifetime extension on one hand or premature permanent shut-down due to unforeseen events on the other hand. For this chapter SGBM is used as the underlying stochastic optimization method.

Chapter 6 presents a decision-support tool, which takes into account the major uncertainties in the cost elements of a nuclear power plant, to ultimately provide an *optimal portfolio* of nuclear reactors. Capital costs, fuel, operation and maintenance (O&M) costs, and electricity prices play a key role in the economics of nuclear power plants, where especially capital costs are known to be highly uncertain. Different nuclear reactor types compete economically by having either lower and less uncertain construction costs, increased efficiencies, lower and less uncertain fuel cycles and O&M costs, etc. The proposed decision tool uses a holistic approach that takes into account the key economic factors and their uncertainties to compute an optimal portfolio of nuclear reactors. The portfolio so obtained, under the model assumptions, maximizes the combined returns for a given level of risk or uncertainty. The decisions are made using a combination of real option theory, which uses SGBM as the underlying pricing method, and mean-variance portfolio optimization.

Chapter 7 summarizes the findings from the other chapters and reflects upon the broader implication of this research.

Part I

Method Development

CHAPTER

2 Pricing Bermudan Options Using The Stochastic Grid Method

The contents of this chapter have appeared in [40]. The *stochastic grid method* SGM, which serves as a basis for the SGBM, to be introduced in the next chapter, is a dynamic programming based option pricing method which recursively computes the option price, moving backwards in time. A functional approximation, obtained using regression, of the option price at a given time step is used to compute the option price at the previous time step. The dimensionality of the problem is recursively reduced using the intrinsic value as a mapping function. Although numerical results are given for high-dimensional problems, the error for SGM is shown to be bounded for a one-dimensional problem.

Several simulation-based methods have been proposed to price options with early-exercise features, which combine Monte Carlo path generation and dynamic programming techniques to determine optimal exercise policies. The class of *regression-based methods* has been developed by Carriere (1996)[20], Tsitsiklis and Van Roy (1999)[83], containing the *Least Squares Method* (LSM) by Longstaff and Schwartz (2001)[53] as its most prominent member. A detailed analysis of regression-based methods can be found in Glasserman (2003)[34].

Another approach is based on approximating the transition probabilities using either *bundling*, as in Tilley (1990) [81], *partitioning*, as in Barraquand and Martineau (1997)[7] and Jin *et al.* (2007)[45], or *quantization*, as in Bally and Pages (2004)[5], of the state space; or computing weights to approximate these conditional probabilities, as in the *stochastic mesh method* Broadie and Glasserman (2004)[17].

Other than the above two approaches, there exist *duality-based methods* proposed by Haugh and Kogan (2004)[36] and Rogers (2002)[71]. By a duality-based method an upper bound on the option value for a given exercise policy can be obtained, by adding a non-negative quantity that penalizes potentially incorrect exercise decisions made by a sub-optimal policy.

SGM has certain advantages over existing methods. LSM [53], although computationally fast and simple to implement, uses a large number of paths to obtain a good exercise policy. Also the number of basis functions required for regression grows almost exponentially with the dimensions of the problem. SGM on the other

hand can be used to obtain a good exercise policy using far fewer paths. The number of basis functions used in SGM is almost independent of the dimensions of the problem. SGM uses sub-simulation when moments required to approximate the transition density function are unavailable, which can make the method computationally rather expensive in this situation.

This chapter is organized as follows, Section 2.2 is devoted to the description of the SGM. In Section 2.3 a basic error analysis for a one-dimensional problem is reported and some of the results for the single asset case are discussed. In Section 2.4 the results for high-dimensional problems are discussed and compared with the other available Monte Carlo techniques. Section 2.5 provides some conclusions, and points out some of the limitations of the method.

2.1 Problem Formulation

A complete probability space $(\Omega, \mathscr{F}, \mathbb{P})$ and finite time horizon $[0,T]$ are assumed here. Ω is the set of all possible realizations of the stochastic economy between 0 and T. The information structure in this economy is represented by an augmented filtration $\mathscr{F}_t : t \in [0,T]$, with $\mathscr{F}_t$ the sigma field of distinguishable events at time t, and $\mathbb{P}$ is the risk-neutral probability measure on elements of $\mathscr{F}$. It is assumed that $\mathscr{F}_t$ is generated by W_t, a d-dimensional standard Brownian motion, and the state of the economy is represented by an $\mathscr{F}_t$-adapted Markovian process, $\mathbf{S}_t = (S_t^1, \ldots, S_t^d) \in \mathbb{R}^d$, where $t \in [t_0 = 0, \ldots, t_m, \ldots, t_M = T]$. Let $h_t := h(\mathbf{S}_t)$ be an adapted process representing the intrinsic value of the option, i.e. the holder of the option receives $\max(h_t, 0)$, if the option is exercised at time t. With the risk-less savings account process $B_t = \exp(\int_0^t r_s \, ds)$, where r_t denotes the instantaneous risk-free rate of return, we define

$$D_{t_m} = \frac{B_{t_m}}{B_{t_{m+1}}}.$$

The special case where r_t is constant is considered. The problem is then to compute

$$V_{t_0}(\mathbf{S}_{t_0}) = \max_{\tau} \mathbb{E}\left[\frac{h(\mathbf{S}_\tau)}{B_\tau}\right],$$

where τ is a stopping time, taking values in the finite set $\{0, t_1, \ldots, T\}$. The value of the option at the terminal time T is equal to the product's pay-off,

$$V_T(\mathbf{S}_T) = \max(h(\mathbf{S}_T), 0). \tag{2.1}$$

The conditional continuation value Q_{t_m}, i.e. the expected pay-off at time t_{m+1}, is given by:

$$Q_{t_m}(\mathbf{S}_{t_m}) = D_{t_m} \mathbb{E}\left[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}) | \mathbf{S}_{t_m}\right]. \tag{2.2}$$

The Bermudan option value at time t_m and state $\mathbf{S}_{t_m}$ is given by

$$V_{t_m}(\mathbf{S}_{t_m}) = \max(h(\mathbf{S}_{t_m}), Q_{t_m}(\mathbf{S}_{t_m})). \tag{2.3}$$

We are interested in finding the value of the option at the initial state $\mathbf{S}_{t_0}$, i.e. $V_{t_0}(\mathbf{S}_{t_0})$.

2.2 The Stochastic Grid Method

The stochastic grid method (SGM) solves a general optimal stopping problem using a hybrid of dynamic programming and Monte Carlo methods. The method first computes the optimal exercise policy and a direct estimator of the true option price. The lower bound values are computed by discounting the pay-off obtained by following this exercise policy. The details on how these bounds are obtained are described in the sections to follow.

Method Details of SGM

A (Markovian) discretization scheme which is easy to simulate, *e.g.* the Euler scheme, is used to generate N sample paths originating from the initial state $\mathbf{S}_{t_0}$. When the diffusion process appears in closed form, such as the case of the commonly used multi-dimensional Black - Scholes model, the sample paths can be generated directly. The stochastic grid points $\mathbf{S}_{t_m}(n)$, $n = 1,\ldots,N$, can be interpreted as the intersections of the sample paths with a plane representing different intermediate time steps t_m. Figure 2.1 shows the grid points for an option with two underlying assets $\mathbf{S}_{t_m} = (S^1_{t_m}, S^2_{t_m})$ starting from the initial state $\mathbf{S}_{t_0} = (100, 100)$ at two different time intervals, where one is closer to the initial time, while the other is closer to the final exercise time T.

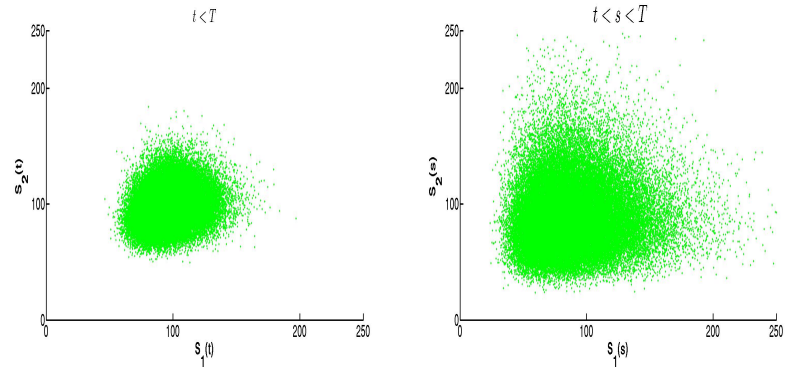

(a) (b)

Figure 2.1: Grid Points (30000 × 30000), Figure (a) at t , Figure (b) at s where $t < s < T$.

The value of the option at the expiration time $t_M = T$ will be equal to its pay-off given by $\max(h(\mathbf{S}_T), 0)$. Only financial derivatives whose pay-off are an element of the space of square integrable or finite variance functions are considered. Examples of $h(\mathbf{S}_t)$ on multiple assets include, for a basket call option,

$$h(\mathbf{S}_t) = (w_1 S^1_t + \cdots + w_d S^d_t - \mathcal{K}),$$

for an out-performance option

$$h(\mathbf{S}_t) = (\max(S_t^1, \ldots, S_t^d) - \mathscr{K}),$$

where $\mathscr{K}$ is the option strike price.

Computing the Optimal Exercise Policy

The main obstacle in pricing Bermudan options using Monte Carlo methods is the fact that the optimal exercise policy is not known. SGM computes the continuation value at each grid point, starting from the grid points at the expiration time $t_M = T$ and moving backwards in time. The option is exercised if the immediate pay-off is greater than the discounted continuation value.

The direct estimator is defined recursively starting with $\widehat{V}_T(\mathbf{S}_T) = \max(h(\mathbf{S}_T), 0)$, and for $m = M-1, \ldots, 1$, by

$$\widehat{V}_{t_m}(\mathbf{S}_{t_m}) = \max\left(h(\mathbf{S}_{t_m}), D_{t_m}\mathbb{E}\left[\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right)|\mathbf{S}_{t_m}\right]\right), \tag{2.4}$$

where,

$$\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right) = \mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right]. \tag{2.5}$$

Mapping function $g(\cdot)$ maps the high-dimensional $\mathbf{S}_{t_{m+1}}-$ space to a low-dimensional $g(\mathbf{S}_{t_{m+1}})-$ space.

$\mathbb{E}\left[\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right)|\mathbf{S}_{t_m}\right]$ represents the continuation value for $\mathbf{S}_{t_m}$. Using iterated conditioning it can be shown that,

$$\begin{aligned}\mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m}\right] &= \mathbb{E}\left[\mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right]|\mathbf{S}_{t_m}\right] \\ &= \mathbb{E}\left[\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m})|\mathbf{S}_{t_m}\right]. \end{aligned} \tag{2.6}$$

In the sections to follow we discuss how to approximate $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right)$ and the choice of the mapping function $g(\cdot)$. Once the functional approximation, $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}\right)$, is obtained it can be used to compute the discounted continuation value at the grid points for t_m and thus to make the optimal exercise decision, i.e. exercise if the discounted continuation value is less than the immediate pay-off.

Parametrization of the option values

The continuation value at time t_m and state $\mathbf{S}_{t_m}$, i.e. $Q_{t_m}(\mathbf{S}_{t_m})$ can be computed from Equation (2.2). Instead of using the direct functional approximation of the option price at t_{m+1}, i.e. $\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})$ the law of iterated conditioning, *i.e.*

$$\mathbb{E}\left[\mathbb{E}\left[X|\mathscr{G}\right]|\mathscr{H}\right] = \mathbb{E}\left[X|\mathscr{H}\right],$$

where $\mathscr{H}$ is the sub- $\sigma-$ algebra of $\mathscr{G}$, is used to compute the continuation value. Then, the continuation value can be written as (2.6).

In order to compute $Q_{t_m}(\mathbf{S}_{t_m})$ from Equation (2.6) the functional form of $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right)$ is required. At the expiration time, the option value is given by Equation (2.1).

It is assumed that the unknown functional form of $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right)$ can be represented by a linear combination of a countable set of $\mathscr{F}_{t_{m+1}}$-measurable basis functions, where $\mathscr{F}_{t_{m+1}}$ is the information set at time t_{m+1}.

Similar to the regression-based algorithms (Tsitsiklis & Van Roy, 1999 [83], Longstaff & Schwartz , 2001 [53]) SGM approximates the unknown functional form of $\mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right]$ by projecting it on the first $K(<\infty)$ polynomial basis functions.

Remark 1. *In the examples to follow the function $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right)$ is approximated in SGM by $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}\right)$, as all the grid points at t_{m+1} generated from source $\mathbf{S}_{t_0}$ are used in the regression. The exercise policy obtained is still accurate as shown by the numerical results (lower bound values). An improved approximation will be based on a more sophisticated regression scheme, where grid points at t_m are bundled based on proximity, and only those grid points at t_{m+1} are used for regression to approximate $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right)$ that originate from the bundle containing $\mathbf{S}_{t_m}$, discussed in the next chapter.*

When $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right)$ is approximated by $\widehat{Z}_{t_{m+1}}\left(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}\right)$, an accurate early-exercise policy is obtained when $g(\cdot)$ is equal to $h(\cdot)$. However, also other choices of $g(\cdot)$ can be made. For other choices, it becomes important that the grid points are bundled based on some nearest neighbour rules to get an accurate exercise policy. In the special case when $g(\cdot)$ is chosen to be constant, SGM with bundling would very closely resemble the state space partitioning method by Jin et al. (2007) [45].

We denote this approximation by $Z^K_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0})$. Equation (2.5) is approximated over a set of K polynomial basis functions, as

$$Z^K_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}) = \mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}\right] = \sum_{k=0}^{K-1}\alpha_{t_{m+1}}(k)\phi_k(g(\mathbf{S}_{t_{m+1}})), \tag{2.7}$$

such that at each time step

$$r = \min_{\alpha_{t_m}}\sum_{1}^{N}|Z^K_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}) - V_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|^2, \tag{2.8}$$

where $\{\phi(\cdot)\}_{k=0}^{K-1}$ forms a set of basis functions, and r is the sum of squared residual errors.

This approximation can be justified if it is assumed that $V_{t_m}(\mathbf{S}_{t_m})$ is an element of the L^2 space of square integrable functions relative to some measure and therefore can be written as the linear combination of basis functions. Rather than regressing over the entire $g(\mathbf{S}_{t_{m+1}})-$ space better accuracy is obtained in SGM by *piecewise regression*, as explained in Section 2.3 and shown in the specific examples to follow.

Mapping high-dimensional state onto single-dimensional $g(\cdot)-$ space

In an approach similar to Barraquand and Martineau's SSAP method [7], the dimensionality of the problem is reduced by using $g(\mathbf{S}_{t_{m+1}})$ rather than the cross-products of the underlying states (as in LSM [53]) for regression.

Figure 2.2 shows in a schematic diagram how dimension reduction works in SGM. In order to compute the continuation value at $\mathbf{S}_{t_m}$ directly, a high-dimensional transition density function would be required, as shown in Figure 2.2(a). In SGM, however, first the option value at t_{m+1} is projected on the $g(\mathbf{S}_{t_{m+1}})$-space, see Figure 2.2(b). In other words, conditional expectation,
$\mathbb{E}[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}})]$, is computed using the least squares regression. The continuation value is then computed using the tower property as explained in Equation (2.6), which involves a one-dimensional transition density function. When all grid points at t_{m+1} are used for regression, $\mathbb{E}\left[\widehat{V}(t_{m+1},\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}\right]$, is computed instead of $\mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_m}\right]$. A better approximation is obtained by bundling the grid points at t_m based on proximity and using only those grid points at t_{m+1} that originate from the bundle containing $\mathbf{S}_{t_m}$, for regression (see the next chapter for details). However, it is found that in the case that all grid points at t_m are in a single bundle, a very satisfactory exercise policy (as is reflected in the lower bounds) can still be obtained when the mapping function $g(\cdot)-$ is of the form of the payoff function for the problems considered.These latter results for higher-dimensional problems are reported in the numerical section.

Boyle *et al.* (1997)[12] and Broadie and Detemple (1996) [16] show that the pay-off value is not a sufficient statistic for determining the optimal exercise decision for options on the maximum of several assets for SSAP. This argument however is specific to the SSAP and does not apply to SGM. In the SSAP method the state space is first mapped to the partitions (cells) along the pay-off space $h(\mathbf{S}_t)$ and then the same exercise decision is applied for all underlying states that fall into a particular cell or partition. This results in seemingly far off state points (like (100,90),(100,100) and (100,50)) to have the same exercise decision. In SGM *first* the exercise decision is made for each underlying state $\mathbf{S}_{t_m}$ (or grid point) at time step t_m and only *then* the state space is reduced to $g(\mathbf{S}_{t_m})$.

In order to give a better intuition of the method and allay the concerns raised by Boyle *et al.* [12], consider the same example given by them. Figures 2.3 to 2.6 show the evolution of two asset prices $\mathbf{S}_t = (S_t^1, S_t^2)$ with two exercise time steps. The option pay-off, $h(\mathbf{S}_t = (S_t^1, S_t^2)) = g(\mathbf{S}_t) = \max(S_t^1, S_t^2)$ and for convenience the risk-free interest rate is taken to be zero. It is straight forward to see that the option value

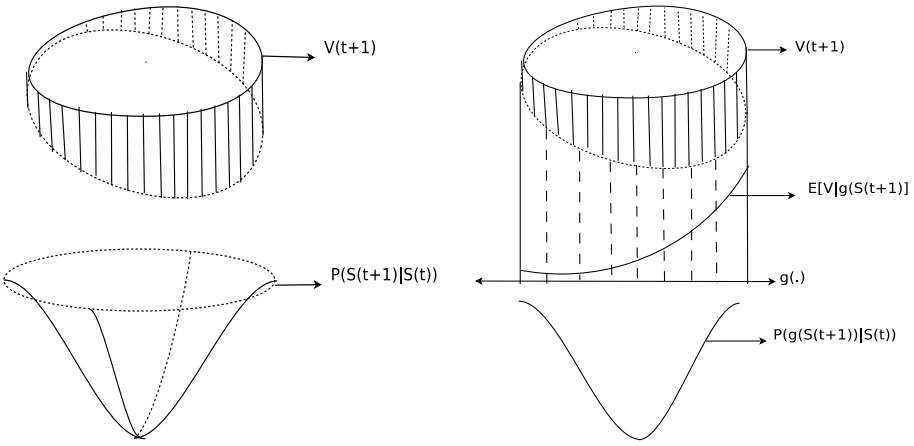

Figure 2.2: Schematic diagram showing how dimension reduction works in SGM. The option value at step t_{m+1} is given, Figure (a) shows the conventional way of computing the continuation value at $\mathbf{S}_{t_m}$, based on $\mathbb{P}(\mathbf{S}_{t_{m+1}}|\mathbf{S}_{t_m})$; Figure (b) shows how the continuation value is computed in SGM by means of projection $\mathbb{E}[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}})]$ onto $g(\cdot)$ and $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$.

at t_0 is 11. The steps followed at each time step starting from the final expiration time t_2 are

- Step 1: Compute the continuation value at each state point.
- Step 2: Make the exercise decision, based on the greater of immediate exercise $h(\mathbf{S}_t)$ or continuation value $Q_t(\mathbf{S}_t)$.
- Step 3: Regress the option value obtained over $g(\mathbf{S}_t) = \max(S_t^1, S_t^2)$, to be used in the previous exercise time step (as one moves backwards in time) to compute the continuation value.
- Step 4: In the previous exercise time step, compute the transition probability from each state point to the $g(\cdot)$- space in the next time step, i.e. $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$.
- Step 5: Compute the continuation value $\widehat{Q}_{t_m}(\mathbf{S}_{t_m})$ and the option value $\widehat{V}_{t_m}(\mathbf{S}_{t_m})$ using Equation (2.4).

Focusing on the example, Figure 2.3 shows that at time t_2 the option values $V_{t_2}(\mathbf{S}_{t_2} = (14,2))$ and $V_{t_2}(\mathbf{S}_{t_2} = (2,14))$ are 14 and $V_{t_2}(\mathbf{S}_{t_2} = (4,2))$ is 4. On regressing these values over $\max(S_t^1, S_t^2)$ gives $\widehat{Z}_{t_2}(g(\mathbf{S}_{t_2}) = 14, \mathbf{S}_{t_0}) = 14$ and $\widehat{Z}_{t_2}(g(\mathbf{S}_{t_2}) = 4, \mathbf{S}_{t_0}) = 4$ are obtained, as shown in Figure 2.4. Moving to exercise time step t_1 first the transition probability $\mathbb{P}[g(\mathbf{S}_{t_2})|\mathbf{S}_{t_1}]$, is computed. In the present example the state $\mathbf{S}_{t_1} = (8,8)$ transitions to $g(\mathbf{S}_{t_2}) = 14$ with probability 1. Similarly, the conditional transition probability for $\mathbf{S}_{t_1} = (8,4)$ equals $\mathbb{P}(g(\mathbf{S}_{t_2}) = 4|\mathbf{S}_{t_1} = (8,4)) = 1$. Together with these conditional transition probabilities and the approximation of the option values at

t_2, the continuation values for the state points at t_1 are computed. The continuation value at $\mathbf{S}_{t_1} = (8,8)$ equals 14, computed by:

$$\widehat{Q}_{t_1}(\mathbf{S}_{t_1}) = \sum_x \widehat{Z}_{t_2}(g(\mathbf{S}_{t_2}) = x, \mathbf{S}_{t_0}) \cdot \mathbb{P}(g(\mathbf{S}_{t_2}) = x | \mathbf{S}_{t_1} = (8,8)).$$

The continuation value at $\mathbf{S}_{t_1} = (8,4)$ is 4, as:

$$\widehat{Q}_{t_1}(\mathbf{S}_{t_1}) = \sum_x \widehat{Z}_{t_2}(g(\mathbf{S}_{t_2}) = x, \mathbf{S}_{t_0}) \cdot \mathbb{P}(g(\mathbf{S}_{t_2}) = x | \mathbf{S}_{t_1} = (8,4)).$$

Figure 2.5 shows that the option value at $\mathbf{S}_{t_1}$ is the *maximum* of immediate exercise and continuation, i.e. $\max(8,14)$ for $\mathbf{S}_{t_1} = (8,8)$ and $\max(8,4)$ for $\mathbf{S}_{t_1} = (8,4)$. Thus it is optimal to exercise in state $\mathbf{S}_{t_1} = (8,4)$ and to continue in the state $\mathbf{S}_{t_1} = (8,8)$. On regressing these values over $\max(S^1_{t_1}, S^2_{t_1})$, $\widehat{V}_{t_1}(g(\mathbf{S}_{t_1}) = 8)$ value 11 is obtained, as shown in Figure 2.6. Finally, for time step t_0 state $(8,6)$ evolves to $g(\mathbf{S}_{t_1}) = 8$ with probability 1. Therefore, the conditional continuation value is 11, i.e.

$$\sum_x \widehat{Z}_{t_1}(g(\mathbf{S}_{t_1}) = x, \mathbf{S}_{t_0}) \cdot \mathbb{P}(g(\mathbf{S}_{t_1}) = x | \mathbf{S}_{t_0} = (8,6)) = 11,$$

and the option value $\widehat{V}_{t_0}(\mathbf{S}_{t_0} = (8,6)) = \max(8,11)$, which gives the correct value.

Although this example is over simplified, it gives a basic understanding of the approach. Figure 2.7 plots the shape of typical exercise regions for a Bermudan call option on the max of two underlying assets obtained using SGM. The figures are in agreement with those deduced by Broadie and Detemple (1996) [16]. Interestingly, it can be seen, as was found by Broadie *et al.* that, prior to maturity exercise is not optimal when the prices of the underlying assets are equal.

Computing the Continuation Value

The continuation value for grid point $\mathbf{S}_{t_m}$ is the discounted conditional expectation of the option values in the next time step t_{m+1} given $\mathbf{S}_{t_m}$. This is given by,

$$Q_{t_m}(\mathbf{S}_{t_m}) = D_{t_m} \mathbb{E}[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}) | \mathbf{S}_{t_m}].$$

As mentioned in Section 2.2 we first approximate the conditional expectation of the option values at t_{m+1} given $g(\mathbf{S}_{t_{m+1}})$ as a polynomial function of $g(\mathbf{S}_{t_{m+1}})$, Equation (2.7). The continuation value can then be approximated using iterated conditioning as

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = D_{t_m} \mathbb{E}\left[\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0}) | \mathbf{S}_{t_m}\right]. \tag{2.9}$$

Here $\widehat{Z}_{t_{m+1}}$ is a polynomial function of the adapted process $g(\mathbf{S}_{t_{m+1}})$ and hence the conditional probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}}) | \mathbf{S}_{t_m})$ needs to be determined in order to compute its expectation. Using (2.7), Equation (2.9) can be written as

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = D_{t_m} \int_{\mathbf{S}_{t_{m+1}} \in \mathscr{R}^d} \left(\sum_{k=0}^{K-1} \alpha_{t_m}(k) \phi_k(g(\mathbf{S}_{t_{m+1}})) \right) d\mathbb{P}(g(\mathbf{S}_{t_{m+1}}) | \mathbf{S}_{t_m}). \tag{2.10}$$

There are three possibilities for computing $\mathbb{P}[g(\mathbf{S}_{t_{m+1}}) | \mathbf{S}_{t_m}]$:

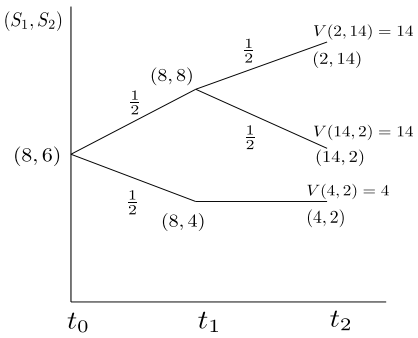

Figure 2.3: Step I: Compute the option values at t_2 as function of $(S^1_{t_2}, S^2_{t_2})$

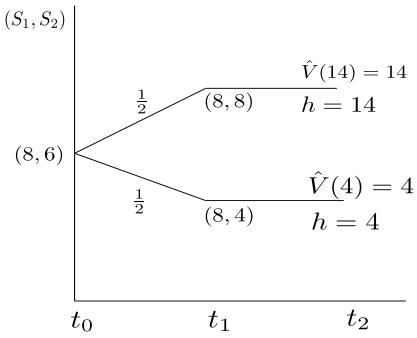

Figure 2.4: Step II: Map the option prices to $\max(S^1_{t_2}, S^2_{t_2})$

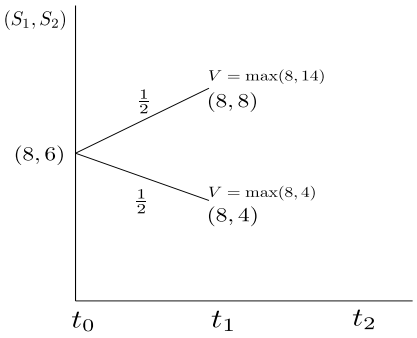

Figure 2.5: Step III: Compute the option values at t_1 as function of $(S^1_{t_1}, S^2_{t_1})$

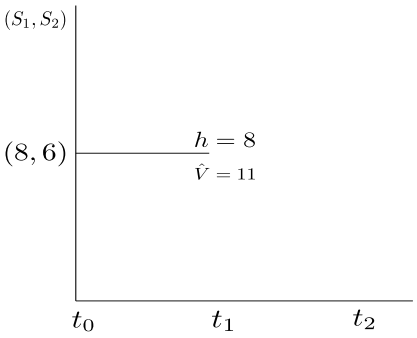

Figure 2.6: Step IV: Map the option price to $\max(S^1_{t_1}, S^2_{t_1})$

1. The exact transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ is known, for example for a call or put on a single asset in the Black-Scholes framework, a call or put on the geometric mean of d assets.

2. The transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ is unknown, however, the moments of the distribution are known, for example for a call or put on the *Max* or *Min* of d assets in the Black-Scholes framework.

3. The transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ and its moments are unknown.

Case 1 is the trivial case where the density function is already known. This case can also be handled efficiently by Fourier techniques, particularly when the conditional density function is not known but when the characteristic function (the Fourier transform of the conditional density) is (Fang and Oosterlee (2008) [29]). Case 3 can be reduced to Case 2, by computing the moments with the help of Monte Carlo sub-simulations. For each grid point at time step t_m, generate sub-paths until time

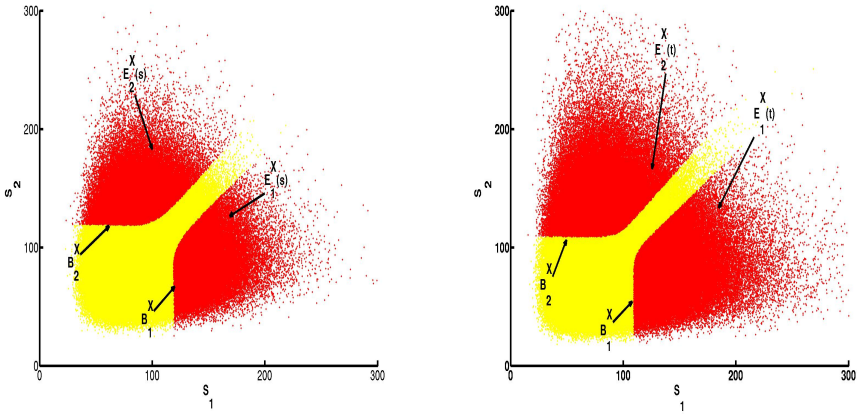

$t < T$ $t < s < T$

Figure 2.7: Exercise regions for a max call option

t_{m+1} and compute the first four central moments, for the distribution of $g(S_{t_{m+1}})$, conditioned on $\mathbf{S}_{t_m}$. The computational effort required for such a sub-simulation is of order $O(N_G N_S)$, where N_G is the number of grid points and N_S is the number of sub-paths simulated. In the examples considered, when sub-simulation was required, the computational time was a few minutes. The computational time can further be reduced by using GPUs and generating sub-paths for a group of nearest neighbour grid points, rather than for each one of them.

Once these moments for $g(\mathbf{S}_{t_{m+1}})$ corresponding to the grid points at t_m are known, the conditional density function $f(x|\mathbf{S}_{t_m})$ is approximated using the *Gram Charlier series* (See Kendall and Stuart (1969) [48]). Given the moments of a distribution, the Gram Charlier series approximates the density function $f(x)$ as,

$$\widehat{f}(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(x-\mu)^2}{2\sigma^2}\right]\left[1 + \frac{\kappa_3}{3!\sigma^3} H_3\left(\frac{x-\mu}{\sigma}\right) + \frac{\kappa_4}{4!\sigma^4} H_4\left(\frac{x-\mu}{\sigma}\right)\right], \qquad (2.11)$$

where $H_3(x) = x^3 - 3x$ and $H_4(x) = x^4 - 6x^2 + 3$ are Hermite polynomials. $\kappa_1 = \mu$, $\kappa_2 = \sigma^2$, $\kappa_3 = \mu_3$, $\kappa_4 = \mu_4 - 3\mu_2^2$ are the first four cumulants. More details about computing the probability density function are given in the specific examples in the sections to follow.

Convergence of Gram Charlier series

The convergence of the Gram Charlier series has been discussed by Milne (1929) [60]. If a distribution satisfies the conditions given by Equations (2.16) and (2.17) then Milne shows the order of convergence for Gram Charlier series approximation of the distribution is $O(n^{-\frac{1}{2}})$, where n is number of terms in the series expansion. Here we give some numerical results to show the effect of,

- error in the moment estimates in the case of sub-simulation,
- non-random error in the Gram Charlier series approximation,

on the SGM estimator.

The effect of error in the moment estimates from sub-simulation is illustrated in Figure 2.8. It plots the standard error for the direct estimator when an increasing number of sub-paths are used. We plot the standard error vs $\frac{1}{\sqrt{N_S}}$, where N_S is the number of paths used in the sub-simulation. When the exact values of the higher moments are used (as can be computed using the Clark algorithm), the standard error should be the intercept with the Y-axis of the fitted function (as it corresponds to the case where $N_S \to \infty$). We find that this is indeed the standard error for the direct SGM estimator when we use the exact moments from the Clark algorithm. Also we find that the mean of the direct SGM estimator and SGM lower bound values obtained by sub-simulation are close to those obtained using the exact moments.

As the error from the Gram Charlier series is independent of the regression error, we look at the case of a European option price for a max option on three assets. In this case the error in the approximation of $\mathbb{E}\left[V(T,S_T)|g(S_T)\right]$ is zero. Then the error in the option price is only due to the Gram Charlier series. We compare the results with those from Boyle (1990) [13] as reference values.

Table 2.1 gives the max European call option values for a 3-d case, when the first two, three and four moments are used in the Gram Charlier series. We find that the error due to exclusion of higher moments while approximating the Gram Charlier series in the case of max option is significant only when the volatilities of the underlying assets are not the same.

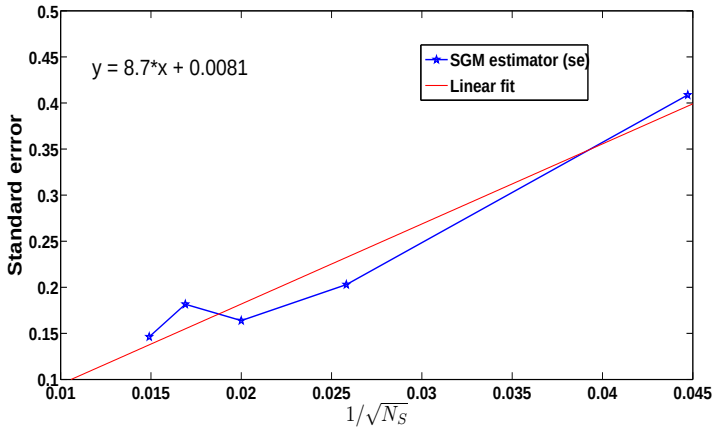

Figure 2.8: Standard error vs $\frac{1}{\sqrt{N_S}}$ for a Bermudan max call option with 10 equally spaced exercise opportunities and 30000 grid points. The parameters are same as in Table 2.3 for 2 assets and S_0 = [100, 100].

Strike	Boyle	2 moments GC	3 moments GC	4 moments GC
30	16.703	16.705	16.700	16.705
40	9.235	9.249	9.251	9.237
50	4.438	4.375	4.458	4.439

Table 2.1: European call option on max of 3 underlying assets: The results are compared with Boyle (1990). The parameters are: $S_0 = [40, 40, 40]$, $r = 10\%$, $\rho = 0.9$, $T = 1$, $\sigma = [25, 30, 35]$ %.

Need for Peripheral Paths

It is noticed that in high-dimensional problems the exercise policy obtained is better if additional paths are generated from points on the periphery of source point $\mathbf{S}_{t_0}$. This idea is not new and was originally proposed by N.S. Rasmussen (2005) [68] as an improvement for the LSM, which he calls *initial state dispersion* where instead of using the original initial state $\mathbf{S}_{t_0}$ for generating the state variables one starts with some fictitious initial time point $-T_D < 0$ and the original state for generating the state variables. More recently, K.H. Kan *et al.* (2010) [46] propose a scheme to disperse the points around the initial source point without starting from a fictitious initial time point. In our examples however two additional point sources around the initial point are used and an equal number of paths from these three source points are generated.

Lower Bound Values

The solution from the SGM can be validated by computing the lower bound on the option price, using the exercise policy obtained from it. To compute the lower bound on the option price, we simulate a number of sample paths (fresh set of paths should be used) originating from $\mathbf{S}_{t_0}$ using the same discretization scheme. The continuation value at the new grid points is then obtained using

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = D_{t_m}\mathbb{E}\left[\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})|\mathbf{S}_{t_m}\right],$$

where the functional approximations of the conditional option values $\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})$ are obtained by the SGM algorithm. For each sample path, the first exercise period t_m, if it exists, is determined for which $h(\mathbf{S}_{t_m}) \geq \widehat{Q}_{t_m}(\mathbf{S}_{t_m})$. The option is then exercised and its discounted pay-off is given by $h(\mathbf{S}_{t_m})/B_{t_m}$. The lower bound on the option price is then obtained as

$$\underline{V_0} = \mathbb{E}\left[\frac{h(\mathbf{S}_{\tilde{\tau}})}{B_{\tilde{\tau}}}\right], \tag{2.12}$$

where $\tilde{\tau} = \min\{t \in [0, T] : \widehat{Q}_t(\mathbf{S}_t) \leq h(\mathbf{S}_t)\}$. The option value obtained by following any exercise strategy is dominated by the optimal strategy. In other words, as the option value is obtained by following a stopping rule $\tilde{\tau}$ it gives a lower bound on the true price (see Andersen and Broadie, 2004 [2]).

Algorithm

The SGM algorithm is briefly summarized as:

- Step I: Generate N sample paths $\{\mathbf{S}_{t_0}, \ldots, \mathbf{S}_{t_M}\}$, where $[t_0 = 0, \ldots, t_M = T]$ and $\mathbf{S}_{t_m} \in \mathscr{R}^d$, all starting from initial state $\mathbf{S}_{t_0}$. The paths are discretized in time using some discretization scheme (e.g. *Euler's discretization scheme*). The N asset prices $\mathbf{S}_{t_m}$ represent the grid points at t_m;
- Step II: Compute the option value for grid points at $t_M = T$ as

$$V_T(\mathbf{S}_T) = max(h(\mathbf{S}_T), 0);$$

- Step III: Compute the approximate functional form,

$$\widehat{Z}_T\left(g(\mathbf{S}_T), \mathbf{S}_{t_0}\right) = \mathbb{E}\left[\widehat{V}_T(\mathbf{S}_T)|g(\mathbf{S}_T), \mathbf{S}_{t_0}\right],$$

  by *regressing* the option value at the grid points over polynomial basis functions of $g(\mathbf{S}_T)$;
- Step IV: Perform the following steps for each exercise time t_m moving backwards in time, starting from t_{M-1} until one reaches t_0 to obtain the direct SGM estimator value $V_{t_0}(\mathbf{S}_{t_0})$:
  - Compute the continuation value for grid points at t_m using the functional approximation of $\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})$,

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = D_{t_m}\mathbb{E}\left[\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})|\mathbf{S}_{t_m}\right];$$

  - Compute the option value for grid points at t_m as

$$\widehat{V}_{t_m}(\mathbf{S}_{t_m}) = \max(h(\mathbf{S}_{t_m}), \widehat{Q}_{t_m}(\mathbf{S}_{t_m}));$$

  - Compute the functional approximation for the conditional expectation, *i.e.*

$$\widehat{Z}_{t_m}(g(\mathbf{S}_{t_m}), \mathbf{S}_{t_0}) = \mathbb{E}\left[\widehat{V}_{t_m}(\mathbf{S}_{t_m})|g(\mathbf{S}_{t_m}), \mathbf{S}_{t_0}\right]$$

    by *regressing* the option value obtained at each grid point in t_m over a set of polynomial basis function of $g(\mathbf{S}_{t_m})$;
  - Go to the previous time step ($m \Rightarrow m-1$).
- Step V: Using the exercise strategy obtained while computing the direct SGM estimator. For each path (from a set of new paths) determine the earliest time to exercise $\tilde{\tau} = \min\{t \in [0, T] : \widehat{Q}_t(\mathbf{S}_t) \leq h(\mathbf{S}_t)\}$. Obtain the lower bound option value as $\mathbb{E}\left[\frac{h(\mathbf{S}_{\tilde{\tau}})}{B_{\tilde{\tau}}}\right]$.

2.3 Error Analysis for the Single Asset Case

A basic error analysis is performed for the single asset case. It can be readily seen that for a single asset case for the choice of $g(x) = x$,

$$Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0}) = Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_m}) = V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}).$$

SGM has two main sources of error in the penultimate exercise opportunity, *i.e.* at $t_{m=M-1}$. They are,

- $\epsilon^z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}))$: error in the approximation of

$$Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0}) = \mathbb{E}\left[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}) | g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0}\right],$$

- $\epsilon^f_{t_m}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$: error in the approximation of the transition density function,

$$f(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$$

The approximation of the continuation value at t_m, is given by,

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = \int (Z_{t_{m+1}}(x, \mathbf{S}_{t_0}) + \epsilon^z_{t_{m+1}}(x))(f(x|\mathbf{S}_{t_m}) + \epsilon^f_{t_m}(x|\mathbf{S}_{t_m}))dx. \tag{2.13}$$

The error in the estimation of the continuation value, $\epsilon^Q_{t_m}(\mathbf{S}_{t_m})$, thus comes from errors in the approximation of $Z_{t_{m+1}}(\mathbf{S}_{t_{m+1}}, \mathbf{S}_{t_0})$ and the transition density function $f(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$. The error, $\epsilon^Q_{t_m}(\mathbf{S}_{t_m})$, can be split into an error due to approximation of the transition density function, $\epsilon^{Q_f}_{t_m}(\mathbf{S}_{t_m})$, and an error due to the approximation of $Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})$, i.e. $\epsilon^{Q_z}_{t_m}(\mathbf{S}_{t_m})$, i.e. ,

$$\begin{aligned}
\epsilon^Q_{t_m}(\mathbf{S}_{t_m}) &\approx \int \epsilon^f_{t_m}(x) Z_{t_{m+1}}(x, \mathbf{S}_{t_0})dx + \int \epsilon^z_{t_{m+1}}(x) f(x|\mathbf{S}_{t_m})dx \\
&\leq \int |\epsilon^f_{t_m}(x)| Z_{t_{m+1}}(x, \mathbf{S}_{t_0})dx + \int |\epsilon^z_{t_{m+1}}(x)| f(x|\mathbf{S}_{t_m})dx \qquad (2.14) \\
&= \epsilon^{Q_f}_{t_m}(\mathbf{S}_{t_m}) + \epsilon^{Q_z}_{t_m}(\mathbf{S}_{t_m}). \qquad (2.15)
\end{aligned}$$

It is now shown that these two errors are bounded.

Error Due to Gram Charlier Approximation

Milne (1929) [60] showed that if $f(x)$ satisfies a condition of the form

$$|e^{\frac{x_1^2}{4}} f(x_1) - e^{\frac{x_2^2}{4}} f(x_2)| < L|x_1 - x_2|, \tag{2.16}$$

and if

$$|xe^{\frac{x^2}{4}} f(x)| < L, \tag{2.17}$$

with L constant, then the error of a Gram-Charlier series as in (2.11) with n terms is bounded by

$$|f(x) - f_n(x)| := |\epsilon^f(x)| < BLn^{-\frac{1}{2}} e^{-\frac{x^2}{4}}, \tag{2.18}$$

where B is a constant independent of n. Assuming that the conditions above are satisfied the error in the continuation value due to the Gram Charlier approximation can be bounded by

$$\epsilon_{t_m}^{Q_f}(\mathbf{S}_{t_m}) < BLn^{-\frac{1}{2}} \int e^{-\frac{x^2}{4}} Z_{t_{m+1}}(x, \mathbf{S}_{t_0}) dx. \tag{2.19}$$

Error Due to Parametrization of Option price

$Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})$ is approximated by piecewise interpolation. If a single high-degree polynomial is used for regression it can lead to significant errors if one of the derivatives of $Z_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}), \mathbf{S}_{t_0})$ is discontinuous. A robust alternative is to replace the single high degree polynomial for regression in $[x_0, x_n]$, here $x_0 = \min_{\mathbf{S}_{t_{m+1}}}(g(\mathbf{S}_{t_{m+1}}))$ and $x_n = \max_{\mathbf{S}_{t_{m+1}}}(g(\mathbf{S}_{t_{m+1}}))$, by several low-degree polynomials by appropriately dividing the regression domain $[x_0, x_n]$. An extreme case of this would be to use a linear polynomial to interpolate between adjacent data points. In such a case the maximum error due to regression is bounded by

$$\max_{x\in[x_0,x_n]} |Z_{t_{m+1}}(x, \mathbf{S}_{t_0}) - \widehat{Z}_{t_{m+1}}(x, \mathbf{S}_{t_0})| = |\epsilon^z|_{\max} \leq \max_{x\in[x_0,x_n]} \frac{1}{2} |\frac{\partial^2 Z_{t_{m+1}}(x, \mathbf{S}_{t_0})}{\partial x^2}| \Delta^2, \tag{2.20}$$

where Δ denotes the largest space between interpolation points.

In practice, however, dividing the domain in upto six regions with four polynomial basis functions for each region already gives a small regression error. The break points for dividing the domain $[x_0, x_n]$ are chosen as the early-exercise point and the critical points for $\frac{\partial^2 Z_{t_{m+1}}(x, \mathbf{S}_{t_0})}{\partial x^2}$, Figure 2.9 compares the maximum and mean regression error with different numbers of pieces (keeping the number of grid points constant) and with different numbers of grid points (keeping the number of pieces constant) for a call option on a single asset. It can be seen that for the same number of grid points, significantly smaller errors in regression can be obtained using more partitions.

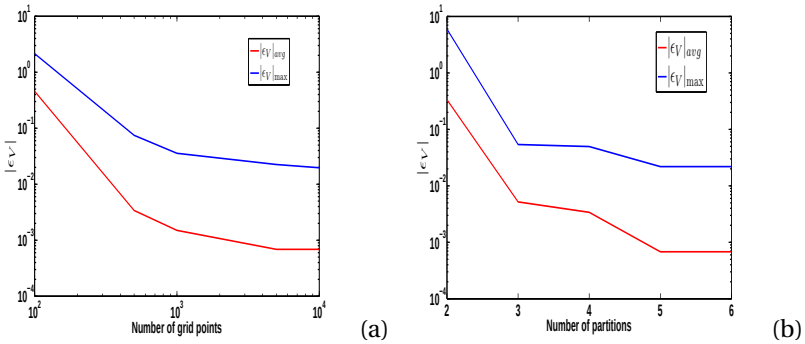

(a) Grid Points (b) Partitions

Figure 2.9: Maximum and average squared residual errors due to parametrization of the option price when, (a) when the number of segments in the piecewise regression is constant = 6, (b) the number of grid points used in the regression is constant = 10,000.

Assuming the conditions above are satisfied the error in the continuation value due to parametrization of the option price is then bounded by

$$\epsilon^{Q_z}_{t_m}(\mathbf{S}_{t_m}) \leq |\epsilon^z_{t_{m+1}}|_{\max} \int f(x|\mathbf{S}_{t_m})dx \tag{2.21}$$

Under the assumption that the conditions for convergence of Gram Charlier series expansions are satisfied and a large number of local regression functions are used, the error in the continuation value is bounded by

$$\epsilon^{Q}_{t_m}(\mathbf{S}_{t_m}) \leq BLn^{-\frac{1}{2}} \int e^{-\frac{x^2}{4}} Z_{t_{m+1}}(x,\mathbf{S}_{t_0})dx + |\epsilon^z_{t_{m+1}}|_{\max} \int f(x|\mathbf{S}_{t_m})dx. \tag{2.22}$$

Here it is assumed that $\int e^{-\frac{x^2}{4}} Z_{t_{m+1}}(x,\mathbf{S}_{t_0})dx$ is bounded.

Error Due to Recursion

From (2.22) the error in continuation value at t_m is bounded. At t_m the error in the option price $V_{t_m}(\mathbf{S}_{t_m})$ can be determined using

$$\begin{aligned} \widehat{V}_{t_m}(\mathbf{S}_{t_m}) &= \max(Q_{t_m}(\mathbf{S}_{t_m}) + \epsilon^Q_{t_m}(\mathbf{S}_{t_m}), h(\mathbf{S}_{t_m})) \\ &\leq \max(Q_{t_m}(\mathbf{S}_{t_m}), h(\mathbf{S}_{t_m})) + |\epsilon^Q_{t_m}(\mathbf{S}_{t_m})|. \end{aligned} \tag{2.23}$$

The continuation value at t_{m-1} will have an error described by

$$\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}) \leq \int (Z_{t_m}(x,\mathbf{S}_{t_0}) + |\epsilon^z_{t_m}(x)| + |\epsilon^Q_{t_m}(x)|)(f(x|\mathbf{S}_{t_{m-1}}) + \epsilon^f_{t_m}(x))dx. \tag{2.24}$$

The additional term in Equation (2.24) when compared to (2.13), is the error due to recursion, ϵ_R :

$$\epsilon_R \leq \int |\epsilon^Q_{t_m}(x)|(f(x|\mathbf{S}_{t_{m-1}}) + \epsilon^f_{t_{m-1}}(x))dx,$$

which is bounded by

$$\epsilon_R \leq \max_{\mathbf{S}_{t_m}}(|\epsilon^Q_{t_m}(\mathbf{S}_{t_m})|).$$

It can be shown that the error due to recursion at time step t_0 is bounded by

$$\epsilon_{R_0} \leq \sum_m \max_{\mathbf{S}_{t_m}}(|\epsilon^Q_{t_m}(\mathbf{S}_{t_m})|).$$

Numerical Results for Bermudan Put on a Single Asset

The error analysis for a put on a single asset is illustrated using numerical results, where the risk-neutral asset price follows the stochastic differential equation

$$dS_t = rS_t dt + \sigma S_t dW_t, \tag{2.25}$$

r being the continuously compounded risk-free interest rate, σ the annualized volatility. Here r and σ are assumed to be constant. W_t is the standard Brownian motion. It is assumed that the option is exercisable a finite number of times (M) per year, at a strike price of $\mathscr{K}$, up-to and including the final expiration time T. N sample paths, $\{\mathbf{S}_{t_0}, \ldots, S_{t_M}\}$, are generated using the closed form solution for SDE (2.25). The asset values $\mathbf{S}_{t_m}$ represent the grid points at t_m.

Parametrization of the Option Value for a Single Asset

The option price at any time t_m prior to the expiration time T is given by

$$V_{t_m}(\mathbf{S}_{t_m}) = \max(h(\mathbf{S}_{t_m}), Q_{t_m}(\mathbf{S}_{t_m})).$$

To compute the functional approximation of the option value at time t_m, the option values obtained at the grid points are regressed on polynomial basis functions of $g(\mathbf{S}_{t_m}) = \mathbf{S}_{t_m}$. A piecewise least squares regression with one of the break points at $\mathscr{X}^*_{t_m} = \mathbf{S}^*_{t_m}$, where $\mathbf{S}^*_{t_m}$ is the early-exercise point, is performed. For better approximation the continuation region can be further divided into pieces with break

points selected at the critical points for $Z''_{t_m}(\mathbf{S}_{t_m}, \mathbf{S}_{t_0})$. For the two segment case, we regress the option value as

$$\widehat{Z}_{t_m}(\mathbf{S}_{t_m}, \mathbf{S}_{t_0}) = \mathbf{1}_{\{g(\mathbf{S}_{t_m}) < \mathscr{X}^*_{t_m}\}} \sum_{k=0}^{K-1} \alpha^1_{t_m}(k)(|g(\mathbf{S}_{t_m})|)^k + \mathbf{1}_{\{g(\mathbf{S}_{t_m}) \geq \mathscr{X}^*_{t_m}\}} \sum_{k=0}^{K-1} \alpha^2_{t_m}(k)(|g(\mathbf{S}_{t_m})|)^k, \tag{2.26}$$

with the coefficients $\alpha^1_{t_m}$ and $\alpha^2_{t_m}$ chosen so that residuals r_1 and r_2 are minimized,

$$r_1 = \min_{\alpha^1_{t_m}} \left(\mathbf{1}_{\{g(\mathbf{S}_{t_m}) < \mathscr{X}^*_{t_m}\}} \sum |V_{t_m}(\mathbf{S}_{t_m}) - \widehat{Z}_{t_m}(g(\mathbf{S}_{t_m}), \mathbf{S}_{t_0})|^2 \right),$$

$$r_2 = \min_{\alpha^2_{t_m}} \left(\mathbf{1}_{\{g(\mathbf{S}_{t_m}) \geq \mathscr{X}^*_{t_m}\}} \sum |V_{t_m}(\mathbf{S}_{t_m}) - \widehat{Z}_{t_m}(g(\mathbf{S}_{t_m}), \mathbf{S}_{t_0})|^2 \right).$$

The first four polynomials (including the constant) are chosen as basis functions. Increasing the number of basis functions does not significantly improve the approximation, however increasing the number of pieces does improve the solution.

Continuation Value for the Single Asset Case

In order to compute the continuation value for the grid points at t_m using Equation (2.9) the transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ is required. For a single asset following a stochastic process given by Equation (2.25), the conditional transition density function is given by

$$\mathbb{P}(g(\mathbf{S}_{t_{m+1}}) = x|\mathbf{S}_{t_m}) = \mathbf{S}_{t_m} e^{\left((r-\frac{\sigma^2}{2})\Delta t + \sigma\sqrt{\Delta t}\, Y\right)} \mathbb{P}(Y = x^*), \tag{2.27}$$

where $\Delta t = t_{m+1} - t_m$, $Y \sim \mathscr{N}(0,1)$ and

$$x^* := \frac{1}{\sigma\sqrt{\Delta t}} \left[\log\left(\frac{x}{\mathbf{S}_{t_m}}\right) - (r - \frac{\sigma^2}{2})\Delta t \right].$$

Equation (2.10) can then be written as

$$\begin{aligned} \widehat{Q}_{t_m}(\mathbf{S}_{t_m}) \quad = \quad & D_{t_m} \left(\int_{-\infty}^{K^*} \sum_{k=0}^{K-1} \alpha^1_{t_m}(k)(f(Y))^k d\mathbb{P}(Y) \right. \\ & \left. + \int_{K^*}^{\infty} \sum_{k=0}^{K-1} \alpha^2_{t_m}(k)(f(Y))^k d\mathbb{P}(Y) \right), \end{aligned} \tag{2.28}$$

where

$$K^* = \frac{1}{\sigma\sqrt{\Delta t}} \left[\log\left(\frac{|\mathscr{X}^*_{t_{m+1}}|}{\mathbf{S}_{t_m}}\right) - (r - \frac{\sigma^2}{2})\Delta t \right],$$

$$f(Y) = \mathbf{S}_{t_m} e^{\left((r-\frac{\sigma^2}{2})\Delta t + \sigma\sqrt{\Delta t}\, Y\right)}$$

$$d\mathbb{P}(Y) = \frac{1}{\sqrt{2\pi}} e^{-\frac{Y^2}{2}}\, dY.$$

Solving Equation (2.28) gives the continuation value at each grid point as

$$\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) = D_{t_m}\left[\sum_{k=0}^{K-1}\varphi_{t_m}^k\left((\alpha_{t_m}^1(k)-\alpha_{t_m}^2(k))\Phi\left(K^*-k\sigma\sqrt{\Delta t}\right)+\alpha_{t_m}^2(k)\right)\right], \quad (2.29)$$

where

$$\varphi_{t_m}^k = \left(\mathbf{S}_{t_m}^k e^{k\left(\left(r-\frac{\sigma^2}{2}\right)+\frac{k}{2}\sigma^2\right)\Delta t}\right),$$

and

$$\Phi(x) = \frac{1}{2}\left[1+erf\left(\frac{x}{\sqrt{2}}\right)\right].$$

In order to compute the value of $\mathscr{X}_{t_m}^* = |g((\mathbf{S}_{t_m})|$, the non-linear equation

$$h(\mathbf{S}_{t_m}) = Q_{t_m}(\mathbf{S}_{t_m}), \quad (2.30)$$

needs to be solved, where the value of $Q_{t_m}(\mathbf{S}_{t_m})$, is obtained from Equation (2.29). The value of $\mathscr{X}_{t_m}^*$ can be approximated as,

$$\mathscr{X}_{t_m}^* = \max\left(|h(\mathbf{S}_{t_m})|\mathbf{1}_{h(\mathbf{S}_{t_m})\geq Q_{t_m}(\mathbf{S}_{t_m})}\right),$$

i.e., find the maximum value of the asset price for the grid points lying in the *early-exercise* region, or alternatively the minimum value of the asset price for grid points in the *continuation* region.

Results for Single Asset Put option

To illustrate the results, Table 2.2 reports the value of the early-exercise option implied by both the COS method and SGM. The COS method with $N = 2^{10}$ terms in the Fourier expansion is used to generate our reference values. The lower bound values, which are obtained by following the exercise policy from SGM on a fresh set of paths, are sometimes greater than the true option price. The lower bound values are taken as the mean of 30 simulation results. True lower bound values can be obtained by computing the mean over a large number of simulation results.

The SGM estimates are based on 10000 (5000 plus and 5000 antithetic) paths using 50 exercise points per year, while the LSM estimates are based on 100000 (50000 plus and 50000 antithetic) paths. Figure 2.10 compares the SGM direct estimator with the true option price for different numbers of grid points. Figure 2.11 compares the lower bound values obtained from SGM with the lower bound from the LSM algorithm for different numbers of paths. The exercise policy obtained using SGM is better and more stable compared to the one obtained using LSM, as can be deduced from the standard errors for the lower bounds for the two algorithms.The direct estimator value converges fast to the reference price as the number of partitions and grid points increase. The standard errors of the direct estimator are small compared to those of the SGM lower bound values and much lower than those of the LSM values.

The time taken for each simulation is a few seconds on a system with Intel(R) Duo-Core 2.13 GHz processors and 2 GB RAM.

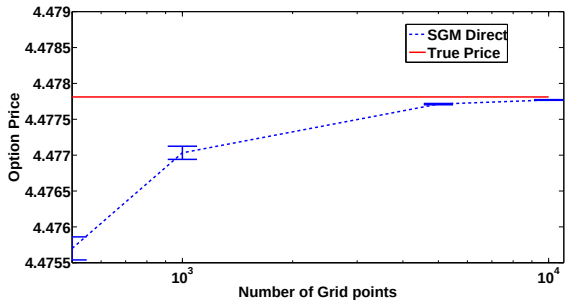

Figure 2.10: SGM direct estimator with confidence interval for different numbers of grid points. The regression is performed on 6 different pieces.

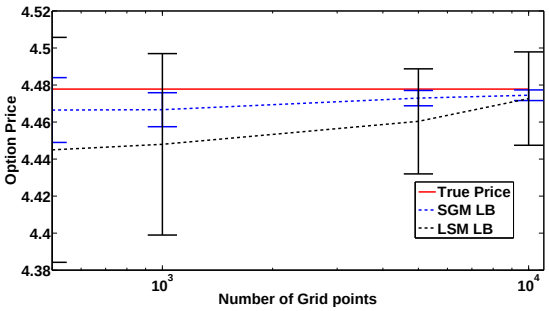

Figure 2.11: Comparison between lower bounds and confidence interval obtained using the exercise policy from SGM and LSM for different numbers of grid points (paths for latter).

2.4 Numerical results in high dimensions

This section illustrates the methodology by pricing Bermudan options on the max of two, three and five assets, and a basket option on an arithmetic mean of four and five assets. The underlying assets are assumed to follow the standard single and multi-asset Black-Scholes model (geometric Brownian motion, GBM).

Bermudan Call on Maximum of *d* Assets

A Bermudan max-option is a discretely-exercisable option on multiple underlying assets whose pay-off depends on the maximum among all asset prices. It is assumed that the asset prices follow correlated geometric Brownian motion processes, i.e.,

$$\frac{dS_t^\delta}{S_t^\delta} = (r - q_\delta)dt + \sigma_\delta dW_t^\delta, \tag{2.31}$$

where each asset pays a dividend at a continuous rate of q_δ. $W_t^\delta, \delta = 1,\ldots,d$, are standard Brownian motions and the instantaneous correlation between W_t^i and

S_0	σ	T	COS Method Bermudan	SGM Lower Bound(s.e)	SGM Direct Estimator(s.e)	LSM (s.e)	Closed form European
36	0.4	2	8.508	8.512 (0.56)	8.509 (0.010)	8.488 (0.51)	7.700
38	0.4	2	7.670	7.665 (0.53)	7.670 (0.011)	7.669 (0.50)	6.979
40	0.4	2	6.920	6.913 (0.59)	6.919 (0.011)	6.921 (0.55)	6.326
42	0.4	2	6.248	6.252 (0.59)	6.246 (0.013)	6.243 (0.51)	5.736
44	0.4	2	5.647	5.632 (0.66)	5.642 (0.014)	5.622 (0.51)	5.202

Table 2.2: Comparison of the SGM direct estimator and lower bound values with the LSM and COS methods results for a Bermudan put option on a single asset, where the option is exercisable 50 times per year. The strike price of the put is 40, the short term interest rate is 0.06. The simulation for SGM is based on 10000 (5000 plus 5000 antithetic) paths for the asset price process, and for LSM is based on 100000 (50000 plus and 50000 antithetic) paths. The standard error for the simulation (s.e) is in cents while the option values are in dollars.

W_t^j is ρ_{ij}. It is assumed that the option expires at time T and there are M equally spaced exercise dates in the interval $[0,T]$.

The method starts by generating N sample grid points $(S_{t_m}^1,\ldots,S_{t_m}^d)$ at each time step t_m, using the discretization scheme

$$S_{t_m}^{\delta} = S_{t_{m-1}}^{\delta}\exp\left((r-q_{\delta}-\frac{1}{2}|\sigma_{\delta}|^2)\Delta t+\sum_{1\le k\le d}\sigma_{\delta k}W_{\Delta t}^k\right), 1\le\delta\le d, \tag{2.32}$$

where $\Delta t = t_m - t_{m-1}$. As explained in Section 2.2, for high-dimensional options additional peripheral paths are required to obtain better lower bound values. In the present example additional sample paths from two points around initial source point $\mathbf{S}_{t_0}$, the points selected as $\mathbf{S}_{t_0}e^{0.3\sigma\sqrt{\Delta t}}$ and $\mathbf{S}_{t_0}e^{-0.1\sigma\sqrt{\Delta t}}$, are generated which already significantly improves the lower bound values. The peripheral paths are used only to obtain the exercise-policy from the direct SGM estimator and are not used to obtain the lower bound values. Additional peripheral paths are required because in their absence the regressed function values around peripheral grid points becomes a source of error.

The Clark Algorithm

In order to compute the continuation value for grid points at t_m using Equation (2.9), the transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ is required. For a call on the max of d underlying assets, $(g(\mathbf{S}_{t_{m+1}}) = \max(S_{t_{m+1}}^1,\ldots,S_{t_{m+1}}^d))$, it is difficult to compute the exact transition density function. Like Boyle and Tse (1990) [13], Clark's algorithm is used to approximate the first four moments of this distribution. The approximation of the transition probability density function can be obtained from these moments using the *Gram Charlier expansion* [60].

The Clark algorithm (1961) [22], calculates the first four moments of the maximum of a pair of jointly normal variates as well as the correlation coefficient between the maximum of the pair and a third normal variate. Let X_1 and X_2 have a bivariate normal distribution, with means μ_1 and μ_2, and standard deviations σ_1 and σ_2, respectively. The correlation coefficient between the two is ρ. Y denotes the maximum of (X_1, X_2). Let ν_i denote the ith non-central moment for the distribution of Y, then

$$\nu_1 = \mu_1 \Phi(\alpha) + \mu_2 \Phi(-\alpha) + a\phi(\alpha), \tag{2.33}$$

$$\nu_2 = \left(\mu_1^2 + \sigma_1^2\right)\Phi(\alpha) + \left(\mu_2^2 + \sigma_2^2\right)\Phi(-\alpha) + \left(\mu_1 + \mu_2\right)a\phi(\alpha), \tag{2.34}$$

$$\begin{aligned} \nu_3 \quad = \quad & \left(\mu_1^3 + 3\mu_1\sigma_1^2\right)\Phi(\alpha) + \left(\mu_2^3 + 3\mu_2\sigma_2^2\right)\Phi(-\alpha) \\ & + \left[\left(\mu_1^2 + \mu_1\mu_2 + \mu_2^2\right)a + \left(2\sigma_1^4 + \sigma_1^2\sigma_2^2 + 2\sigma_2^4\right.\right. \\ & \left.\left. -2\sigma_1^3\sigma_2\rho - 2\sigma_1\sigma_2^3\rho - \sigma_1^2\sigma_2^2\rho^2\right)a^{-1}\right]\phi(\alpha), \end{aligned} \tag{2.35}$$

$$\begin{aligned} \nu_4 \quad = \quad & \left(\mu_1^4 + 6\mu_1^2\sigma_1^2 + 3\sigma_1^4\right)\Phi(\alpha) + \left(\mu_2^4 + 6\mu_2^2\sigma_2^2 + 3\sigma_2^4\right)\Phi(-\alpha) \\ & + \{\left(\mu_1^3 + \mu_1^2\mu_2 + \mu_1\mu_2^2 + \mu_2^3\right)a - 3\alpha\left(\sigma_1^4 - \sigma_2^4\right) \\ & + 4\mu_1\sigma_1^3\left[3\left(\frac{\sigma_1 - \sigma_2\rho}{a}\right) - \left(\frac{\sigma_2 - \sigma_1\rho}{a}\right)^3\right] \\ & + 4\mu_2\sigma_2^3\left[3\left(\frac{\sigma_2 - \sigma_1\rho}{a}\right) - \left(\frac{\sigma_1 - \sigma_2\rho}{a}\right)^3\right]\}\phi(\alpha). \end{aligned} \tag{2.36}$$

If X_3 is a random variable with normal distribution, and the correlation coefficients between X_3 and X_1, X_2 are ρ_1, ρ_2, respectively, then the correlation coefficient $\rho_{X_3 Y}$ between X_3 and $Y = \max(X_1, X_2)$ is given by

$$\rho_{X_3 Y} = \left[\sigma_1\rho_1\Phi(\alpha) + \sigma_2\rho_2\Phi(-\alpha)\right] / \left(\nu_2 - \nu_1^2\right)^{\frac{1}{2}}, \tag{2.37}$$

where

$$a^2 = \sigma_1^2 + \sigma_2^2 - 2\sigma_1\sigma_2\rho$$

$$\alpha = \frac{\mu_1 - \mu_2}{a}$$

$$\phi(x) = (2\pi)^{-\frac{1}{2}} \exp\left(-\frac{x^2}{2}\right)$$

$$\Phi(x) = \int_{-\infty}^{x} \phi(t)\, dt.$$

Clark's method can be used to obtain the exact moments of Y and its correlation with X_3, however, as the distribution of Y is not exactly normal, the method can be only used to obtain the *approximation* of the first four moments of the maximum of a set of $d \geq 3$ normal variates. If $X_1, \ldots, X_d$ are the d jointly normal variates, and Y is the maximum of these d variates then by using the recursive scheme

$$Y_i = \max(X_1, X_2, \ldots, X_{i+1}) = \max(Y_{i-1}, X_{i+1})$$

and applying Clark's approximation at each step we can compute the approximation of the first four moments for the distribution of Y. It is easy to deduce how Clark's method can be used to obtain the moments for the *minimum* of d assets as well (see Boyle *et al.* (1990) [13]).

Parametrization of the Option Value for Max Options

In order to compute the functional form of the option value at t_{m+1}, the option values obtained at the grid points are regressed over the polynomial basis functions of $g(\mathbf{S}_{t_{m+1}})$. Piecewise regression, with the break points at $\mathscr{X}^*_{t_{m+1}} = h(\mathbf{S}^*_{t_{m+1}})$, where $h(\mathbf{S}^*_{t_{m+1}}) = Q_{t_{m+1}}(\mathbf{S}^*_{t_{m+1}})$, is used. The regression scheme can be written as,

$$\begin{aligned}\widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}) &= \mathbf{1}_{\{g(\mathbf{S}_{t_{m+1}})<\mathscr{X}^*_{t_{m+1}}\}}\sum_{k=0}^{K-1}\alpha^1_{t_m}(k)(\phi_k(g(\mathbf{S}_{t_{m+1}}))) + \\ &\quad \mathbf{1}_{\{g(\mathbf{S}_{t_{m+1}})\geq\mathscr{X}^*_{t_{m+1}}\}}\sum_{k=0}^{K-1}\alpha^2_{t_m}(k)(\phi_k(g(\mathbf{S}_{t_{m+1}}))), \qquad (2.38)\end{aligned}$$

where ϕ are the basis functions. The coefficients $\alpha^1_{t_m}$ and $\alpha^2_{t_m}$ are chosen such that residuals r_1 and r_2 are minimized,

$$r_1 = \min_{\alpha^1_{t_m}}\left(\mathbf{1}_{\{g(\mathbf{S}_{t_{m+1}})<\mathscr{X}^*_{t_{m+1}}\}}\sum|V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}) - \widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0})|^2\right),$$

$$r_2 = \min_{\alpha^2_{t_m}}\left(\mathbf{1}_{\{g(\mathbf{S}_{t_{m+1}})\geq\mathscr{X}^*_{t_{m+1}}\}}\sum|V_{t_{m+1}}(\mathbf{S}_{t_{m+1}}) - \widehat{Z}_{t_{m+1}}(g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0})|^2\right).$$

A set of four (including the constant) *Hermite polynomial* basis functions of $g(\mathbf{S}_{t_{m+1}})$ are used for regression in our example.

Computing the Continuation Value for Max Options

With $\mathbf{S}_{t_m}$ being a log-normal process given by Equation (2.32), it can be written as,

$$\begin{aligned}\mathbb{P}\left(g(\mathbf{S}_{t_{m+1}}) = X|\mathbf{S}_{t_m}\right) &= \mathbb{P}\left(\max_{1\leq j\leq d}(S^j_{t_{m+1}}) = X|\mathbf{S}_{t_m}\right) \\ &= \mathbb{P}\left(\max_{1\leq j\leq d}(Y^j_{t_{m+1}}) = \log(X)|\mathbf{S}_{t_m}\right), \qquad (2.39)\end{aligned}$$

where $Y^j_{t_{m+1}}, 1\leq j\leq d$ has a multivariate normal distribution. Using Clark's algorithm the first four moments of the random variable

$$Y = \max(Y^1_{t_{m+1}},\ldots,Y^d_{t_{m+1}}),$$

can be obtained. If ${\kappa_m}^1 (1\leq i\leq 4)$ are the first four cumulants of Y then using the *Gram Charlier Expansion* the approximate probability density function of Y can

[1] $\kappa_1 = \mu = \mu'_1$

be written as Equation (2.11). The continuation value given by Equation (2.10) can then be written as,

$$\begin{aligned}\widehat{Q}_{t_m}(\mathbf{S}_{t_m}) &= D_{t_m}\left(\int_{-\infty}^{K^*}\sum_{k=0}^{K-1}\alpha^1_{t_m}(k)\phi_k(e^x)d\mathbb{P}(Y=x|\mathbf{S}_{t_m})\right.\\ &\left.+\int_{K^*}^{\infty}\sum_{k=0}^{M-1}\alpha^2_{t_m}(k)\phi_k(e^x)d\mathbb{P}(Y=x|\mathbf{S}_{t_m})\right),\end{aligned} \tag{2.40}$$

where $K^* = \log(\mathscr{X}^*_{t_m})$.

The term,

$$\int_{-\infty}^{K^*}\sum_{m=0}^{M-1}a_m\Psi_m(e^x)d\mathbb{P}(Y=x|S_{t_{i-1}}),$$

where $\Psi(\cdot)$ is polynomial basis function, can be written as the linear combination of,

$$\int_{-\infty}^{K^*}(e^{mx})d\mathbb{P}(Y=x|S_{t_{i-1}}),$$

$m=(0,\ldots,(M-1))$, and $\mathbb{P}(Y=x|S_{t_{i-1}})$ is

$$\phi\left(\frac{x-\mu}{\sigma}\right)\left[1+\frac{\kappa_3}{3!\sigma^3}H_3\left(\frac{x-\mu}{\sigma}\right)+\frac{\kappa_4}{4!\sigma^4}H_4\left(\frac{x-\mu}{\sigma}\right)\right].$$

Here $\phi(x)$ is,

$$\phi(x)=\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}.$$

We need to compute,

$$\int_{-\infty}^{K^*}(e^{mx})\phi\left(\frac{x-\mu}{\sigma}\right)\left[1+\frac{\kappa_3}{3!\sigma^3}H_3\left(\frac{x-\mu}{\sigma}\right)+\frac{\kappa_4}{4!\sigma^4}H_4\left(\frac{x-\mu}{\sigma}\right)\right]dx. \tag{2.41}$$

Equation (2.41) can be written as,

$$A\int_{-\infty}^{K^*}\phi\left(\frac{x-\theta}{\sigma}\right)\left[1+\frac{\kappa_3}{3!\sigma^3}H_3\left(\frac{x-\mu}{\sigma}\right)+\frac{\kappa_4}{4!\sigma^4}H_4\left(\frac{x-\mu}{\sigma}\right)\right]dx, \tag{2.42}$$

where,

$$A=e^{\left(\mu m+\frac{m^2\sigma^2}{2}\right)},$$

and

$$\theta=\left(\mu+m\sigma^2\right).$$

$$\kappa_2=\sigma^2=\mu'_2-{\mu'_1}^2$$

$$\kappa_3=\mu'_3-3\mu'_2\mu'_1+2{\mu'_1}^3$$

$$\kappa_4=\mu'_4-4\mu'_3\mu'_1-3{\mu'_2}^2+12\mu'_2{\mu'_1}^2-6{\mu'_1}^4$$

where μ'_m is the *i th* non-central moment

This can be written in a form easy to integrate,

$$A \quad \int_{-\infty}^{K^*} \phi\left(\frac{x-\theta}{\sigma}\right)\left[1+\frac{\kappa_3}{3!\sigma^3}\sum_{j=0}^{3}\binom{3}{j}\left(\frac{\theta-\mu}{\sigma}\right)^j H_{3-j}\left(\frac{x-\theta}{\sigma}\right)\right.$$
$$\left.+\frac{\kappa_4}{4!\sigma^4}\sum_{k=0}^{4}\binom{4}{k}\left(\frac{\theta-\mu}{\sigma}\right)^k H_{4-k}\left(\frac{x-\theta}{\sigma}\right)\right] dx, \tag{2.43}$$

using the property,

$$\int_{-\infty}^{x} \phi(y)H_n(y)dy = -\phi(x)H_{n-1}(x).$$

Results for Bermudan Call on Max of several Assets

To illustrate the results Table 2.3 compares the results of a Bermudan max options on 2, 3 and 5 underlying assets. The results reported in Table 2.3 are fairly remarkable given the simplicity of the method. The values obtained by SGM are close to the values reported in the literature. The number of paths required to obtain an accurate exercise policy (as reflected by the lower-bound values) is far less than required to obtain the exercise policy for the reference methods. Also the time for each simulation is less than a minute on a system with Intel(R) Duo-Core 2.13 GHz processors and 2 GB RAM. The number of basis functions required for regression, irrespective of the dimensions of the problem, is upto 4 (including the constant). The exercise policy obtained using SGM, especially when the dimensions of the problem are large, is farther from the optimal policy when compared to those obtained using LSM, as can be inferred from the lower bound values.

Bermudan Put on Arithmetic Mean of *d* Assets

A Bermudan basket option is a discretely-exercisable option on multiple underlying assets whose pay-off depends on the weighted average of the underlying asset prices. It is assumed that the asset prices follow correlated geometric Brownian motion processes given by Equation (2.31).

The payoff for this option is given by,

$$h(\mathbf{S}_t) = \mathscr{K} - (w_1 S_t^1 + \cdots + w_d S_t^d), \tag{2.44}$$

subject to

$$\sum_{\delta=1}^{d} w_\delta = 1.$$

The discretization and parametrization scheme for a Bermudan put on a basket is the same as for a Bermudan call on the max of several assets. However, the case of the *basket option* is used to show how the conditional continuation value can be computed in the general case.

S_0	SGM LB(s.e)	SGM Direct (s.e)	Binomial Value	95% CI AB	95 % CI BC
n = 2 assets					
90	8.069	8.088	8.075	[8.053, 8.082]	-
	(0.026)	(0.003)			
100	13.892	13.900	13.902	[13.892, 13.934]	-
	(0.024)	(0.004)			
110	21.282	21.290	21.345	[21.316, 21.359]	-
	(0.028)	(0.003)			
n = 3 assets					
90	11.228	11.253	11.29	[11.265, 11.308]	-
	(0.023)	(0.003)			
100	18.665	18.625	18.69	[18.661, 18.728]	-
	(0.031)	(0.005)			
110	27.463	27.413	27.58	[27.512, 27.663]	-
	(0.036)	(0.006)			
n = 5 assets					
90	16.527	16.644	-	[16.602, 16.655]	[16.620, 16.653]
	(0.028)	(0.005)			
100	25.992	26.141	-	[26.109, 26.292]	[26.115, 26.164]
	(0.033)	(0.006)			
110	36.590	36.725	-	[36.704, 36.832]	[36.710, 36.798]
	(0.047)	(0.005)			

Table 2.3: Bermudan Max-Call options on 2, 3 and 5 underlying assets: The results are compared with Andersen and Broadie (2004) [2] and Broadie and Cao (2008) [15]. The parameters are: $K = 100, r = 5\%, q = 10\%, \rho = 0, T = 3, \sigma = 20\%$. There are ten exercise opportunities equally spaced in time. Values in parentheses are standard errors. The total number of grid points at each time step was 30,000 with an equal number of paths generated from the 3 source grid points (two peripheral and one initial point).

Computing the Continuation Value for Bermudan Basket Options

In order to compute the continuation value for grid points at t_m using Equation (2.9), the transition probability density function $\mathbb{P}(g(\mathbf{S}_{t_{m+1}})|\mathbf{S}_{t_m})$ is required. For a put on the weighted mean of d underlying assets, the exact transition density function is unknown. The moments for the distribution of $g(\mathbf{S}_{t_{m+1}})$ can be obtained using *sub-simulations*, which can be used to approximate the density function using the *Gram Charlier series* (Equation (2.11)). For each grid point at t_m, sub-paths are generated until the next time step t_{m+1}, and the first four non-central moments of the distribution of $g(\mathbf{S}_{t_{m+1}})$ so obtained are computed. In order to re-use the results obtained for the Bermudan max option, the distribution $\mathbb{P}(\log(|g(\mathbf{S}_{t_{m+1}})|) = x|\mathbf{S}_{t_m})$ is approximated, rather than determining $\mathbb{P}(|g(\mathbf{S}_{t_{m+1}})| = x|\mathbf{S}_{t_m})$. The continuation value is then given by Equation (2.40).

Results for Bermudan Basket Option

To illustrate the results Table 2.4 compares the results of a Bermudan put option on 4 underlying assets. In order to compute the continuation value, $N_S = 1000$, sub-paths are generated for each of the underlying assets. The computational ef-

fort increases linearly with the number of exercise opportunities M, the number of paths N_S in the sub-simulations, and the dimension of the problem d. Although computationally more expensive than the case where the moments of the distribution can be computed analytically, this example shows a generic case when it is not easy to compute the transition probability density or its moments directly. The time taken for each simulation was a few (< 5) minutes. Table 2.5 compares the results of a Bermudan put option on 5 underlying assets with those reported by Bender *et al.* (2006). The LSM values and confidence intervals reported by Bender *et al.* are close to our values.

S_0	SGM LB(s.e)	SGM Direct(s.e)	FFT Value	LSM (s.e.)
40	1.739 (0.37)	1.740 (0.16)	1.739	1.739 (0.08)

Table 2.4: Bermudan put option on arithmetic mean of 4 underlying assets: The results are compared with CONV method of Lord *et al.* (2008) [54] and the LSM values. The parameters are: $\mathscr{K} = 40, r = 6\%, q = 2\%, \rho = 0.25, T = 1, \sigma = 20\%$. There are ten exercise opportunities equally spaced in time. Values in parentheses are standard errors. The total number of grid points at each time step is 30,000 with equal numbers of paths generated from the 3 source grid points (two peripheral and one initial point). For the LSM algorithm there were 300,000 paths for each asset, and 18 basis functions.

S_0	SGM LB(s.e)	SGM Direct (s.e)	LSM (s.e)	BKS 95 % CI
90	10.000 (0.00)	10.000 (0.00)	10.000 (0.00)	[10.000, 10.004]
100	2.134 (0.012)	2.141 (0.008)	2.163 (0.001)	[2.154, 2.164]
110	0.540 (0.010)	0.550 (0.006)	0.540 (0.001)	[0.535, 0.540]

Table 2.5: Bermudan put option on arithmetic mean of 5 underlying assets: The results are compared with the intervals reported by Bender *et al.* (2006) [10] and the LSM values. The parameters are: $\mathscr{K} = 100, r = 5\%, \rho = 0, T = 3, \sigma = 20\%$. There are four exercise opportunities (including t_0) equally spaced in time . Values in parentheses are standard errors. The total number of grid points at each time step is 3000 with an equal numbers of paths generated from the 3 source grid points (two peripheral and one initial point). For the LSM algorithm there were 120,000 paths for each asset, and 24 basis functions.

2.5 Conclusion

This chapter presents the stochastic grid method for pricing and exercising Bermudan options. SGM uses dynamic programming and linear least squares regression for option pricing. One of the main achievements of the algorithm is its ability to reduce a multi-dimensional problem to a single-dimensional setting, and yet avoid some of the associated short comings as were discussed by Boyle *et al.* [12].

SGM follows the same dynamic programming approach, by approximating the option price at exercise times t_{m+1} and moving backwards in time using the information at t_{m+1}, to approximate the continuation value and hence the option price at exercise step t_m. SGM uses regression to approximate the conditional expectation $\mathbb{E}\left[\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}},\mathbf{S}_{t_0})\right]$.

The regression in SGM differs from the LSM algorithm as SGM does not approximate the functional form of the *continuation value*, rather it uses regression to approximate the functional form of $\mathbb{E}\left[V_{t_{m+1}}(\mathbf{S}_{t_{m+1}})|g(\mathbf{S}_{t_{m+1}}),\mathbf{S}_{t_0}\right]$ at the exercise dates t_{m+1}.

In SSAP [7] before pricing the option, the entire (i.e. at all time steps) state space is reduced to a one-dimensional state (using the pay-off function as the mapping function). Option pricing is then done on this reduced state space. This scheme can result in an incorrect exercise policy as was shown by Boyle *et al.* [12]. In SGM, the option price and exercise policy for grid points at t_{m+1} are first computed in the high-dimensional space. This is followed by reducing the state space at t_{m+1}. The continuation value at t_m is then computed (by iterated conditioning) using a one-dimensional probability density function, rather than multivariate distributions for the transition $\mathbf{S}_{t_{m+1}}|\mathbf{S}_{t_m}$.

Despite its advantages, SGM suffers from following main disadvantages:

- SGM can be computationally expensive when sub-simulations are required, especially when there are many early-exercise dates.
- The early-exercise policy obtained using SGM is sub-optimal when compared to that obtained using LSM, especially in higher dimensions and results in lower bound values which are farther away from the true price.
- SGM requires peripheral paths to improve the early-exercise policy, but the approach to generate the peripheral paths is not well defined.

In order to show convergence of the SGM algorithm in high dimensions with an increasing number of paths, the method should additionally include *bundling* for more accurately computation of the conditional expectation given by Equation (2.5). The next chapter introduces the *Stochastic Grid Bundling Method*, which, although motivated by SGM, overcomes much of the drawbacks of the latter by use of bundling or state space partitioning.

CHAPTER

3 The Stochastic Grid Bundling Method: Efficient Pricing of Bermudan Options and their Greeks

The contents of this chapter have appeared in [41]. The *Stochastic Grid Bundling Method* (SGBM) for pricing of Bermudan options with several underlying assets is a hybrid of *regression*- and *bundling*- based approaches. It can be seen as an improvement of SGM presented in Chapter 2. The method uses regressed value functions, together with bundling of the state space to approximate continuation values at different time steps. A high-biased *direct estimator* and an early-exercise policy are first computed using SGBM. The early-exercise policy is then used to determine a lower bound to the true option price. SGBM can also be used to compute a duality-based high-biased estimator. Compared to LSM, the approximate option values computed using SGBM, have lower numerical noise, not just at the initial step but also at intermediate time steps; which makes it a good candidate for computations that require option values at intermediate times steps (for example, computing future exposures within the credit value adjustment, CVA context).

Efficient calculation of price sensitivities continues to be among the greatest practical challenges facing users of Monte Carlo methods in the early-exercise derivatives pricing industry. Computing Greeks is essential for hedging and risk management, but it typically requires substantially more computing time than pricing the derivative. A favourable property of SGBM is that it can be used to get fast approximations of the sensitivities or Greeks of the option price, a feature illustrated through numerical examples, upto fifteen-dimensional basket option problems.

The chapter is organized as follows. Section 3.1 describes the details of SGBM. In Section 3.2 various numerical examples of increasing complexity are used to discuss various aspects of the method and finally Section 3.3 gives some concluding remarks.

3.1 Stochastic Grid Bundling Method

The Stochastic Grid Bundling Method (SGBM), which is introduced here, is a simulation-based dynamic programming method, which first generates paths forward in time,

followed by determining the optimal early-exercise policy moving backwards in time. The problem formulation used here is the same as stated in Chapter 2 (Section 2.1).The steps involved in the SGBM algorithm, which are detailed in sections to follow are:

Step I: Generating grid points

The grid points in SGBM are generated by simulating independent copies of sample paths, $\{\mathbf{S}_{t_0}(n),\ldots,\mathbf{S}_{t_M}(n)\},\ n = 1,\ldots,N$, of the underlying process $\mathbf{S}_t$, all starting from the same initial state $\mathbf{S}_{t_0}$. The n-th grid point at time step t_m is then $\mathbf{S}_{t_m}(n),\ n = 1,\ldots,N$. Depending upon the underlying process an appropriate discretization scheme, e.g. the Euler scheme, is used to generate sample paths. Sometimes the diffusion process can be simulated directly, essentially because it appears in a closed form, as an example, for the regular multi-dimensional Black-Scholes model.

Step II: Option value at terminal time

The option value at terminal time is given by:

$$V_{t_M}(\mathbf{S}_{t_M}) = \max(h(\mathbf{S}_{t_M}), 0).$$

This relation is used to compute the option value for all grid points at the final time step.

The following steps are subsequently performed for each time step, t_m, $m \leq M$, recursively, moving backwards in time, starting from t_M.

Step III: Bundling

The grid points at t_{m-1} are *bundled* into $\mathscr{B}_{t_{m-1}}(1),\ldots,\mathscr{B}_{t_{m-1}}(\nu)$ non-overlapping sets or partitions. Three different approaches for partitioning are considered in this chapter, they are:

- *k-means clustering algorithm,*
- *Recursive bifurcation,*
- *Recursive bifurcation of reduced state space.*

These techniques are detailed in a subsequent section.

Step IV: Mapping high-dimensional state space to a low-dimensional space

Corresponding to each bundle $\mathscr{B}_{t_{m-1}}(\beta),\ \beta = 1,\ldots,\nu$, a parametrized value function $Z : \mathbb{R}^d \times \mathbb{R}^K \mapsto \mathbb{R}$, which assigns values $Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})$ to states $\mathbf{S}_{t_m}$, is computed. Here $\alpha^{\beta}_{t_m} \in \mathbb{R}^K$ is a vector of free parameters[1]. The objective is then to choose, for each t_m and β, a parameter vector $\alpha^{\beta}_{t_m}$ so that

$$Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m}) \approx V_{t_m}(\mathbf{S}_{t_m}).$$

Step V: Computing the continuation and option values at t_{m-1}

The continuation values for $S_{t_{m-1}}(n) \in \mathscr{B}_{t_{m-1}}(\beta),\ n = 1,\ldots,N,\ \beta = 1,\ldots,\nu$, are approximated by

$$\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)) = \mathbb{E}[Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}}(n)]$$

The option value is then given by:

$$\widehat{V}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)) = \max(h(\mathbf{S}_{t_{m-1}}(n)), \widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n))).$$

Bundling in SGBM

SGBM employs bundling to approximate the conditional distribution using simulation. Bundling in SGBM is used to cluster grid points based on proximity.

The distribution of $\mathbf{S}_{t_m}$ conditional on the state $\mathbf{S}_{t_{m-1}} = X$, can be sampled by simulating paths from the state $\mathbf{S}_{t_{m-1}} = X$ until time step t_m. Such an approach, however, is computationally expensive, as the number of paths grows exponentially over time. Another approach for sampling this distribution is to bundle the grid points at t_{m-1}, using some measure of proximity, into ν non-overlapping partitions, and then using those paths that originate from the bundle that contains $\mathbf{S}_{t_{m-1}} = X$ to sample $\mathbf{S}_{t_m}$. As will be shown in the discussions to follow, under our model assumptions, with increasing numbers of paths and bundles, this sampled distribution approaches the true conditional distribution.

From here on, the bundle that contains $\mathbf{S}_{t_{m-1}} = X$, where X be any point in $\mathbb{R}^d$ or a simulated grid point $\mathbf{S}_{t_{m-1}}(n)$, will be indicated by $\mathscr{B}_{t_{m-1}}(\beta)$, where β can be $1,\ldots,\nu$.

The different bundling techniques employed are first explained:

[1] The notation here is slightly different from that used in Chapter 2

k-means clustering

Given N grid points, $(\mathbf{S}_{t_{m-1}}(1),\ldots,\mathbf{S}_{t_{m-1}}(N))$, *k-means* clustering aims to partition these N observations into ν non-overlapping sets, $\mathscr{B}_{t_{m-1}} = \{\mathscr{B}_{t_{m-1}}(1),\ldots,\mathscr{B}_{t_{m-1}}(\nu)\}$ so as to minimize the sum of squares within clusters, i.e:

$$\underset{\mathscr{B}_{t_{m-1}}}{\operatorname{argmin}} \sum_{\beta=1}^{\nu} \left(\sum_{\mathbf{S}_{t_{m-1}}(n)\in\mathscr{B}_{t_{m-1}}(\beta)} \| \mathbf{S}_{t_{m-1}}(n) - \mu_\beta \|^2 \right),$$

where μ_β is the mean of the points in $\mathscr{B}_{t_{m-1}}(\beta)$.

Lloyd's algorithm (Lloyd 1982)[51] can be used to bundle the grid points using *k-means clustering*. Briefly, the algorithm uses an iterative refinement technique, where, given the initial guess of cluster means, $\mu_1^{(1)},\ldots,\mu_\nu^{(1)}$, the algorithm performs the following two steps alternately:

Step 1 Assign grid points to the set whose mean is closest to it.

$$\mathscr{B}^{(l)}_{t_{m-1}}(\beta) = \{\mathbf{S}_{t_{m-1}}(n) : \| \mathbf{S}_{t_{m-1}}(n) - \mu^{(l)}_\beta \|^2 \leq \| \mathbf{S}_{t_{m-1}}(n) - \mu^{(l)}_j \|^2, \forall\, 1 \leq j \leq \nu\},$$

where grid point $\mathbf{S}_{t_{m-1}}(n)$ is assigned to only one bundle, even though it could be assigned to more than one.

Step 2 The results have converged if the assignment of the grid points doesn't change anymore from a previous iteration, else the centroids are updated for the new clusters as:

$$\mu^{(l+1)}_\beta = \frac{1}{|\mathscr{B}^{(l)}_{t_{m-1}}(\beta)|} \sum_{\mathbf{S}_{t_{m-1}}(n)\in\mathscr{B}^{(l)}_{t_{m-1}}(\beta)} \mathbf{S}_{t_{m-1}}(n).$$

Figure 3.1 shows the bundles obtained using the *k means clustering* algorithm for a two asset case at a particular time step.

Recursive bifurcation

The aim of bundling in SGBM is to cluster grid points based on proximity and as will become evident through numerical examples, it's not so important how optimally the grid points are allocated to the different bundles, therefore, a fast practical scheme to do this is proposed here. To bundle the grid points, $\{\mathbf{S}_{t_{m-1}}(1),\ldots,\mathbf{S}_{t_{m-1}}(N)\}$, the following steps can be performed:

Step 1: Compute the *mean* of the given set of grid points, along each dimension, i.e.

$$\mu_\delta = \frac{1}{N}\sum_{n=1}^{N} S^\delta_{t_{m-1}}(n), \delta = 1,\ldots,d.$$

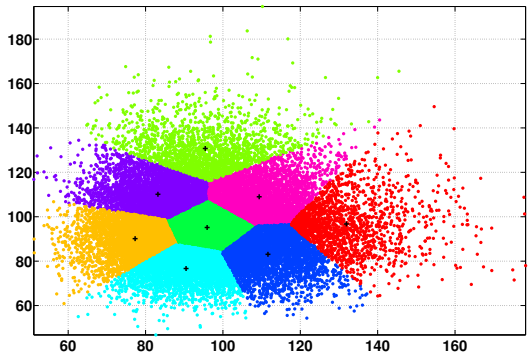

Figure 3.1: Bundling of grid points in a two-dimensional space using *k-means* clustering. The grid points are bundled into 8 non-overlapping partitions.

Step 2: The grid points are bundled separately along each dimension. This is done by dividing the grid points into 2^d sets as:

$$A_\delta = \{\mathbf{S}_{t_{m-1}}(n) : S^\delta_{t_{m-1}}(n) > \mu_\delta,\, n = 1,\ldots,N\},$$
$$\overline{A}_\delta = \{\mathbf{S}_{t_{m-1}}(n) : S^\delta_{t_{m-1}}(n) \leq \mu_\delta,\, n = 1,\ldots,N\},$$

where $\delta = 1,\ldots,d$.

Step 3: The 2^d unique *non-overlapping* bundles are then obtained using the following intersections of these sets :

$$\begin{aligned}
\mathscr{B}_{t_{m-1}}(1) &= A_1 \cap A_2 \cap \ldots \cap A_d, \\
\mathscr{B}_{t_{m-1}}(2) &= \overline{A}_1 \cap A_2 \cap \ldots \cap A_d, \\
\mathscr{B}_{t_{m-1}}(3) &= A_1 \cap \overline{A}_2 \cap \ldots \cap A_d, \\
&\vdots \\
\mathscr{B}_{t_{m-1}}(2^d) &= \overline{A}_1 \cap \overline{A}_2 \cap \ldots \cap \overline{A}_d,
\end{aligned} \tag{3.1}$$

Step 4: Bundles $\mathscr{B}_{t_{m-1}}(1),\ldots$ can be split further, in the next iteration, again following the steps above.

The number of partitions, or bundles, after p iterations, where each of the bundles obtained is split further, would be equal to $(2^d)^p$. Figure 3.2 shows an example of bundling of grid points in a two-dimensional space. First, the single large bundle is halved along each dimension, resulting in a total of four partitions. Then, each of these partitions undergoes the same process, resulting in 16 partitions in the second iteration, and 64 bundles in the third iteration. The number of computations is linear in the total number of grid points, N, the number of dimensions, d, and the

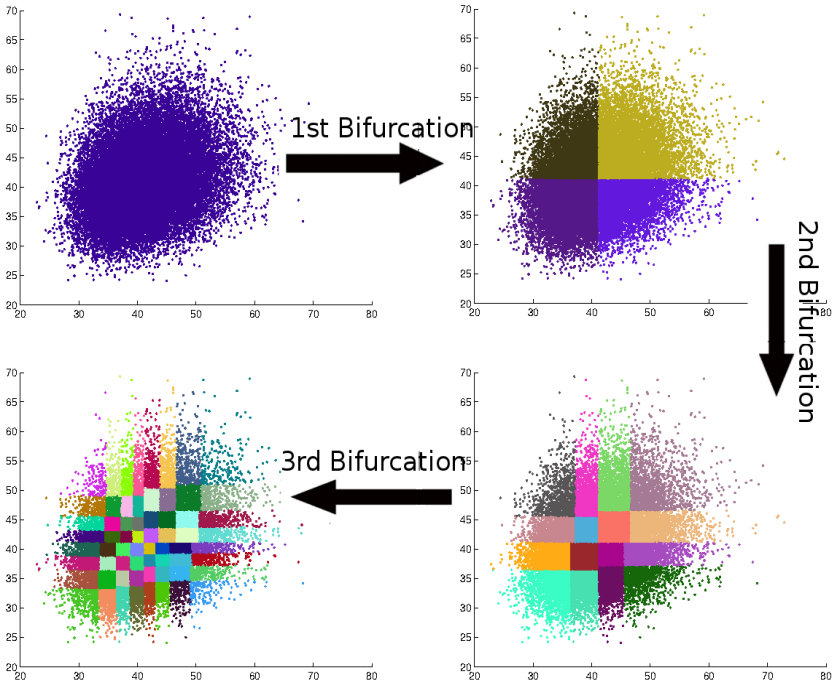

Figure 3.2: Bundling of grid points in a two-dimensional space. After the first bifurcation 4 bundles are obtained. A second bifurcation for each of these 4 bundles results in 16 bundles and after the third bifurcation 64 bundles are obtained.

number of iteration steps, p, which makes this method of bundling practical and fast. However, this approach will be less attractive with increasing dimensions of the problem, as the number of bundles obtained after each iteration would be too large.

Recursive bifurcation of reduced state space

This method is motivated by the *stratified state aggregation method* by Barraquand and Martineau (1995)[7], where rather than partitioning the actual state space in which the value function resides, a *reduced state space* obtained by some mapping function is partitioned. Like the authors of [7], the payoff is used as the mapping function and the grid points are bundled based on proximity of the reduced state space $h(\mathbf{S}_{t_{m-1}})$. The bundling scheme is then similar to the *recursive bifurcation*, except that now the effective dimension d will be equal to 1. The number of bundles obtained after p iterations in this case will be 2^p.

Assumption 1. *$\mathbf{S}_{t_m}$, $m = 1,\dots,M$, is an everywhere dense set of $\mathbb{R}^d$ valued vectors. Furthermore, the probability density function of $\mathbf{S}_{t_m}$ is assumed to be continuous everywhere.*

Assumption 2.

$$\lim_{\nu\to\infty}\lim_{N\to\infty}|\mathscr{B}_{t_{m-1}}(\beta)|\to\infty,\ m=2,\dots,M,\ \beta=1,\dots,\nu.$$

Definition 1.

$$\mathbb{P}_{N,\nu}(\mathbf{S}_{t_m}\le y|\mathbf{S}_{t_{m-1}}=X):=\frac{\frac{1}{N}\sum_{n=1}^{N}\mathbf{1}_{\mathbf{S}_{t_m}(n)\le y}(\mathbf{S}_{t_m}(n)).\mathbf{1}_{\mathscr{B}_{t_{m-1}}(\beta)}(\mathbf{S}_{t_{m-1}}(n))}{\frac{1}{N}|\mathscr{B}_{t_{m-1}}(\beta)|},$$

where, $(\mathbf{S}_{t_{m-1}}=X)\in\mathscr{B}_{t_{m-1}}(\beta)$.

When Assumptions 1 and 2 hold true, then

Proposition 1.

$$\lim_{\nu\to\infty}\lim_{N\to\infty}|\mathbb{P}_{N,\nu}(\mathbf{S}_{t_m}\le y|\mathbf{S}_{t_{m-1}}=X)-\mathbb{P}(\mathbf{S}_{t_m}\le y|\mathbf{S}_{t_{m-1}}=X)|=0.$$

Proof for Proposition 1

The following definitions will assist us in the proof of Proposition 1.

Definition 2. *Let $x:=(x_1,\dots,x_\nu)\in(\mathbb{R}^d)^\nu$. A Borel partition $\mathscr{C}_\beta(x)$, $\beta=1,\dots,\nu$ of $\mathbb{R}^d$ is a Voronoi tessellation of x if, for every $\beta\in\{1,\dots,\nu\}$, $\mathscr{C}_\beta(x)$ satisfies*

$$\mathscr{C}_\beta(x)\subset\{y\in\mathbb{R}^d|\ \|\ x_\beta-y\ \|=\min_{1\le j\le\nu}\|\ y-x_j\ \|\}.$$

Definition 3. *Let $X \in L^2_{\mathbb{R}^d}(\Omega, \mathscr{F}, \mathbb{P})$. The random vector*

$$\widehat{X}^x = \sum_{\beta=1}^{\nu} x_\beta \mathbf{1}_{\mathscr{C}_\beta(x)}(X),$$

is called a Voronoi quantization of X.

When $\mu := (\mu_1, \ldots, \mu_\nu)$ are the centroids obtained from the *k-means* clustering algorithm then $\mathscr{B}_{t_{m-1}}(\beta) \subset \mathscr{C}_\beta(\mu)$ and

$$\widehat{\mathbf{S}}^\mu_{t_{m-1}} = \sum_{\beta=1}^{\nu} \mu_\beta \mathbf{1}_{\mathscr{B}_{t_{m-1}}(\beta)}(\mathbf{S}_{t_{m-1}}),$$

where $\mathbf{1}_{\mathscr{B}_{t_{m-1}}(\beta)}(\mathbf{S}_{t_{m-1}})$ is the indicator function which returns 1 if $\mathbf{S}_{t_{m-1}}$ belongs to the bundle $\mathscr{B}_{t_{m-1}}(\beta)$, and 0 otherwise.

Lemma 1. *When Assumptions 1 and 2 hold, then*

$$\lim_{\nu\to\infty} \lim_{N\to\infty} \| \mathbf{S}_{t_{m-1}} - \widehat{\mathbf{S}}^\mu_{t_{m-1}} \| = 0,$$

Proof: We assume that in the limiting case, when the number of grid points, N, and bundles, ν, go to infinity, $\widehat{\mathbf{S}}^\mu_{t_{m-1}}$ to be an everywhere dense set of $\mathbb{R}^d$ valued vectors. An intuitive explanation for the assumption is that for a finite N when the number of bundles ν are equal to N the bundle centroids would coincide with the grid points, and will have the same distribution as the grid points. As, when N goes to infinity, $\mathbf{S}_{t_{m-1}}$ is an everywhere dense set of $\mathbb{R}^d$ valued vectors, then it is safe to assume that when the number of grid points and bundles go to infinity, the bundle centroids also constitute an everywhere dense set. It follows then by Lebesgue dominated convergence theorem, $\| \mathbf{S}_{t_{m-1}} - \widehat{\mathbf{S}}^\mu_{t_{m-1}} \|$ goes to zero.

For the case of *recursive bifurcation*, we sketch the proof of Lemma 1 as following:

Assume that $\mathbf{S}_{t_{m-1}} = X$ belongs to bundle $\mathscr{B}_{t_{m-1}}(\beta)$, which is bounded as follows

$$\mathscr{B}_{t_{m-1}}(\beta) = \{\mathbf{S}_{t_{m-1}} | \mathbf{s}_{\min} < \mathbf{S}_{t_{m-1}} \leq \mathbf{s}_{\max}\}.$$

Let $\epsilon = \max(|X - \mathbf{s}_{\min}|, |X - \mathbf{s}_{\max}|)$ be the maximum width of a given bundle. When N goes to infinity, it is easy to see that the grid points can be recursively partitioned until maximum width, ϵ, of the bundle is less than an arbitrarily small ε. Lemma 1 then follows from dominated convergence.

Proof for Proposition 1

Event A is defined as $A := \{\mathbf{S}_{t_m} | \mathbf{S}_{t_m} \leq y\}$, and event B, is defined as $(\mathbf{S}_{t_{m-1}} = X) \in \mathscr{B}_{t_{m-1}}(\beta)$, i.e. $\widehat{X}^{\mu} = \mu_{\beta}$.

As the distribution of $\mathbf{S}_{t_{m-1}}$ is continuous everywhere, by the law of large numbers we find,

$$\mathbb{P}(B) = \lim_{N\to\infty} \frac{|\mathscr{B}_{t_{m-1}}(\beta)|}{N}.$$

Using Lemma 1 we have,

$$\mathbb{P}(\mathbf{S}_{t_{m-1}} = X) = \lim_{\nu\to\infty} \mathbb{P}(B) = \lim_{\nu\to\infty} \lim_{N\to\infty} \frac{|\mathscr{B}_{t_{m-1}}(\beta)|}{N}. \tag{3.2}$$

Similarly, it can be shown that:

$$\begin{aligned}
\mathbb{P}((\mathbf{S}_{t_m} \leq y) \cap (\mathbf{S}_{t_{m-1}} = X)) &= \lim_{\nu\to\infty} \mathbb{P}(A \cap B) \\
&= \lim_{\nu\to\infty} \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \mathbf{1}_{\mathbf{S}_{t_m} \leq y}(\mathbf{S}_{t_m}(n)).\mathbf{1}_{\mathscr{B}_{t_{m-1}}(\beta)}(\mathbf{S}_{t_{m-1}}(n))
\end{aligned}$$

Again, assuming that the conditional distribution is continuous everywhere, we obtain

$$\begin{aligned}
\mathbb{P}(\mathbf{S}_{t_m} \leq y | \mathbf{S}_{t_{m-1}} = X) &= \lim_{\nu\to\infty} \frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)} \\
&= \lim_{\nu\to\infty} \lim_{N\to\infty} \frac{\frac{1}{N}\sum_{n=1}^{N} \mathbf{1}_{\mathbf{S}_{t_m} \leq y}(\mathbf{S}_{t_m}(n)).\mathbf{1}_{\mathscr{B}_{t_{m-1}}(\beta)}(\mathbf{S}_{t_{m-1}}(n))}{\frac{1}{N}|\mathscr{B}_{t_{m-1}}(\beta)|}.
\end{aligned}$$

Parameterizing the option values

As the dimension of the state space is usually large, the pricing problem becomes intractable and requires the approximation of the value function. This can be achieved by introducing a *parametrized value function* $Z : \mathbb{R}^d \times \mathbb{R}^K \mapsto \mathbb{R}$, which assigns a value $Z(\mathbf{S}_{t_m}, \alpha)$ to state $\mathbf{S}_{t_m}$, where $\alpha \in \mathbb{R}^K$ is a vector of free parameters. The objective is to choose, corresponding to each bundle β at epoch t_{m-1}, a parameter vector $\alpha_{t_m}^{\beta} := \alpha$ so that,

$$V_{t_m}(\mathbf{S}_{t_m}) \approx Z(\mathbf{S}_{t_m}, \alpha_{t_m}^{\beta}),$$

This chapter follows Tsitsiklis and van Roy (2001)[83] in defining the approximation function. A basis function[2] that maps the state space from $\mathbb{R}^d$ to $\mathbb{R}$, is used to

[2] Basis functions are sometimes referred to as *features*

approximate the value functions. For a particular problem it may be required to define several basis functions, $\phi_1,\ldots,\phi_K$. Then, to each state $\mathbf{S}_{t_m}$ a vector, $\phi(\mathbf{S}_{t_m}) = (\phi_1(\mathbf{S}_{t_m}),\ldots,\phi_K(\mathbf{S}_{t_m}))'$, of basis functions is associated. The vector of basis functions is chosen such that it represents the most salient properties of a given state.

In our approximation $Z(\mathbf{S}_{t_m},\alpha^{\beta}_{t_m})$ depends on $\mathbf{S}_{t_m}$ only through $\phi(\mathbf{S}_{t_m})$. Hence, for some function $f:\mathbb{R}^K\times\mathbb{R}^K\mapsto\mathbb{R}$, we can write $Z(\mathbf{S}_{t_m},\alpha^{\beta}_{t_m}) = f(\phi(\mathbf{S}_{t_m}),\alpha^{\beta}_{t_m})$. Usually, the basis functions ϕ, are selected based on the problem and relies on human experience, as in the case of the *least squares method* (LSM) (Longstaff and Schwartz (2001)[53]). The function f, which maps the option values onto the span of ϕ, in our discussion is restricted to the form:

$$Z(\mathbf{S}_{t_m},\alpha^{\beta}_{t_m}) = \sum_{k=1}^{K}\alpha^{\beta}_{t_m}(k)\phi_k(\mathbf{S}_{t_m}), \tag{3.3}$$

i.e., the value function is a linear combination of basis functions. Define a weighted quadratic norm $\| V_{t_m}(\mathbf{S}_{t_m})\|_{\pi}$, as

$$\| V_{t_m}(\mathbf{S}_{t_m})\|_{\pi} = \left(\int_{\mathbf{S}_{t_m}\in\mathbb{R}^d} V^2_{t_m}(\mathbf{S}_{t_m})d(\pi(\mathbf{S}_{t_m}))\right)^{\frac{1}{2}},$$

where $\pi(\mathbf{S}_{t_m})$ is the conditional distribution of $\mathbf{S}_{t_m}$, with a conditional density function given by $\mathbb{P}(\mathbf{S}_{t_m}|\mathscr{B}_{t_{m-1}}(\beta))$. Here, it should be emphasized that $\pi(\mathbf{S}_{t_m})$ is the distribution of $\mathbf{S}_{t_m}$ conditional on information $\mathbf{S}_{t_{m-1}}\in\mathscr{B}_{t_{m-1}}(\beta)$ and *not* conditional on $\mathbf{S}_{t_0}$, which is a major difference compared to SGM in Chapter 2. It is assumed that $\| V_{t_m}(\mathbf{S}_{t_m})\|_{\pi}<\infty$. The problem is then to project the value function onto the span of $\phi_1,\ldots,\phi_K$, which is characterized by:

$$\underset{\alpha^{\beta}_{t_m}}{\operatorname{argmin}}\| V_{t_m}(\mathbf{S}_{t_m}) - \sum_{k=1}^{K}\alpha^{\beta}_{t_m}(k)\phi_k(\mathbf{S}_{t_m})\|_{\pi}.$$

Exact computation of the above projection is not generally viable, however, it can be *approximated* by sampling a collection of states $\mathbf{S}_{t_m}$, according to probability measure $\mathbb{P}(\mathbf{S}_{t_m}|\mathscr{B}_{t_{m-1}}(\beta))$. The approximate distribution $\tilde{\pi}(\mathbf{S}_{t_m})$, of grid points $\mathbf{S}_{t_m}(n)$, whose paths originate from the bundle $\mathscr{B}_{t_{m-1}}(\beta)$, would in the limiting case of the number of paths N have a transition density function equal to $\mathbb{P}(\mathbf{S}_{t_m}|\mathscr{B}_{t_{m-1}}(\beta))$, i.e. $\lim_{N\to\infty}\tilde{\pi}(\mathbf{S}_{t_m})\equiv\pi(\mathbf{S}_{t_m})$. The projection (3.3) can now be approximated by:

$$Z(\mathbf{S}_{t_m},\widehat{\alpha}^{\beta}_{t_m}) = \sum_{k=1}^{K}\widehat{\alpha}^{\beta}_{t_m}(k)\phi_k(\mathbf{S}_{t_m}), \tag{3.4}$$

and satisfies,

$$\underset{\widehat{\alpha}^{\beta}_{t_m}}{\operatorname{argmin}}\sum_{n=1}^{|\mathscr{B}_{t_{m-1}}(\beta)|}\left(V_{t_m}(\mathbf{S}_{t_m}(n)) - \sum_{k=1}^{K}\widehat{\alpha}^{\beta}_{t_m}(k)\phi_k(\mathbf{S}_{t_m}(n))\right)^2. \tag{3.5}$$

Proposition 2. *Assume that Assumptions 1 and 2 hold, then it follows that,*

$$\lim_{N\to\infty} \| Z(\mathbf{S}_{t_m}, \widehat{\alpha}^{\beta}_{t_m}) - Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m}) \|_{\pi} = 0.$$

Proof: Convergence is by the law of large numbers.

To summarize, corresponding to each bundle $\mathscr{B}_{t_{m-1}}(\beta)$, a parametrized function $Z(\mathbf{S}_{t_m}, \widehat{\alpha}^{\beta}_{t_m})$ is computed using ordinary least squares regression, so that:

$$V_{t_m}(\mathbf{S}_{t_m}(n)) = Z(\mathbf{S}_{t_m}(n), \widehat{\alpha}^{\beta}_{t_m}) + \epsilon^{\beta}_{t_m}, \tag{3.6}$$

where $\mathbf{S}_{t_{m-1}}(n) \in \mathscr{B}_{t_{m-1}}(\beta)$.

Ordinary least squares gives us an unbiased estimator, and it is further assumed that the following assumption holds:

Assumption 3. $\mathbb{E}[\epsilon^{\beta}_{t_m} | \mathbf{S}_{t_{m-1}}(n)] = 0, \mathbf{S}_{t_{m-1}}(n) \in \mathscr{B}_{t_{m-1}}(\beta)$.

Function $Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})$ can also be seen as the linear unbiased estimator of the conditional expectation

$$Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m}) = \mathbb{E}[V_{t_m}(\mathbf{S}_{t_m}) | \phi(\mathbf{S}_{t_m}), \mathscr{B}_{t_{m-1}}(\beta)].$$

Proposition 3.

$$\lim_{\nu\to\infty} Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m}) = \mathbb{E}[V_{t_m}(\mathbf{S}_{t_m}) | \phi(\mathbf{S}_{t_m}), \mathbf{S}_{t_{m-1}} = X].$$

Proof: As the number of sample paths N and bundles ν go to infinity, Lemma 1 shows that the maximum distance between any two grid points within a bundle approaches zero. Therefore, for the bundle $\mathscr{B}_{t_{m-1}}(\beta)$, that contains $\mathbf{S}_{t_{m-1}} = X$, it can then be stated that,

$$\lim_{\nu\to\infty} \mathbb{E}[V_{t_m}(\mathbf{S}_{t_m}) | \phi(\mathbf{S}_{t_m}), \mathscr{B}_{t_{m-1}}(\beta)] \quad = \quad \mathbb{E}[V_{t_m}(\mathbf{S}_{t_m}) | \phi(\mathbf{S}_{t_m}), \mathbf{S}_{t_{m-1}} = X].$$

Computing the continuation value

Using the parametrized option value function $Z(\mathbf{S}_{t_m}, \widehat{\alpha}^{\beta}_{t_m})$ corresponding to bundle $\mathscr{B}_{t_{m-1}}(\beta)$, the continuation values for the grid points that belong to this bundle are approximated by:

$$\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)) = D_{t_{m-1}}\mathbb{E}[Z(\mathbf{S}_{t_m}, \widehat{\alpha}^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}} = \mathbf{S}_{t_{m-1}}(n)], \tag{3.7}$$

where $\mathbf{S}_{t_{m-1}}(n) \in \mathscr{B}_{t_{m-1}}(\beta)$. Using Equation (3.4), this can be written as:

$$\begin{aligned}\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)) &= D_{t_{m-1}}\mathbb{E}\left[\left(\sum_{k=1}^{K}\widehat{\alpha}^{\beta}_{t_m}(k)\phi_k(\mathbf{S}_{t_m})\right)|\mathbf{S}_{t_{m-1}} = \mathbf{S}_{t_{m-1}}(n)\right] \\ &= D_{t_{m-1}}\sum_{k=1}^{K}\widehat{\alpha}^{\beta}_{t_m}(k)\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = \mathbf{S}_{t_{m-1}}(n)\right]. \end{aligned} \tag{3.8}$$

The vector of basis functions ϕ should ideally be chosen such that $\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X\right]$, is known in a closed form, or has an analytic approximation. It is observed that $h(\cdot)$, is usually an important basis function, and, as a rule of thumb, if analytic solutions or approximations for a single time period European equivalent of an option are available, then the Bermudan option pricing problem at hand can be solved efficiently using SGBM. Examples are given in Section 3.2.

Theorem 1. *When Assumptions 1 and 2 hold, then,*

$$\lim_{\nu\to\infty}\lim_{N\to\infty}|\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}) - Q_{t_{m-1}}(\mathbf{S}_{t_{m-1}})| = 0.$$

Proof:

The continuation value, using the law of iterated conditioning, can be written as,

$$\begin{aligned}Q_{t_{m-1}}(\mathbf{S}_{t_{m-1}} = X) &= \mathbb{E}[V_{t_m}(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X] \\ &= \mathbb{E}[\mathbb{E}[V_{t_m}(\mathbf{S}_{t_m})|\phi(\mathbf{S}_{t_m}), \mathbf{S}_{t_{m-1}} = X]|\mathbf{S}_{t_{m-1}} = X].\end{aligned}$$

Using Propositions 2 and 3 it can then be stated that:

$$\begin{aligned}Q_{t_{m-1}}(\mathbf{S}_{t_{m-1}} = X) &= \lim_{\nu\to\infty}\mathbb{E}[Z(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}} = X],\ \beta = 1,\ldots,\nu \\ &= \lim_{\nu\to\infty}\lim_{N\to\infty}\mathbb{E}[Z(\mathbf{S}_{t_m}, \widehat{\alpha}^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}} = X] \\ &= \lim_{\nu\to\infty}\lim_{N\to\infty}\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}} = X).\end{aligned}$$

Definition 4. *The* direct estimator *of the option values for the grid points at t_{m-1}, is defined as*

$$\widehat{V}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)) = \max(h(\mathbf{S}_{t_{m-1}}(n), \widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}}(n)),$$

where $n = 1, \ldots, N$.

Theorem 2. *Under the condition that $\mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X]$ is known exactly, and Assumption 3 holds, the* direct estimator *is biased high, i.e.*

$$\mathbb{E}[\widehat{V}_{t_0}(\mathbf{S}_{t_0})] \geq V_{t_0}(\mathbf{S}_{t_0}).$$

Proof: The proof is by induction. At the terminal time $\widehat{V}_{t_M}(\mathbf{S}_{t_M}) = V_{t_M}(\mathbf{S}_{t_M}) = h(\mathbf{S}_{t_M})$, for all $\mathbf{S}_{t_M}$. Take as induction hypothesis $\mathbb{E}[\widehat{V}_{t_m}(\mathbf{S}_{t_m})] \geq V_{t_m}(\mathbf{S}_{t_m})$ for all $\mathbf{S}_{t_m}$. Now,

$$\begin{aligned}
\mathbb{E}[\widehat{V}_{t_{m-1}}(\mathbf{S}_{t_{m-1}} = X)] &= \mathbb{E}\left[\max\left(h(\mathbf{S}_{t_{m-1}} = X), \mathbb{E}\left[\widehat{Z}(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}} = X\right]\right)\right] \\
&\geq \max\left(h(X), \mathbb{E}\left[\mathbb{E}\left[\widehat{Z}(\mathbf{S}_{t_m}, \alpha^{\beta}_{t_m})|\mathbf{S}_{t_{m-1}} = X\right]\right]\right) \\
&= \max\left(h(X), \mathbb{E}\left[\mathbb{E}\left[\left(\widehat{V}_{t_m}(\mathbf{S}_{t_m}) - \epsilon^{\beta}_{t_m}\right)|\mathbf{S}_{t_{m-1}} = X\right]\right]\right) \\
&= \max\left(h(X), \mathbb{E}\left[\mathbb{E}\left[\mathbb{E}[\widehat{V}_{t_m}(\mathbf{S}_{t_m})|\mathbf{S}_{t_m}]|\mathbf{S}_{t_{m-1}} = X\right]\right]\right) \\
&\geq \max\left(h(X), \mathbb{E}\left[\mathbb{E}\left[V_{t_m}(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X\right]\right]\right) \\
&= \max\left(h(X), \mathbb{E}\left[V_{t_m}(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X\right]\right) \\
&= V_{t_{m-1}}(\mathbf{S}_{t_{m-1}} = X).
\end{aligned}$$

The first step uses Equation (3.7), the second Jensen's inequality and the third uses Equation (3.6). The fourth step is based on Assumption 3 and uses the basic property of a conditional expectation. The fifth step employs the induction hypothesis and the sixth is again based on a basic property of the expectation.

Corollary 1. *When Assumptions 1 to 3 hold, we have,*

$$\lim_{\nu \to \infty} \lim_{N \to \infty} |\widehat{V}_{t_0}(\mathbf{S}_{t_0}) - V_{t_0}(\mathbf{S}_{t_0})| = 0$$

Proof: The proof is an immediate outcome of Theorem 1 and the dynamic programming formulation given by Equations (1.3) and (1.4).

Lower bounds using path estimator

Once the early-exercise policy and direct estimator values have been obtained, an estimator based on the simulated paths which is biased low can be developed. Together, the high-biased direct estimator and the low-biased estimator, can generate a valid confidence interval for the option price. In order to compute the low-biased estimates, we generate a new set of paths $\mathbf{S}(n) = \{\mathbf{S}_{t_1}(n), \ldots, \mathbf{S}_{t_M}(n)\}$, $n = 1, \ldots, N_L$, using the same scheme as followed for generating the paths for direct estimator. Along each path, the approximate optimal policy exercises at,

$$\widehat{\tau}^*(\mathbf{S}(n)) = \min\{t_m : h(\mathbf{S}_{t_m}(n)) \geq \widehat{Q}_{t_m}(\mathbf{S}_{t_m}(n)),\, m = 1, \ldots, M\},$$

where $\widehat{Q}_{t_m}(\mathbf{S}_{t_m}(n))$ is computed using Equation (3.8). The path estimator is defined by $v(n) = h(\mathbf{S}_{\widehat{\tau}^*(\mathbf{S}(n))})$.

Theorem 3. *A low-biased estimate, $\underline{V}_{t_0}(\mathbf{S}_{t_0})$, to the true option value, $V_{t_0}(\mathbf{S}_{t_0})$, can be computed as:*

$$\underline{V}_{t_0}(\mathbf{S}_{t_0}) = \lim_{N_L} \frac{1}{N_L} \sum_{n=1}^{N_L} v(n) \leq V_{t_0}(\mathbf{S}_{t_0}).$$

Under the assumption that Proposition 1 holds, additionally it can be show that

$$\underline{V}_{t_0}(\mathbf{S}_{t_0}) \to V_{t_0}(\mathbf{S}_{t_0}). \tag{3.9}$$

The proof for the bias of the path estimator, i.e. Theorem 3, and the convergence of the path estimator is the same as the proofs for the Theorems 3 and 4 in Broadie and Glasserman (2001)[17].

Variance Reduction

The *direct estimator* to the option price usually has lower variance than the *path estimator*, because the parametrized option value function at t_m uses basis functions, $\phi_k(\mathbf{S}_{t_m})$, whose expectations, $\mathbb{E}[\phi(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}}]$, are either known, or an accurate numerical estimate of them can be obtained. Therefore, $\phi_k(\mathbf{S}_{t_m})$ *additionally serves as a control variate*. To elaborate further, our interest is in computing $\widehat{Q}_{t_{m-1}}(X) = \mathbb{E}[\widehat{V}_{t_m}(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X]$, when the expectation $\mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X]$, is known. For simplicity, let's assume that the sample $\widehat{V}_{t_m}(\mathbf{S}_{t_m}(n)), \phi(\mathbf{S}_{t_m}(n))$, $n = 1, \ldots, N$, is generated from $\mathbf{S}_{t_{m-1}} = X$. The usual procedure to form the *controlled estimator* $\tilde{Q}_{t_{m-1}}(X)$ is (see Rasmussen 2005 [68] for details),

$$\frac{1}{N}\sum_{m=1}^{N} \widehat{V}(\mathbf{S}_{t_m}(n)) - \sum_{k=1}^{K} \alpha_{t_m}(k) \left(\frac{1}{N} \sum_{n=1}^{N} \phi_k(\mathbf{S}_{t_m}(n)) - \mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X] \right), \tag{3.10}$$

where $\alpha_{t_m}(k)$ is chosen to solve:

$$\min_{\alpha_{t_m}} \frac{1}{N} \sum_{n=1}^{N} \left[\widehat{V}_{t_m}(\mathbf{S}_{t_m}(n)) - \left(\sum_{k=1}^{K} \alpha_{t_m}(k) \phi_k(\mathbf{S}_{t_m}(n)) \right) \right]^2. \tag{3.11}$$

Note that Equation (3.11) is the same as (3.5). Reordering Equation (3.10), gives us

$$\begin{aligned} \tilde{Q}_{t_{m-1}}(X) &= \frac{1}{N} \sum_{n=1}^{N} \left(\widehat{V}(\mathbf{S}_{t_m}(n)) - \sum_{k=1}^{K} \alpha_{t_m}(k) \phi_k(\mathbf{S}_{t_m}(n)) \right) \\ &+ \sum_{k=1}^{K} \alpha_{t_m}(k) \mathbb{E}[\phi_k(\mathbf{S}_{t_m}) | \mathbf{S}_{t_{m-1}} = X] \\ &= \frac{1}{N} \sum_{n=1}^{N} \epsilon_{t_m}(n) + \sum_{k=1}^{K} \alpha_{t_m}(k) \mathbb{E}[\phi_k(\mathbf{S}_{t_m}) | \mathbf{S}_{t_{m-1}} = X] \\ &= \widehat{Q}_{t_{m-1}}(X), \end{aligned}$$

where the mean of $\epsilon_{t_m}(n)$ is zero as ordinary least squares gives us an unbiased estimator. Therefore, the direct estimator can also be seen as controlled estimator. The effectiveness of the procedure depends on the correlation of $\widehat{V}_{t_m}(\mathbf{S}_{t_m}(n))$ and $\phi(\mathbf{S}_{t_m}(n)))$.

Computing the Greeks

An advantage of SGBM is that it can be *used directly* to approximate the first- and second-order derivatives of the option price with respect to the underlying assets. Existing methods for computing Greeks have been discussed in Glasserman(2004)[34], and can broadly be classified into methods that employ finite-difference approximations, and methods that use information about the simulated stochastic process to replace numerical differentiation by exact differentiation. The path-wise derivative method and the likelihood ratio method belong to the second category, and are found to be computationally more efficient than the finite-difference approach. Wang and Caflisch (2009)[87] propose a modified least squares method for estimating the Greeks, and they report a performance comparable to the path-wise method, the likelihood ratio method as well as the likelihood ratio and duality (LRD) method [47]. However, the choice of initial distribution is arbitrary, and may have a significant effect in extreme cases. Other recent methods for computing the sensitivity of Bermudan options include, Belomestny *et al.* (2010)[8], who use a regression-based approach for computing the Greeks; Capriotti and Giles (2010)[18], who use the Adjoint method for computing the option price sensitivities.

The option delta is defined as,

$$\Delta^{\delta}_{t_0} := \frac{\partial V_{t_0}(\mathbf{S}_{t_0})}{\partial S^{\delta}_{t_0}}, \delta = 1, \ldots, d,$$

and can be computed as,

$$\frac{\partial V_{t_0}(\mathbf{S}_{t_0})}{\partial S^{\delta}_{t_0}} = \lim_{x\to 0} \frac{V_{t_0}(\overline{\mathbf{S}}_{t_0}) - V_{t_0}(\mathbf{S}_{t_0})}{(S^{\delta}_{t_0} + x) - S^{\delta}_{t_0}}, \tag{3.12}$$

where $\overline{\mathbf{S}}_{t_0} = \{S^1_{t_0}, \ldots, S^{\delta}_{t_0} + x, \ldots, S^d_{t_0}\}$.

SGBM approximates Equation (3.12) by:

$$\begin{aligned}
\widehat{\Delta}^{\delta}_{t_0} &= \lim_{x\to 0} \frac{\widehat{V}_{t_0}(\overline{\mathbf{S}}_{t_0}) - \widehat{V}_{t_0}(\mathbf{S}_{t_0})}{(S^{\delta}_{t_0} + x) - S^{\delta}_{t_0}} \\
&= \lim_{x\to 0} \frac{\mathbb{E}[Z(\mathbf{S}_{t_1}, \widehat{\alpha}^1_{t_1})|\overline{\mathbf{S}}_{t_0}] - \mathbb{E}[Z(\mathbf{S}_{t_1}, \widehat{\alpha}^1_{t_1})|\mathbf{S}_{t_0}]}{x} \\
&= \lim_{x\to 0} \frac{\mathbb{E}\left[\sum_{k=1}^{K} \widehat{\alpha}^1_{t_m}(k)\phi_k(\mathbf{S}_{t_m})|\overline{\mathbf{S}}_{t_0}\right] - \mathbb{E}\left[\sum_{k=1}^{K} \widehat{\alpha}^1_{t_m}(k)\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_0}\right]}{x} \\
&= \lim_{x\to 0} \sum_{k=1}^{K} \widehat{\alpha}^1_{t_m}(k) \frac{\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\overline{\mathbf{S}}_{t_0}\right] - \mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_0}\right]}{x} \\
&= \sum_{k=1}^{K} \widehat{\alpha}^1_{t_m}(k) \frac{\partial \mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_0}\right]}{\partial S^1_{t_0}}.
\end{aligned} \tag{3.13}$$

At time t_0, there is only one grid point, which is equal to the initial state $\mathbf{S}_{t_0}$, so there is just one bundle, i.e $v = 1$, at t_0. As x can be arbitrarily small, it's safe to assume that $\overline{\mathbf{S}}_{t_0}$ and $\mathbf{S}_{t_0}$ will lie in this same bundle, and therefore the same approximation $Z(\mathbf{S}_{t_1}, \widehat{\alpha}^1_{t_1})$ can be used to compute the continuation value for these two states. Additionally, it is assumed that $\alpha^1_{t_1}$ is independent of $\mathbf{S}_{t_0}$. Only the case when $\mathbf{S}_{t_0}$ is not in the early-exercise region is considered, as the case otherwise is trivial. The derivative $\frac{\partial \mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_0}]}{\partial S^1_{t_0}}$ is either known, or is computed using numerical methods. SGBM can compute the delta (and similarly gamma) simultaneously with the direct estimator, at no additional computational cost.

Proposition 4. *Under the assumptions considered,*

$$\lim_{v\to\infty} \lim_{N\to\infty} |\widehat{\Delta}^1_{t_0} - \Delta^1_{t_0}| = 0.$$

Proof: The numerator in the r.h.s. of Equation (3.13) approaches (3.12) in the limiting case following Theorem 1.

Duality

Haugh and Kogan (2004)[36] and Rogers (2002)[71] proposed the dual formulation for pricing Bermudan options. For an arbitrary adapted super-martingale process $\mathscr{M}_t$, it follows that,

$$\begin{aligned} V_{t_0}(\mathbf{S}_{t_0}) &= \sup_{\tau} \mathbb{E}\left[\frac{h_\tau}{B_\tau}|\mathbf{S}_{t_0}\right] \\ &= \sup_{\tau} \mathbb{E}\left[\frac{h_\tau}{B_\tau} + \mathscr{M}_\tau - \mathscr{M}_\tau|\mathbf{S}_{t_0}\right] \\ &\leq \mathscr{M}_{t_0} + \sup_{\tau} \mathbb{E}\left[\frac{h_\tau}{B_\tau} - \mathscr{M}_\tau|\mathbf{S}_{t_0}\right] \\ &\leq \mathscr{M}_{t_0} + \mathbb{E}\left[\max_t\left(\frac{h_t}{B_t} - \mathscr{M}_t\right)|\mathbf{S}_{t_0}\right], \end{aligned}$$

which gives us the upper bound of the option price $V_{t_0}(\mathbf{S}_{t_0})$. Thus, the dual problem is to minimize the upper bound with respect to all adapted super-martingale processes, i.e.,

$$\overline{V}_{t_0}(\mathbf{S}_{t_0}) = \inf_{\mathscr{M}\in\Pi}\left(\mathscr{M}_{t_0} + \mathbb{E}\left[\max_t\left(\frac{h_t}{B_t} - \mathscr{M}_t\right)|\mathbf{S}_{t_0}\right]\right), \tag{3.14}$$

where Π is the set of all adapted super-martingale processes. Haugh and Kogan (2004)[36] show that when the super-martingale process, $\mathscr{M}_t$ in (3.14) coincides with the discounted option value process, $\frac{V_t(\mathbf{S}_t)}{B_t}$, the upper bound $\overline{V}_{t_0}(\mathbf{S}_{t_0})$ equals the true price for the Bermudan option. This suggests that a tight upper bound can be obtained by an accurate approximation of $V_t(\mathbf{S}_t)$, i.e. by defining $\mathscr{M}_t$ so that when the approximate option price, $\widehat{V}_t(\mathbf{S}_t)$, coincides with the exact price $V_t(\mathbf{S}_t)$, $\mathscr{M}_t$ equals the discounted process $\frac{V_t(\mathbf{S}_t)}{B_t}$. An obvious choice for $\mathscr{M}_t$ is then:

$$\mathscr{M}_{t_0} = \underline{V}_{t_0}(\mathbf{S}_{t_0}), \tag{3.15}$$

$$\mathscr{M}_{t_{m+1}} = \mathscr{M}_{t_m} + \frac{\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}})}{B_{t_{m+1}}} - \frac{\widehat{Q}_{t_m}(\mathbf{S}_{t_m})}{B_{t_m}}. \tag{3.16}$$

Therefore, corresponding to each simulated path, the martingale $\mathscr{M}_t$ is constructed as,

$$\mathscr{M}_{t_{m+1}}(n) = \mathscr{M}_{t_m}(n) + \frac{\widehat{V}_{t_{m+1}}(\mathbf{S}_{t_{m+1}}(n))}{B_{t_{m+1}}} - \frac{\widehat{Q}_{t_m}(\mathbf{S}_{t_m}(n))}{B_{t_m}}.$$

The upper bound, $\overline{V}_{t_0}$, corresponding Equations (3.15) and (3.16) is then given by

Set 1 :
$S^1_{t_0} = 40$, $K = 40$, $r = 0.06$, $\sigma = 0.2$, $T = 1$, $M = 50$.
Set 2 :
$S^\delta_{t_0} = 40$, $K = 40$, $r = 0.06$, $q_\delta = 0$, $\sigma_\delta = 0.2$, $\rho_{ij} = 0.25$, $T = 1$, $M = 10$.
Set 3 :
$S^\delta_{t_0} = 100$, $K = 100$, $r = 0.05$, $q_\delta = 0.1$, $\sigma_\delta = 0.2$, $\rho_{ij} = 0.0$, $T = 3$, $M = 9$.

Table 3.1: Parameter values used in the examples.

$$\begin{aligned}\overline{V}_{t_0}(\mathbf{S}_{t_0}) &= \mathbb{E}\left[\max_t\left(\frac{h_t}{B_t} - \mathscr{M}_t\right)|\mathbf{S}_{t_0}\right] \\ &= \frac{1}{N}\sum_{n=1}^{N}\max_t\left(\frac{h(\mathbf{S}_t(n))}{B_t} - \mathscr{M}_t(n)\right),\ t \in [t_0,\ldots,t_M].\end{aligned}$$

In the limiting case, as the number of paths and bundles go to infinity, the approximations $\widehat{V}_{t_m}(\mathbf{S}_{t_m})$ and $\widehat{Q}_{t_{m-1}}(\mathbf{S}_{t_{m-1}})$ approach their corresponding exact values and then $\overline{V}_{t_0}(\mathbf{S}_{t_0})$ will coincide with $V_{t_0}(\mathbf{S}_{t_0})$. Through numerical examples it becomes evident that in case of SGBM, with an increasing number of paths and bundles, tight upper bounds can be obtained without the need of sub-simulation.

3.2 Numerical experiments

This section illustrates the performance of SGBM by pricing different types of Bermudan options. By means of numerical examples the rate of convergence of the option price, when different bundling schemes are used, is compared. The computational performance of SGBM is also compared against the standard LSM [53] for different options. Numerical results are used to show that the direct estimator has a significantly lower variance when compared to the path estimator.

All underlying assets follow the standard single and multi-asset Black-Scholes model (geometric Brownian motion, GBM). For the examples considered, unless specified otherwise, $N = 50{,}000$ paths are used for computing the direct estimator and the early exercise policy and $N_L = 200{,}000$ paths for computing lower bounds using the path estimator. In case of the *k-means clustering* algorithm, first a set of 5000 training paths is used to obtain the optimal centroids corresponding to each bundle. The code for SGBM is implemented in MATLAB and the computations were performed on an Intel(R) Quad-Core 2 GHz processor with 4 GB RAM.

The parameter sets used for the different problems are listed in Table 3.1

Experiment with bundles

This section discusses by examples the role of bundling in computing the option price. First, with a basic Bermudan put on a single asset is considered, followed by high-dimensional options with different types of payoffs.

Bermudan options on single asset

Consider a Bermudan put on a single asset, where the risk-neutral asset price follows the stochastic differential equation

$$dS_t = rS_t dt + \sigma S_t dW_t, \tag{3.17}$$

r being the continuously compounded risk-free interest rate, σ the annualized volatility (both chosen to be constant), W_t is the standard Brownian motion. The option is exercisable a finite number of times per year, M, up-to and including the final expiration time $t_M = T$. As basis functions $\phi_k(\mathbf{S}_{t_m}) = \mathbf{S}_{t_m}^{k-1}$, where $k = 1,\ldots,4$, are used.

The continuation value, as given by Equation (3.8), requires us to compute

$$\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}}(n)\right] = \mathbb{E}\left[(\mathbf{S}_{t_m})^{k-1}|\mathbf{S}_{t_{m-1}}(n)\right], k = 1,\ldots,4.$$

These moments can be written down as:

$$\mathbb{E}\left[(\mathbf{S}_{t_m})^{k}|\mathbf{S}_{t_{m-1}}(n)\right] = \left(\mathbf{S}_{t_{m-1}}(n)e^{\left(r+\frac{(k-1)\sigma^2}{2}\right)(t_m-t_{m-1})}\right)^k,$$

which can be simply computed. The convergence of the two bundling schemes, i.e. *k-means clustering* and *recursive bifurcation* and their corresponding computational times are compared. Figures 3.3(a) and (b) show the convergence with an increasing number of bundles for the two methods. A highly accurate option reference price is computed using the COS method [29]. Figure 3.3(c) compares the total computational time, i.e. the combined time taken to compute the direct estimator as well as the path estimator using the two bundling approaches. Rapid convergence with increasing bundles for lower computational time makes recursive bifurcation the preferred method in this case.

Geometric Basket Option

Consider the pricing of a Bermudan option on the geometric average of several assets. As is well known, it is possible to reduce this problem to a one-dimensional problem, which can then be priced accurately using the COS method [29], thus

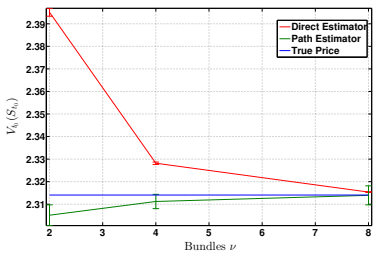

(a)

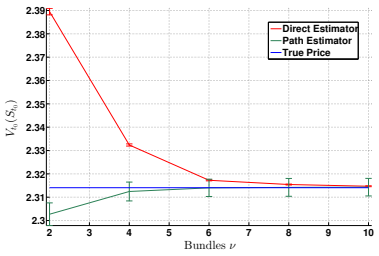

(b)

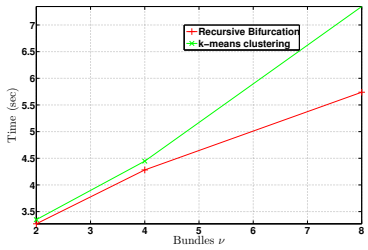

(c)

Figure 3.3: Option price for a put on a single asset, corresponding to different numbers of bundles used, when (a) recursive bifurcation scheme is used and (b) k-means clustering is used to partition the state space. Parameter set 1 in Table 3.1 is employed. The reference option value is 2.3140. (c) total computational time, i.e. time to compute the direct estimator plus time to compute the path estimator.

providing a benchmark result for the algorithm. A geometric average put option on d assets has intrinsic value:

$$h(\mathbf{S}_{t_m}) = K - (\prod_{\delta=1}^{d} S_m^{\delta})^{\frac{1}{d}}.$$

The asset prices are assumed to follow correlated geometric Brownian motion processes, i.e.

$$\frac{dS_t^{\delta}}{S_t^{\delta}} = (r - q_{\delta})dt + \sigma_{\delta} dW_t^{\delta}, \tag{3.18}$$

where each asset pays a dividend at a continuous rate of q_{δ}. $W_t^{\delta}, \delta = 1,\ldots,d$, are standard Brownian motions and the instantaneous correlation coefficient between W_t^i and W_t^j is ρ_{ij}.

As logical basis functions,

$$\phi_k(\mathbf{S}_{t_m}) = \left((\prod_{\delta=1}^{d} S_{t_m}^{\delta})^{\frac{1}{d}} \right)^{k-1}, k = 1,\ldots,5,$$

are used here.

The continuation value, as given by Equation (3.8), requires us to compute,

$$\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}}(n)\right] = \mathbb{E}[(\prod_{\delta=1}^{d} S_{t_m}^{\delta})^{\frac{k-1}{d}}|\mathbf{S}_{t_{m-1}}(n)],\ k = 1,\ldots,5.$$

These moments can directly be computed as:

$$\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}}(n)\right] = \left(P_{t_{m-1}}(n)e^{\left(\bar{\mu}+\frac{(k-1)\bar{\sigma}^2}{2}\right)\Delta t}\right)^{k-1},$$

where,

$$P_{t_{m-1}}(n) \quad = \quad \left(\prod_{\delta=1}^{d} S_{t_{m-1}}^{\delta}(n)\right)^{\frac{1}{d}},\ \bar{\mu} = \frac{1}{d}\sum_{\delta=1}^{d}\left(r - q_\delta - \frac{\sigma_\delta^2}{2}\right),\ \bar{\sigma}^2 = \frac{1}{d^2}\sum_{p=1}^{d}\left(\sum_{q=1}^{d} C_{pq}^2\right)^2.$$

C being the Cholesky factor of the covariance matrix and $C_{pq}-$ matrix element p,q.

Figure 3.4 shows the convergence of the direct estimator and path estimator with an increasing number of bundles for the different bundling schemes for a five-dimensional problem. For the *recursive bifurcation of the reduced state space*, the geometric average of the asset prices is used to map the high-dimensional state space onto a single-dimensional space, which is then partitioned using the recursive bifurcation. Partitioning of the reduced state space leads to better convergence when compared to the other two bundling approaches. Additionally, unlike recursive bifurcation of the high-dimensional state space, which results in 32 bundles in the first bifurcation itself, recursive bifurcation of the reduced state space has a greater flexibility on the choice of the number of bundles that can be created.

Figure 3.4(d) compares the computational time corresponding to different numbers of bundles and the bundling schemes used. While for all three schemes the option price is computed within a few seconds, the recursive bifurcation appears to be computationally most efficient.

Figure 3.5(a) compares the convergence for a geometric basket on 15 assets, when *k-means clustering* and *recursive bifurcation of reduced state space* were used for bundling. Figure 3.5 (b) gives the corresponding total computational time for the two methods. Recursive bifurcation is not used for this case as even with one iteration of the method 2^{15} bundles would be obtained, and a significant number of these bundles will not contain sufficient number of grid points.

Table 3.2 compares the results with LSM for the 10 and 15 assets cases. The results for SGBM correspond to 32 bundles, generated using the different bundling schemes. The standard errors for the direct estimator are much lower than those of the path estimator, even though 4 times more paths are used for computing the path estimator. The variance reduction factor, i.e. ratio of variance of path estimator to the direct estimator ranges between 50 to 100. The LSM results are also quite good for this test case.

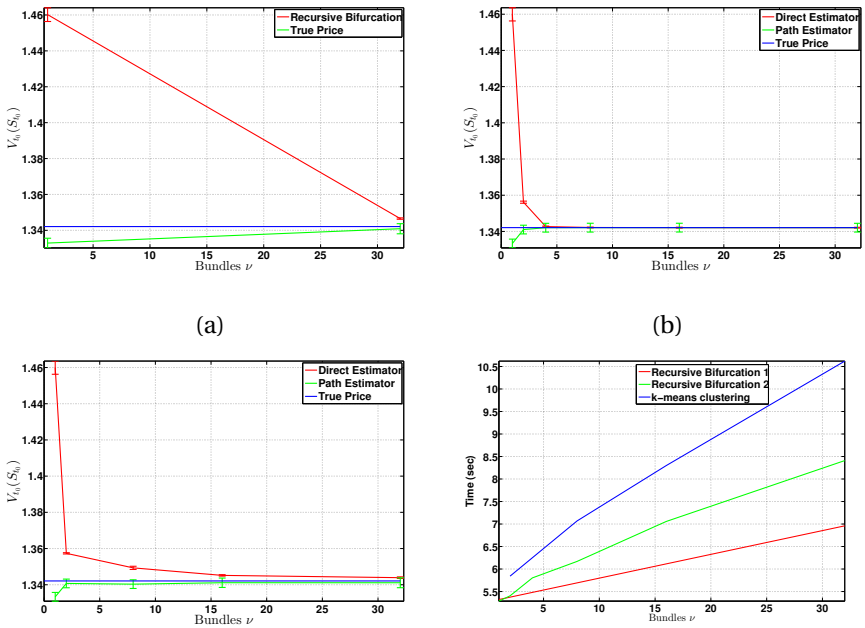

(a) (b)

(c) (d)

Figure 3.4: Option value for a put on geometric average of five assets, when (a) recursive bifurcation in high-dimensions is used and (b) recursive bifurcation on the reduced state space is used, (c) k means clustering is used, to partition the state space. (d) gives the total computational time (recursive bifurcation 2 is recursive bifurcation of reduced state space). Parameter set 2 from Table 3.1 is used. The reference option price is 1.3421.

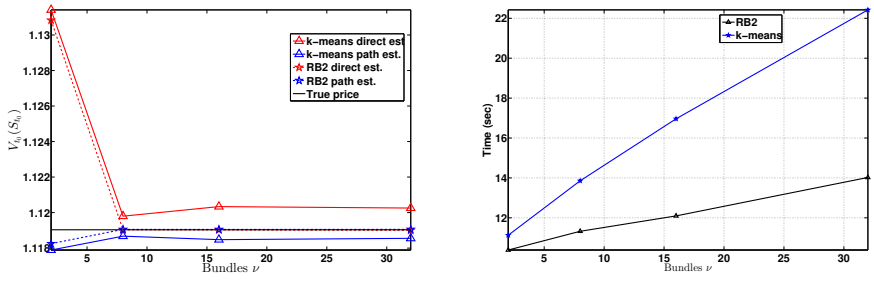

(a) (b)

Figure 3.5: (a) Option value for a put on geometric average of 15 assets, when *k-means clustering* and *recursive bifurcation of the reduced state space* (RB2) are used for bundling. (b) Total computational time for the two approaches. Parameter set 2 from Table 3.1 is used. The reference option price is 1.1190.

	Direct estimator (s.e.)	Path estimator (s.e.)	Computation Time (secs)
d=10 assets:			
SGBM (RB 2)	1.1781 (0.0002)	1.1779 (0.0024)	8.10
SGBM (KM)	1.1795 (0.0004)	1.1777 (0.0027)	16.68
LSM		1.1765 (0.0023)	5.67
d=15 assets:			
SGBM (RB 2)	1.1190 (0.0002)	1.1190 (0.0023)	14.02
SGBM (KM)	1.1202 (0.0003)	1.1185 (0.0027)	22.42
LSM		1.1164 (0.0019)	7.15

Table 3.2: Comparison between SGBM (using different bundling schemes) and LSM for a geometric basket option on 10 and 15 assets. The values in parenthesis are standard errors. Computation time includes the time to compute the policy or direct estimator and the path estimator. RB2 stands for recursive bifurcation of the reduced state space and KM for k-means clustering. The reference option price for the 10 assets case is 1.1779 and for the 15 assets case is 1.1190

Arithmetic Basket Option

Next we consider the case of a Bermudan option on the arithmetic mean of five assets, where the asset prices follow the dynamics given by Equation (3.18).

The arithmetic average put option on d-assets is governed by the intrinsic value function,

$$h(\mathbf{S}_m) = K - \left(\frac{1}{d} \sum_{\delta=1}^{d} S_m^{\delta} \right).$$

As basis functions, the immediate choice,

$$\phi_k(\mathbf{S}_{t_m}) = \left(\frac{1}{d} \sum_{\delta=1}^{d} S_{t_m}^{\delta} \right)^{k-1}, k = 1, \ldots, 5,$$

is used.

The continuation value, as given by Equation (3.8), requires us to compute,

$$\mathbb{E}\left[\phi_k(\mathbf{S}_{t_m}) | \mathbf{S}_{t_{m-1}}(n)\right] = \mathbb{E}\left[\left(\frac{1}{d} \sum_{\delta=1}^{d} S_{t_m}^{\delta} \right)^{k-1} | \mathbf{S}_{t_{m-1}}(n) \right], k = 1, \ldots, 5. \tag{3.19}$$

The expectation in Equation (3.19) can be expressed as a linear combination of moments of the geometric average of the assets, i.e.

$$\left(\sum_{\delta=1}^{d} S_{t_m}^{\delta} \right)^{k} = \sum_{k_1+k_2+\cdots+k_d=k} \binom{k}{k_1, k_2, \ldots, k_d} \prod_{1 \le \delta \le d} (S_{t_m}^{\delta})^{k_\delta},$$

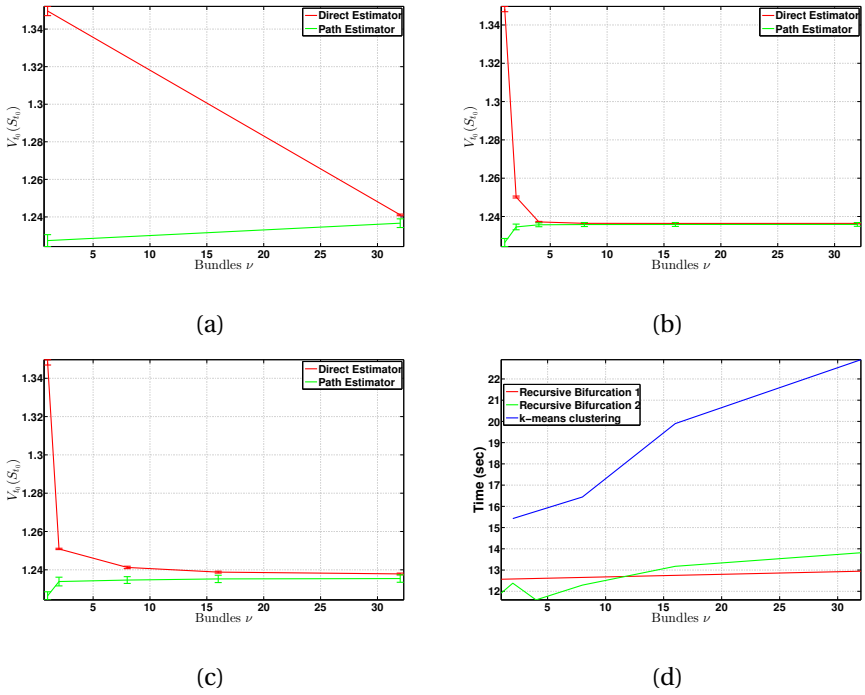

Figure 3.6: Option value for a put on the arithmetic average of five assets, when (a) recursive bifurcation in high-dimensions is used and (b) recursive bifurcation on the reduced state space is used (c) k-means clustering is used. (d) Total computational time corresponding to the different bundling schemes considered. Parameter Set 2 from Table 3.1 is used.

where,

$$\binom{k}{k_1, k_2, \ldots, k_d} = \frac{k!}{k_1!\, k_2! \cdots k_d!},$$

which can be computed in a straightforward way by Equation (3.19).

Figures 3.6(a) to (c) display the direct and path estimator values, for different numbers of bundles and bundling schemes. For *recursive bifurcation of reduced state space*, the arithmetic average of the asset prices is used to map the high-dimensional state space to the single-dimensional space, which is then partitioned using the recursive bifurcation scheme. Again, partitioning of the reduced state space leads to better results when compared to the other two bundling schemes.

Figure 3.6 (d) compares the computational time corresponding to different numbers of bundles. The total computational time is always less than a minute, with the recursive bifurcation being computationally most efficient, while k-means clustering being the most expensive. Figure 3.7 (a) shows the convergence of the method with an increasing number of bundles for an arithmetic basket on 15 assets. Only *k-means clustering* and *recursive bifurcation along the reduced state space* are con-

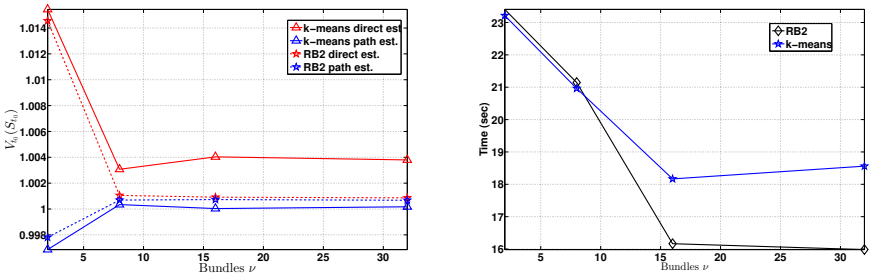

(a) (b)

Figure 3.7: (a) Option value for a basket on the arithmetic mean of 15 assets, corresponding to different numbers of bundles. (b) Computational time corresponding to different numbers of bundles. The parameter values from Set 2 in Table 3.1 are employed.

	Direct estimator (s.e.)	Path estimator (s.e.)	Computation Time (secs)
d=10 assets:			
SGBM (RB 2)	1.0624 (0.0003)	1.0615 (0.0018)	22.60
SGBM (KM)	1.0669 (0.0008)	1.0604 (0.0022)	26.34
LSM		1.0611 (0.0020)	5.13
d=15 assets:			
SGBM (RB 2)	1.0008 (0.0002)	1.0006 (0.0019)	15.98
SGBM (KM)	1.0038 (0.0004)	1.0002 (0.0020)	18.56
LSM		1.0009 (0.0026)	7.20

Table 3.3: Comparision between different bundling schemes used in SGBM and LSM for an arithmetic basket option on 10 and 15 assets. The values in parenthesis are standard errors. Computation time includes the time to compute the policy or direct estimator and the path estimator. RB2 stands for recursive bifurcation of the reduced state space and KM for k-means clustering. Parameter values are taken from Set 2 in Table 3.1.

sidered for bundling, as in the case of *recursive bifurcation* even with a single iteration 2^{15} bundles would be created, with a significant number of bundles having an insufficient number of grid points. Figure 3.7 (c) displays the corresponding computational time, which is still in seconds.

Table 3.3 compares the results with those obtained using LSM for the 10 and 15 assets case. The results reported for SGBM correspond to the case of 32 bundles. The standard error for the direct estimator is significantly lower than that for the path estimator, even though 4 times more paths were used in the latter case. The variance reduction factor, i.e. the ratio of variance of the path estimator to the direct estimator again ranges between 50 to 100.

Duality based upper bounds

Duality-based upper bounds on the option price can be useful when only an approximation for $\mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X]$ can be found and it cannot be computed exactly. This is because for the direct estimator to be an upper bound on the true price, Theorem 2 assumes $\mathbb{E}[\phi_k(\mathbf{S}_{t_m})|\mathbf{S}_{t_{m-1}} = X]$ to be computed *exactly*. An example for such a case is an option on max of more than two assets, discussed below.

Max Option

A Bermudan max-option is a discretely-exercisable option on multiple underlying assets, whose pay-off depends on the maximum among all asset prices. The intrinsic value function for the call option on the max of d assets is given by:

$$h(\mathbf{S}_{t_m}) = \max(S^1_{t_m}, \ldots, S^d_{t_m}) - K,$$

Consider the case where the asset prices follow correlated geometric Brownian motion processes, as given by Equation (3.18).

As basis functions,

$$\phi_k(\mathbf{S}_{t_m}) = \left(\log(\max(S^1_{t_m}, \ldots, S^d_{t_m}))\right)^{k-1}, k = 1, \ldots, 5,$$

$$\phi_6(\mathbf{S}_{t_m}) = \left(\prod_{\delta=1}^{d} S^\delta_{t_m}\right)^{\frac{1}{d}},$$

$$\phi_{6+\delta}(\mathbf{S}_{t_m}) = S^\delta_{t_m}, \delta = 1, \ldots, d,$$

are used. The additional basis functions compared to the previous examples are required because of the non-linear payoff surface.

The continuation value, as given by Equation (3.8), requires us to compute,

$$\mathbb{E}\left[\left(\log\left(\max\left(S^1_{t_m}, \ldots, S^d_{t_m}\right)\right)\right)^{k-1} |\mathbf{S}_{t_{m-1}}(n)\right], \tag{3.20}$$

which can be done using Clark's algorithm [22]. Clark's algorithm (see Chapter 2, Section 2.4 for details on Clark's algorithm) computes the first four moments for the maximum of several correlated normal variates, as well as the correlation coefficient between the maximum of a pair and the third normal variate. The computed values, other than the maximum of two normal variates, will be *approximations*, and therefore for options with more than two assets the direct estimator will not be an upper bound. However, the upper bounds can still be computed using the approach of duality, as discussed in Section 3.1.

Duality-based upper bounds, together with the lower bound computed using the path estimator gives a valid confidence interval within which the true option price lies. The confidence interval is constructed as:

$$\left[\underline{V}_{t_0}(\mathbf{S}_{t_0}) - 1.96\frac{\widehat{s}_L}{\sqrt{N_S}},\ \overline{V}_{t_0}(\mathbf{S}_{t_0}) + 1.96\frac{\widehat{s}_H}{\sqrt{N_S}}\right],$$

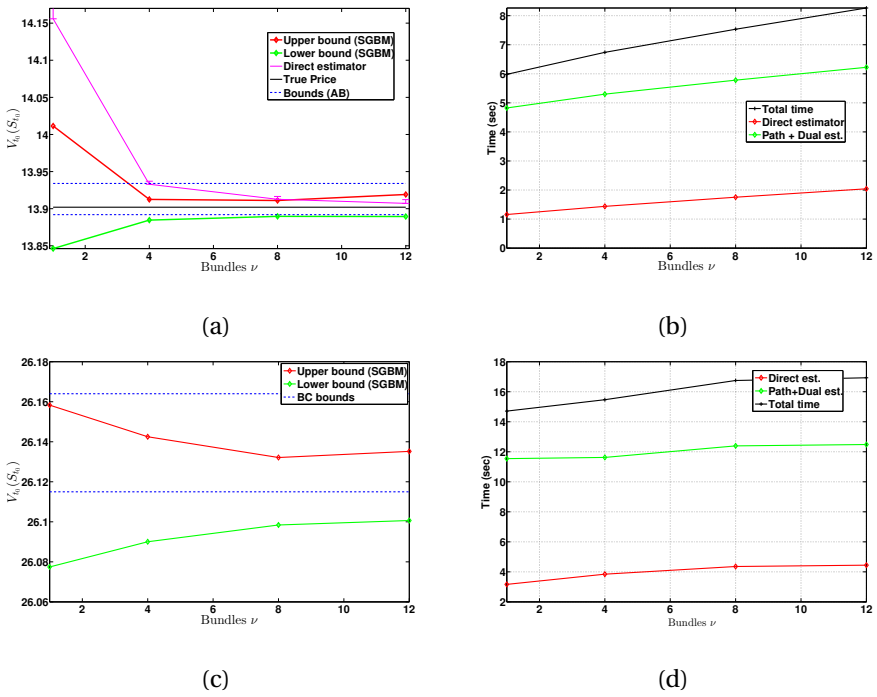

Figure 3.8: (a) Option price corresponding to different numbers of bundles for a call on maximum of two assets. (b) Computational time for the two asset case. (c) Confidence interval corresponding to different numbers of bundles used for a call on the max of five assets. (d) Computational time for the five asset case. Parameter values are taken from Set 3 in Table 3.1.

where $\widehat{s}_L$, is the sample standard deviation for the path estimator and $\widehat{s}_H$, is the sample standard deviation for the duality-based upper bound estimator. These standard deviations are based on N_S independent simulation trials, and for the examples here $N_S = 30$, are taken.

Figure 3.8(a) shows the convergence of the direct estimator with an increasing number of bundles, and the corresponding confidence interval constructed using the duality-based upper bounds and the lower bounds found using the path estimator. For comparison, the confidence interval reported in Andersen and Broadie (2004)[2], which is also based on the dual formulation, is plotted. Figure 3.8(b) shows the time taken to compute the direct estimator, lower (using path estimator) and upper bounds (using duality). For all cases the total computational time is less than a minute.

Figure 3.8(c) displays the confidence interval constructed for the case of the maximum of 5 assets. Also plotted is the confidence interval reported by Broadie and Cao (2009)[15] for the same problem. The corresponding computational times are plotted in Figure 3.8(d). The computation time for duality-based upper bounds, reported in the literature are usually in several minutes, and in comparison SGBM's

S_0	SGBM Direct est. (s.e.)	SGBM Path est. (s.e.)	SGBM Dual est. (s.e.)	SGBM 95% CI	Literature 95% CI	Binomial
d=2 assets:						
90	8.069 (0.013)	8.067 (0.020)	8.105 (0.086)	[8.059 8.136]	[8.053 8.082]	8.075
100	13.907 (0.005)	13.898 (0.023)	13.906 (0.035)	[13.889 13.919]	[13.892 13.934]	13.902
110	21.351 (0.004)	21.338 (0.022)	21.339 (0.023)	[21.329 21.347]	[21.316 21.359]	21.345
d=3 assets:						
90	11.223 (0.006)	11.247 (0.035)	11.483 (0.284)	[11.235 11.585]	[11.265 11.308]	11.29
100	18.650 (0.008)	18.654 (0.037)	18.761 (0.134)	[18.641 18.809]	[18.661 18.728]	18.69
110	27.554 (0.011)	27.537 (0.038)	27.592 (0.158)	[27.523 27.648]	[27.512 27.663]	27.58
d=5 assets:						
90	16.521 (0.009)	16.620 (0.037)	16.625 (0.036)	[16.607 16.637]	[16.620 16.653]	
100	26.086 (0.011)	26.129 (0.044)	26.132 (0.044)	[26.113 26.148]	[26.115 26.164]	
110	36.743 (0.013)	36.753 (0.045)	36.754 (0.045)	[36.737 36.770]	[36.710 36.798]	

Table 3.4: Option values for call on maximum of 2, 3 and 5 assets, with parameter values taken from Set 3, in Table 3.1.The values reported are for 12 bundles created using *k-means clustering* algorithm. The reference confidence interval for the two and three asset case are taken from Andersen and Broadie (2004) [2], and for the five asset case from Broadie and Cao (2009) [15].

time in seconds seems efficient. Results for $d = 2,3$, and 5 assets are summarized in Table 3.4.

Computing Greeks

In this section the sensitivity of the option price is computed using SGBM. As an example a call on the maximum of d assets option, which was discussed above is considered and the Greeks Delta, $\Delta^1_{t_0} (= \partial V_{t_0}(\mathbf{S}_{t_0})/\partial S^1_{t_0})$, and Gamma, $\Gamma^{11}_{t_0} (= \partial^2 V_{t_0}(\mathbf{S}_{t_0})/\partial (S^1_{t_0})^2)$ are computed for it.

First consider the case of the max on two assets, as for this case the exact Greeks can be computed using the 2D COS method of Ruijter and Oosterlee (2012)[74]. Figures 3.9(a) and (b) compare the exact Greeks computed using the 2D COS method with results from SGBM for different numbers of exercise opportunities. The error in the delta values ranges between 0.2% to 0.4%, which is quite modest in comparison to the ones obtained using the traditional bumping method. The errors for gamma values are higher and range between 3% to 7%. Results for SGBM are comparable to those reported by Kaniel *et al.* (2004) [47], however, the computation time for SGBM is less than a minute while for the latter it can be hours.

Figures 3.8(c) and (d) compare the Greeks computed for different numbers of exercise opportunities with the bounds reported in Wang and Calfisch (2009) [87]. It is clear that although the SGBM values reported lie within the confidence interval, the results are not accurate enough when number of exercise opportunities increases.

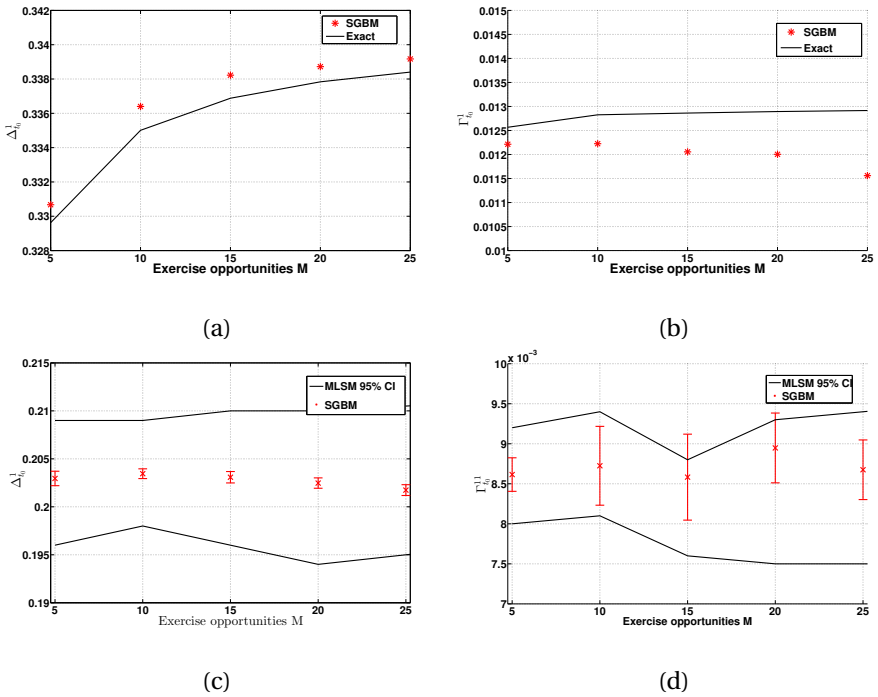

Figure 3.9: (a) Delta and (b) Gamma values, for a call on maximum of two assets, for different number of exercise opportunities. (c) Delta and (b) Gamma values, for a call on maximum of five assets, corresponding to different number of exercise opportunities. Parameter values are taken from Set 3 in Table 3.1.

One of the reasons for this is that with an increasing number of exercise opportunities there is an error build up while the option values are estimated moving backwards in time. SGBM can provide an accurate approximation for the option price sensitivities with minimal computational effort.

3.3 Conclusion

This chapter introduced the Stochastic Grid Bundling Method (SGBM) for approximating the value of Bermudan options by simulation. SGBM is a hybrid of regression-based and bundling-based Monte Carlo methods, and appears to be computationally at least as attractive as existing methods. Basic proofs for convergence are discussed, however, the rate of convergence, especially regarding the number of bundles used is only dealt with by numerical examples.

This chapter illustrates SGBM's performance using a number of realistic examples, including the valuation of options on the geometric and the arithmetic mean, and a maximum of assets option on a basket of assets. The computational time for the

method is comparable to the least squares method [53], but a higher accuracy, not just at the final time step, but also at intermediate time steps makes it a suitable candidate for computing upper bounds using duality-based methods. Another advantage of SGBM is that it can be used for fast approximation of the option price sensitivities. The SGBM method described is intuitive, easy to implement and accurate.

Part II

Real Option Analysis: Investments in Nuclear Power Plants

CHAPTER

4 Valuing Modular Nuclear Power Plants in Finite Time Decision Horizon

The contents of this chapter have appeared in [42]. Deregulation of the electricity market has been driven by the belief in increased cost-efficiency of competitive markets. There is a need for valuation methods to make economic decisions for investment in power plants in these uncertain environments. Kessides (2010) [49] emphasizes the use of real options analysis (ROA) to estimate the option value that arises from the flexibility to wait and choose between further investment in the nuclear plant and other generating technologies as new information emerges about energy market conditions.

There is an increased interest in small and medium sized reactors (SMRs) as an alternative to large Gen III type nuclear reactors [11]. This is primarily because the former have, amongst other benefits, comparatively low upfront costs and flexibility of ordering due to its modular nature [19]. When comparing the economy of large reactors and SMRs, it's necessary to take into account the value of flexibility arising from modular construction, which traditional valuation methods like net present value (NPV) cannot. As the decisions to order new reactors would be planned for finite time horizons, there is a need to adapt the real option valuation for modular construction, as proposed by Gollier *et al.* (2005) [35], to a finite time horizon. The case studies presented here are not only important for the construction of power plants but they are also relevant for a larger class of decision questions in which flexibility due to modularity and economy of scale play an important role.

This chapter focuses on the value of flexibility that arises from the modular construction of SMRs. The approach used here is similar to that used by Gollier *et al.* (2005) [35], where the firm needs to make a choice between a single high capacity reactor (1200 MWe) or a flexible sequence of modular SMRs (4×300 MWe). This chapter, however, considers a *finite time horizon* before which the investment decision should be made. In a competitive market the firms cannot delay an investment decision for ever and need to decide before the anticipated entry of a competitor, or before a technology becomes obsolete. Also utilities need to meet the electricity demand with some minimum reliability, which restricts their decision horizon to finite time. The investment rules, such as the optimal time to start construction

and the real option value of the investment, can differ significantly with changing decision horizons.

Real options can be priced with methods used for pricing American- or Bermudan-style financial options which are discussed in the previous chapters. This chapter uses the stochastic grid method (SGM) [40] (see Chapter 2 for details regarding SGM) , for computing the real option values of modular investment decisions. SGM has been used to price Bermudan options in [40] with results comparable to those obtained using the well-known least squares method (LSM) of Longstaff and Schwartz (2001) [53], but typically with tighter confidence intervals using fewer Monte Carlo paths. The option values are computed by generating stochastic paths for electricity prices, and thus with uncertain future cash flows. The optimal electricity price at which a new module should be ordered is found as an outcome of the real option pricing.

In the sections to follow the problem of modular investment in nuclear power plants is stated and compared with its counter part in the financial world. In Section 4.1 the problem and its real option formulation is described. In Section 4.2 the mathematical formulation behind the problem is discussed. Section 4.3 gives the description of how the stochastic grid method can be used to value a modular real option. Section 4.4 describes in detail the application of the method to the nuclear case. Finally, Section 4.5 gives some concluding remarks and possible future research questions that need to be addressed.

4.1 Problem Context

A competitive electricity market where the price of electricity follows a stochastic process is considered in this chapter. The utility faces the choice of either constructing a single large reactor of 1200 MWe, or sequentially constructing four modules of 300 MWe each. The total number of series units is denoted by n. Unit number i is characterized by discounted averaged cost per KWh equal to θ_i , its construction time is denoted by C_i and the lifetime of its operation by L_i. Both construction and lifetime are expressed in years. It is assumed that different modules are constructed in sequence, where,

1. similar to the case of Gollier [35], construction of module $i+1$ cannot be decided until construction of unit i is over, i.e. no overlap in construction of modules is allowed.

2. a more relaxed constraint where the construction of unit $i+1$ can be decided from any time subsequent to the start of construction of unit i.

A constant discount rate denoted by r is assumed here.

The utility here needs to take a decision to start the construction of the modules within a finite time horizon, denoted by T_i for the ith module. In terms of financial

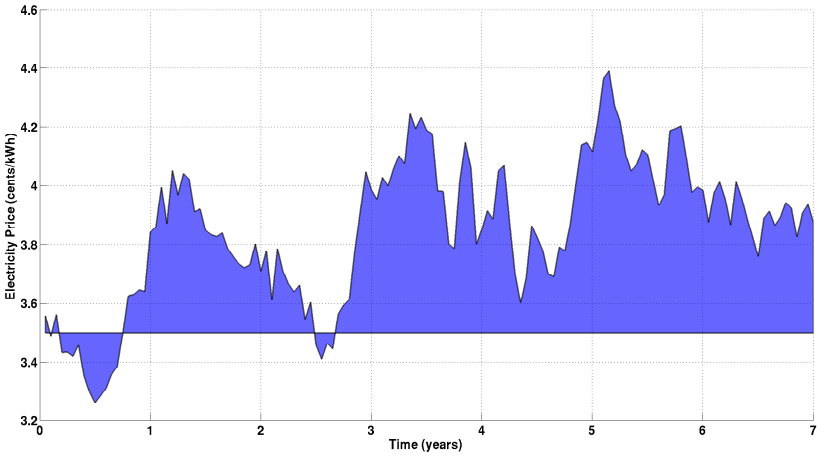

Figure 4.1: The area between the electricity path (starting at 3.5 cents/kWh) and cost of operation = 3.5 cents/kWh, gives cash flow for the reactor.

options, T_i represents the expiration time for the 'option to start the construction of the ith module '. Unlike financial options, it's difficult to quantify the *expiration time* for real options, and it is usually taken as the expected time of arrival of a competitor in the market, or time before which the underlying technology becomes obsolete. In case of an electricity utility, it also represents the time before which the utility needs to set up a plant to meet the electricity demand with certain reliability. Reliability is measured as the probability of the number of unplanned outages in a year with one of the reasons for such an outage being demand exceeding available generation.

The real option formulation

The problem of modular construction can be formulated as a multiple exercise Bermudan option. Consider the stochastic process, P_t, to be the process which models the electricity price. The payoff, $h_i(P_t = x)$, for the real option problem is the expected net cash flows per unit power of electricity sold through the lifetime of module i, when it gets operational at time t and state $P_t = x$.

Figure 4.1 illustrates the profit from the sale of electricity for one realized electricity price path. The cost of operation, θ, in the illustration is 3.5 cents/kWh and the area between the electricity path and θ gives the profit from the sale of electricity. The value of interest is the expected profit, i.e. the mean profit from all possible electricity paths in the future. This expected profit (or net cash flow) is the payoff, $h_i(P_t)$, for the real option.

The *revenue*, R_i, for the ith module, for every unit power of electricity sold through its lifetime L_i, starting construction at time t, when the electricity price is $P_t = x$, can be written as

$$R_i(P_t = x) = \mathbb{E}\left[\int_{t+C_i}^{t+C_i+L_i} e^{-ru} P_u \, du | P_t = x\right]. \tag{4.1}$$

R_i is the discounted expected gross revenue over all possible electricity price paths. The revenue starts flowing in once the construction is over, and therefore the range for the integral starts from $t + C_i$. and lasts as long as the plant is operational, i.e. until $t + C_i + L_i$. Similarly, the *cost* of operating the ith module, K_i, through its lifetime for every unit power of electricity generated, is:

$$K_i = \int_{t+C_i}^{t+C_i+L_i} e^{-ru} \theta_i \, du. \tag{4.2}$$

Here θ_i, the cost of operating the reactor per kWh is assumed to be constant. Therefore, the net discounted cash flow, for module i, is given by:

$$h_i(P_t = x) = R_i(P_t = x) - K_i. \tag{4.3}$$

Equations (4.1) to (4.3) give the expected profit from the sale of electricity through the life of the nuclear reactor. Equation (4.3) is the mean profit from all possible electricity paths in the future.

The expiration time T_n is the time before which the last module should be ordered. The optimal exercise policy $\pi = \{\tau_n, \ldots, \tau_1\}$, is then defined by the determination of the optimal times for starting the construction of different modules, with τ_i, the optimal time for starting the construction of module i, so that the net cash flow from the different modules is maximized.

Electricity price model

The uncertain parameter in our pricing model is the electricity price. Modelling electricity spot prices is difficult primarily due to factors like:

- Lack of effective storage, which implies electricity needs to be continuously generated and consumed.
- The consumption of electricity is often localized due to constraints of relatively poor grid connectivity.
- The prices show other features like daily, weekly and seasonal effects, that vary from place to place.

Models for electricity spot prices have been proposed by Pilipovic (1997) [66], Lucia and Schwartz (2002) [55]. Barlow (2002) [6] develops a stochastic model for electricity prices starting from a basic supply/demand model for electricity. These models

are focused on the short term fluctuations of electricity prices which helps better pricing of electricity derivatives.

As decisions for setting up power plants look at long term evolution of electricity prices, like Gollier [35], the basic Geometric Brownian Motion (GBM) as the electricity price process is used here. However, it should be noted that within our modelling approach it is easy to include other price processes.

Geometric Brownian Motion

If at any time t the electricity price is given by P_t cents/kWh, then the electricity price process is given by

$$dP_t = \alpha P_t dt + \sigma P_t dW_t, \tag{4.4}$$

where α represents the constant growth rate of P_t, σ is the associated volatility and W_t is the standard Brownian motion. In our model it is assumed that α and σ are constant. A closed form solution to the above SDE can be obtained using Ito's lemma, and is given by:

$$P_t = P_0 e^{\left(\left(\alpha - \frac{\sigma^2}{2}\right)t + \sigma\sqrt{t}Z\right)}, \tag{4.5}$$

where Z is a standard normal variable. Also it can be seen that the above process has a log-normal distribution, i.e. $\log(P_t)$ has a Gaussian distribution with mean

$$\mathbb{E}\left[\log(P_t)\right] = \log(P_0) + \left(\alpha - \frac{\sigma^2}{2}\right)t,$$

and variance

$$Var(\log(P_t)) = \sigma^2 t.$$

4.2 Mathematical Formulation

The optimal time to order[1] a new reactor under uncertain electricity price is solved using dynamic programming, where an optimal solution is found recursively moving backwards in time. Here the problem stated above is re-framed as a dynamic programming problem.

[1] The optimal time to order is often called "optimal stopping time". In the case of sequential modular construction optimal stopping time would refer to the time when the option to delay the construction to the next time step terminates.

Dynamic programming formulation

In order to construct all the modules at the optimal time, using Bellman's principle of optimality, optimal decisions need to be taken, starting from the last reactor. The optimal decision time for each of the reactors is computed as well, starting from their respective expiration times and moving backwards in time to the initial state. The expiration time for ordering the ith module is given by

$$T_i = T_n - \sum_{k=i}^{n-1} C_k. \tag{4.6}$$

This constraint comes from the restriction that a new reactor can be ordered once all the prior ordered reactors have been constructed. Here T_n is the expiration time for the option to start the construction of the last module and C_i is the construction time in years for the ith module.

At the expiration time for the last module the firm does not have the option to delay the investment. Therefore, the decision to start the construction is taken at those electricity prices for which the expected NPV of the last module is greater than zero. The option value of the last module at the expiration time is then given by:

$$V_n(t_M = T_n, P_{t_M}) = \max(0, h_n(P_{t_M})). \tag{4.7}$$

At time t_m, $m = M-1, \cdots, 0$, the option value for the last of the series of reactors is the maximum between immediate pay-off h_n and its continuation value Q_n. The continuation value is the expected future cash flow if the decision to construct the reactor is delayed until the next time step. The reactor is constructed if at the given electricity price the net present value is greater than the expected cash flows if the reactor is constructed sometime in the future. This can be written as:

$$V_n(t_m, P_{t_m}) = \max\left(h_n(P_{t_m}), Q_n(t_m, P_{t_m})\right),\ m = 0, \ldots, M-1. \tag{4.8}$$

Given the present state P_{t_m}, the continuation value, or, in other words, the discounted cash flows if the decision to start the construction is delayed for the last reactor is,

$$Q_n(t_m, P_{t_m}) = e^{-r(t_{k+1}-t_m)} \mathbb{E}\left[V_n(t_{k+1}, P_{t_{k+1}}) | P_{t_m}\right]. \tag{4.9}$$

Once the option value at each time step for the last module is known, we move on to modules $n-1, \ldots, 1$. At the expiration time for the i-th module, the decision to start its construction is taken when the combined NPV of the present reactor and the expected future cash flow from the optimally constructed modules $i+1, \ldots, n$ is

greater than zero. Therefore, the option value for the ith module at its expiration time T_i is given by:

$$V_i(T_i, P_{T_i}) = \max\big(0, h_i(P_{T_i}) + Q_{i+1}(T_i, P_{T_i})\big), \tag{4.10}$$

where h_i gives the direct future cash flow from the ith module and $Q_{i+1}(T_i, P_{T_i})$ gives the expected cash flow from the optimal construction of modules $i+1,\ldots,n$, given the information P_{T_i}. The option value for the module at time step t_m, where $t_m < T_i$, is given by

$$V_i(t_m, P_{t_m}) = \max(h_i(P_{t_m}) + Q_{i+1}(t_m, P_{t_m}), Q_i(t_m, P_{t_m})), \tag{4.11}$$

i.e. the decision to start the construction of module i is taken if the cash flow from its immediate construction (given by $h_i(P_{t_m})$) and the expected cash flow from the modules $i+1,\ldots,n$, constructed optimally in the future (modeled by $Q_{i+1}(t_m, P_{t_m})$) , is greater than the expected cash flows from the modules $i,\ldots,n$, if the decision to start its construction is delayed to the next time step (given by $Q_i(t_m, P_{t_m})$). The expected cash flow if the decision to start the construction of modules $i,\ldots,n$ is delayed to the next time step is given by:

$$Q_i(t_m, P_{t_m}) = e^{-r(t_{k+1}-t_m)}\mathbb{E}\left[V_i(t_{k+1}, P_{t_{k+1}})|P_{t_m}\right]. \tag{4.12}$$

The option value, $V_i(t_m, P_{t_m})$, at time t_m for constructing the module i not only carries the information about the cash flows from module i, but also about the cash flows from the optimal construction of the modules $i+1,\ldots,n$ in the future.

For sequential modular construction the payoff for module i is given by $h_i(P_{t_m}) + Q_{i+1}(t_m, P_{t_m})$. The payoff does not only contain h_i, the direct discounted revenue from module i, but also Q_{i+1}, the value of the new option to start or delay the construction of new modules, that opens up with the construction of module i.

4.3 Stochastic Grid Method for multiple exercise options

The real option problems discussed here, have financial counterparts, i.e. the Bermudan options and multiple exercise Bermudan options. A Bermudan option gives the holder the right, but not obligation, to exercise the option once, on a discretely spaced set of exercise dates. A multiple exercise Bermudan option, on the other hand, can be exercised multiple times before the option expires. Pricing of Bermudan options, especially for multi-dimensional processes is a challenging problem owing to its path-dependent settings.

Consider an economy in discrete time defined up to a finite time horizon T_n. The market is defined by the filtered probability space $(\Omega, \mathscr{F}, \mathscr{F}_t, \mathbb{P})$. Let P_t, with $t = t_0, t_1, \ldots, t_m = T_n$, be an $\mathbb{R}^d$-valued discrete time Markov chain describing the state

of the economy, the price of the underlying assets and any other variables that affect the dynamics of the underlying. Here $\mathbb{P}$ is the risk neutral probability measure. The holder of the multiple exercise Bermudan option has n exercise opportunities, that can be exercised at $t_0, t_1, \ldots, t_m$. Let $h_i(P_t)$ represent the payoff from the ith exercise of the option at time t and underlying state P_t. The time horizon for the ith exercise opportunity is given by T_i.

Define a policy, π, as a set of stopping times $\tau_n, \ldots, \tau_1$ with $\tau_n < \ldots < \tau_1$, which takes values in $t_0, \ldots, t_m = T_n$, and τ_i determines the time where the ith remaining exercise opportunity can be used. The option value when there are n early exercise opportunities remaining is then found by solving an optimization problem, i.e. to find the optimal exercise policy, π, for which the expected payoff is maximized. This can be written as:

$$V_n(t_0, P_{t_0} = x) = \sup_{\pi} \mathbb{E}\left[\sum_{k=0}^{n} h_k(P_{\tau_k}) | P_{t_0} = x\right]. \tag{4.13}$$

In simple terms, Equation (4.13) states that of all possible policies for ordering the reactor in the given decision horizon, the real option value is computed using the one which maximizes the expected future cash flows.

The problem of pricing Bermudan options with multiple exercise opportunities has been dealt with by Meinshausen and Hambly (2002) [58], with generalizations by Bender (2008) [9], Aleksandrov and Hambly(2008) [1] and Schoenmaker (2009) [75], who use the dual representation for such pricing problems. Chiara *et al.* (2007) [21] apply the multiple exercise real options in infrastructure projects. They use a multi-least-squares Monte Carlo method for determining the option value.

The problem of sequential modular construction stated above can be solved using the stochastic grid method [40] (see Chapter 2). It is chosen[2], because:

- The stochastic grid method (SGM) can efficiently solve the multiple exercise Bermudan option problem;
- SGM can be used to compute the sensitivities of the real option value;
- The method can be easily extended to higher dimensions;
- The method doesn't depend on the choice of the underlying stochastic process;
- Improved confidence intervals are obtained with fewer paths when compared to LSM.

Although the problem considered here is one-dimensional, with the electricity price as the stochastic variable, a typical real option problem tends to be high-dimensional

[2]We didn't use SGBM here as in the chronology of development of this thesis the method was yet not defined

with several underlying stochastic terms. A proper choice of pricing method would be one which can be extended to higher dimensions in the future.

The stochastic grid method solves a general optimal stopping time problem using a hybrid of dynamic programming and Monte Carlo simulation. The method first determines the *optimal stopping policy* and a direct estimator for the option price. The optimal stopping policy for the ith module at time step t_m involves finding the critical electricity price $P^*_{t_m}$. When the market price of electricity is equal to the critical price, the value of delaying the construction of the module to the next time step is equal to the value of starting the construction immediately, i.e.,

$$Q_i(t_k, P^*_{t_k}) = h_i(P^*_{t_k}).$$

Therefore, the critical price is taken to be the largest "grid point" P_{t_k}, for which $Q_i(t_k, P_{t_k}) > h_i(P_{t_k})$. The module is ordered if the present market price of electricity is greater than the critical price for the given time step. Once the policy for all the time steps is known, SGM computes lower bound values, using a new set of simulated electricity paths, as the mean of the cashflows from each simulated path where the module is ordered following the policy obtained above.

SGM for multiple exercise Bermudan options begins by generating N stochastic paths for the electricity prices, starting from initial state P_0. The electricity prices realized by these paths at time step t_m constitute the grid points at t_m. The electricity price paths can be generated using Equation (4.5) here.

The pricing steps for SGM can be decomposed into two main parts, based on the recursive dynamic programming algorithm from the previous section.

- Parametrization of the option value: The option values at each grid point are converted into a functional approximation using *piece-wise regression*.
- Computation of the continuation value: The continuation value is computed using the conditional probability density function and the functional approximation of the option value at the next time step.

Parametrization of the option value

In order to obtain the continuation value for grid points at t_m, the functional approximations of the option value at t_{k+1}, needs to be determined. Once the option values at the grid points at t_{k+1} are known, the functional approximation is obtained using a *piece-wise least squares regression*. Therefore, the option value at a given time step is divided into two regions, separated by the critical electricity price $P^*_{t_{k+1}}$.

For the two segments the functional approximation is given by,

$$\widehat{V}_i(t_{m+1}, P_{t_{m+1}}) = \mathbf{1}_{\{P_{t_{m+1}} < P^*_{t_{m+1}}\}} \sum_{k=0}^{K-1} a_k P^k_{t_{m+1}} + \mathbf{1}_{\{P_{t_{m+1}} \geq \mathscr{X}^*_{t_{m+1}}\}} \sum_{k=0}^{K-1} b_k P^k_{t_{m+1}} \tag{4.14}$$

The expression $\mathbf{1}_{\{P_{t_{m+1}} < \mathscr{X}^*_{t_{m+1}}\}}$ is an indicator function whose value equals 1, if the argument $\{P_{t_{m+1}} < \mathscr{X}^*_{t_{m+1}}\}$, is true and it is 0 otherwise. Therefore, $\mathbf{1}_{\{P_{t_{m+1}} < \mathscr{X}^*_{t_{m+1}}\}}$ and $\mathbf{1}_{\{P_{t_{m+1}} \geq \mathscr{X}^*_{t_{m+1}}\}}$ group the grid points into two segments, separated by the critical electricity price. The results converge to the reference price when increasing number of segments are used (see [40]).

Computation of the continuation value

Once the functional approximations of the option values for modules i and $(i+1)$ are known for time step t_{m+1}, the continuation value for the ith module at t_m can be computed using Equation (4.12). In order to compute the expectation, $\mathbb{E}\left[V_i(t_{m+1}, P_{t_{m+1}})|P_{t_m}\right]$, the distribution function for $P_{t_{m+1}}$ given P_{t_m}, is required. This conditional distribution function, $f(P_{t_{m+1}}|P_{t_m} = x)$, for the GBM process is known in closed form. Therefore, the continuation value, or the value of the reactor if the decision to order it is delayed to the next time step, as given in Equation (4.12) can be written as:

$$\begin{aligned} \widehat{Q}_i(t_m, P_{t_m}) \quad = \quad & \int_{y \in [0, \mathscr{X}^*]} \left(\sum_{k=0}^{K-1} a_k y^k \right) f(y|P_{t_m} = x)\, dy \\ & + \int_{y \in (\mathscr{X}^*, \infty]} \left(\sum_{k=0}^{K-1} b_k y^k \right) f(y|P_{t_m} = x)\, dy. \end{aligned} \tag{4.15}$$

In a more generic case where the conditional distribution function is unknown, it can be approximated using the *Gram Charlier Series*. For more details on computing the continuation value, refer to [40].

4.4 Numerical Experiments

The case where an investor needs to decide between two projects, one involving a single large reactor of 1200 MWe and the other consisting of four modules of 300 MWe each, is considered here. The construction time and costs for the two projects, given in Table 4.1, are taken from the reference case by Gollier [35]. The discount rate is taken as 8% per annum, which is the OECD average, and the predicted growth rate of electricity price is 0% here. The cost of electricity production for the first unit is relatively expensive when compared to series units, as a large

part of the fixed costs for the modular assembly, like the land rights, access by road and railway, site licensing cost, connection to the electricity grid are carried by the first unit.

In the case of the modular project two different constraints are considered, in the two subsections to follow, i.e., the decision for construction of subsequent units can be made :

1. Once the *construction* of all prior units is completed (similar to the case considered by Gollier),
2. Once the *decision* for the construction of all prior units has been taken. Also, only one unit can be ordered at a given time step.

	Construction Time (months)	Discounted Average cost (cents/KWh)
Large Reactor	60	2.9
Modular Case		
Module 1	36	3.8
Module 2 to 4	24	2.5

Table 4.1: Construction time and discounted averaged costs used for the large reactor and the modular case.

Sequential construction: The case in Gollier [35]

In this test case *constraint 1* is applied for the construction of subsequent modules, i.e., the decision for the construction of a new module will not be made, unless the construction of all previous modules is finalized. By the SGM first an optimal investment policy and a direct estimator of the real option value of the project are obtained. The optimal policy gives the critical electricity price (as a function of time), above which a module should be ordered. At a given time a new module is ordered only when the present electricity price is higher than the corresponding critical price for the module under consideration and when all other constraints are satisfied. Once the optimal investment policy is obtained, a fresh set of electricity paths is generated, and at each of these simulated paths a new module is decided if the following conditions are satisfied:

1. All modules preceding the given module have been constructed.
2. The present electricity price is higher than critical price for ordering the given module.
3. The present time is within the decision horizon for the corresponding module.

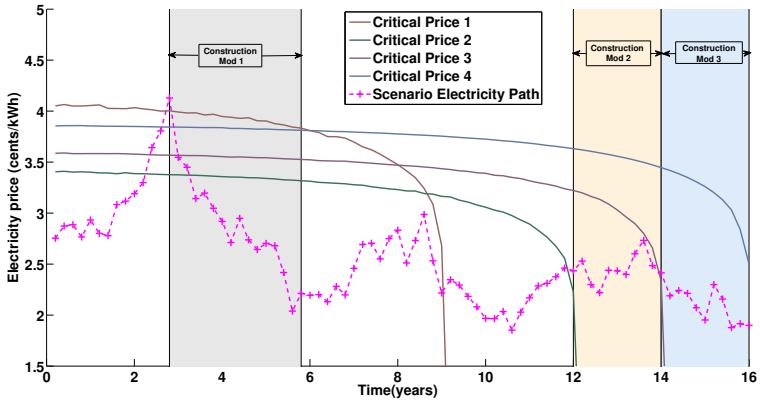

Figure 4.2: 0ptimal investment policy for ordering sequential modular reactors, with a sample scenario path. No overlap between the construction periods of two reactors is possible in this case.

4. The given module hasn't been ordered so far for the given path.

Figure 4.2 illustrates for a sample electricity path when a new module should be ordered. A module is ordered once the above conditions are satisfied and the revenue from this module is discounted back to the initial time. The mean of the discounted revenue for different paths from all the four modules gives the real option value of the project. For a single large reactor the steps followed are the same, except that condition 1 is not required.

As the case in [35] corresponds to an infinite horizon decision problem with exercise opportunities, it is compared with an increasing *finite time* decision horizon. Figure 4.3 compares the real option value of the modular project with the reference value in [35]. The option value of the project doesn't increase much with an increasing number of exercise opportunities per year, however it increases significantly with an increasing decision horizon. From Figure 4.3 it's clear that the real option value of a modular project with a realistic decision horizon is lower than the value obtained in [35], where an infinite decision horizon is assumed. In other simulations, not reported here, it is found that the option value of the modular project with the same parameters, but with a decision horizon of 100 and 200 years and four exercise opportunities per year, has an option value between 390 to 400 Euro/kW, which is already closer to the infinite horizon values.

Figure 4.4 then compares the optimal investment policy for the first module with the corresponding policy in [35]. It shows clearly the effect of a finite decision horizon on the investment policy. As one approaches the final decision time, the value of waiting (given by the continuation value) reduces which lowers the threshold electricity price at which a new module should be ordered. However, in the case of

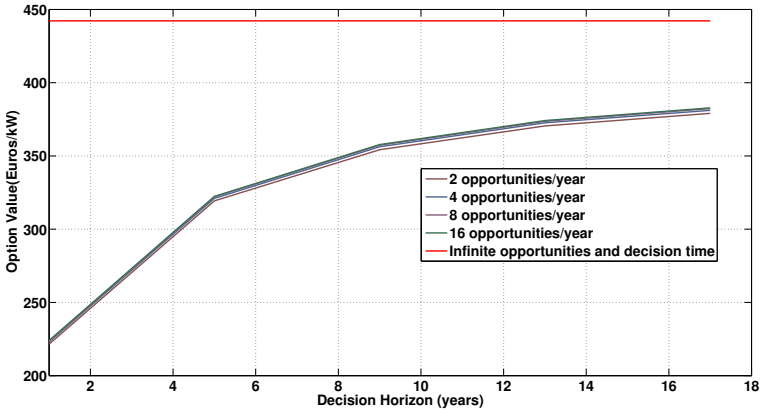

Figure 4.3: Real option value of a modular project with increasing decision horizon and exercise opportunities per year compared with the value obtained by Gollier. The initial price of electricity is 3 cents/kWh and volatility in electricity price is 20%.

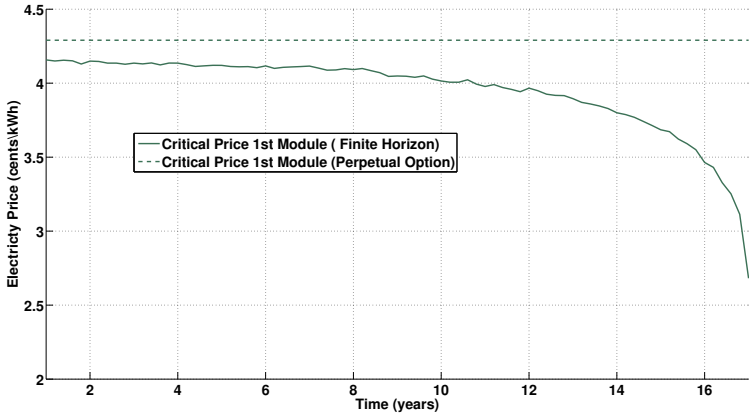

Figure 4.4: Critical price at which the first module should be ordered, comparison between finite time and infinite time horizon.

an infinite decision time horizon, the optimal policy or threshold electricity price remains constant with time.

Comparison of two projects with different decision times

The real option values of the two projects, i.e. the single large reactor and the sequence of small modular units, are compared next for increasing decision time horizon and uncertainty in electricity prices. The construction costs and times for

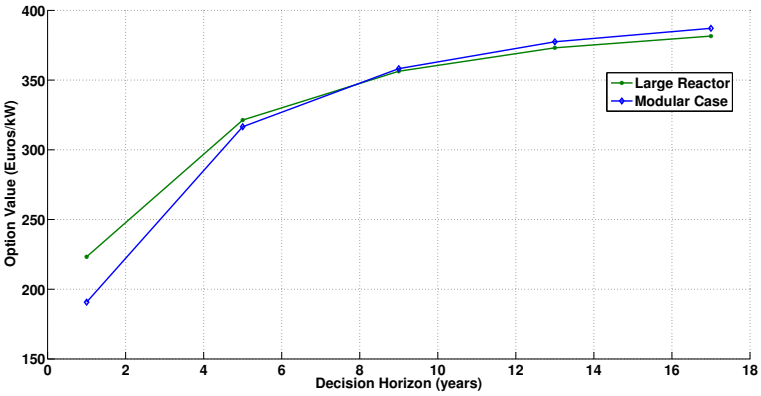

Figure 4.5: Real option value for the large reactor and the modular project for different decision horizons when the initial price of electricity is 3 cents/kWh.

the reactors are taken from Table 4.1. Based on Equation (4.6), the corresponding decision horizon for the construction of the first module is taken between 1 to 17 years. For the single large reactor the decision horizon is taken the same as that for the first module. One of the advantages of a modular construction is that the increasing demand can be met gradually, which allows the spreading of the decisions to a longer time horizon, possibly without significant gaps in demand and supply. Therefore, although the decision horizon for the first module can be small, the decision horizon for the entire project can be longer.

Figure 4.5 compares the option value for the single large reactor with that for the modular project for increasing decision time. When the decision horizon is small it is optimal to opt for the large reactor whereas for longer decision horizons modular projects appear more profitable. An insight into the reason why modular projects are better for longer decision time is given by Figure 4.6, which shows the expected cashflow from the four units for increasing decision time. When the decision horizon is small, the expected cashflow from the first module can be negative, which is the case when the decision horizon for the first module is one year. For short decision times, the first unit needs to be ordered to keep open the option to order more profitable subsequent modules. As the time approaches the final decision time for the first unit, it can be ordered even if the expected revenue from its construction would be negative. With increasing decision horizon the investor can wait longer and order the modules at more profitable electricity prices. Figure 4.7 shows fractions of the scenario paths for which different numbers of modules are ordered. When the decision horizon is small, the project is only partially completed for a large number of scenario paths, making it unprofitable. However, for longer decision times, the fraction of the scenario paths for which all the four modules are constructed increases, while the fraction of partially completed project reduces.

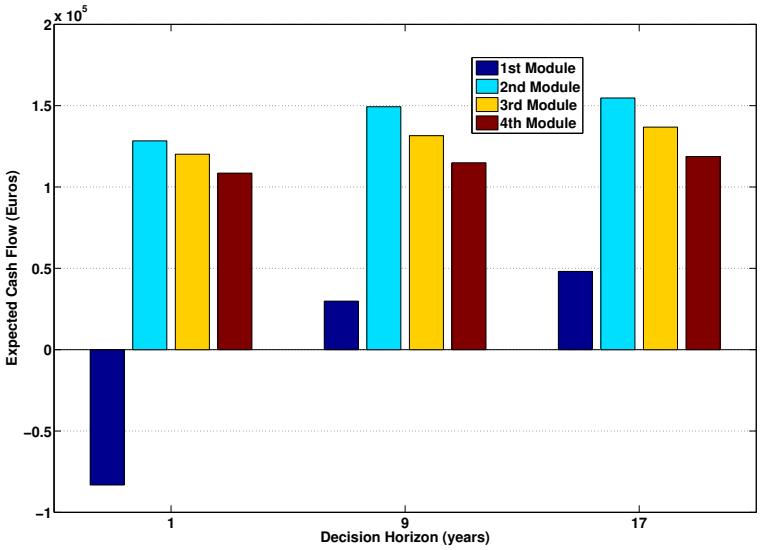

Figure 4.6: Cash flow from different modules with increasing decision time. The initial price of electricity is 3 cents/kWh and the volatility in electricity prices is 20%.

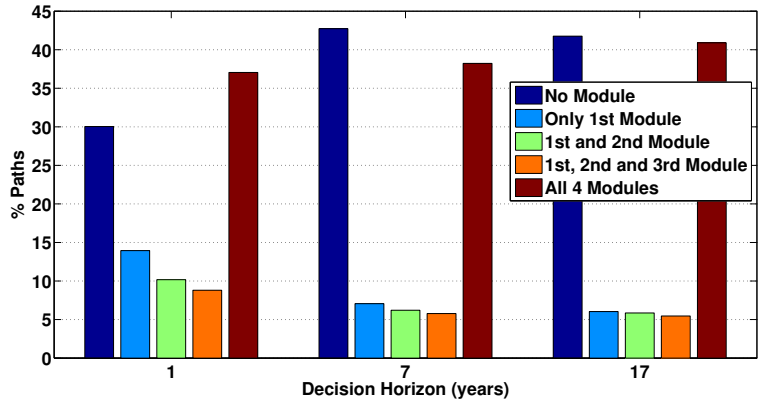

Figure 4.7: Fraction of modules ordered in the end for different scenario paths with increasing decision time. The initial price of electricity is 3 cents/kWh and the volatility in electricity prices is 20%.

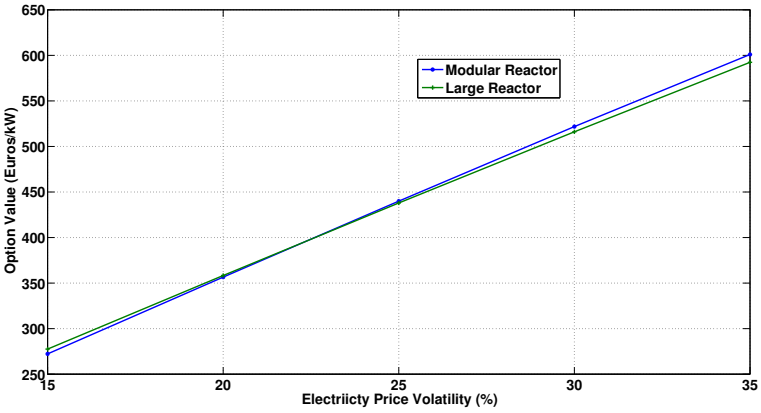

Figure 4.8: Real option value for the large reactor and the modular project for different volatilities when the decision horizon is 9 years and the initial price of electricity is 3 cents/kWh.

Comparison of two projects and different volatilities in electricity prices

A parameter to be considered when deciding between a single large reactor and the modular project is the uncertainty in electricity prices. The real option value of the two projects for increasing volatility in the electricity price is compared in Figure 4.8. When the volatility in the electricity price is low, the single large reactor project is more profitable, while for higher volatilities the modular project seems a better choice. An intuitive answer to this is that, for high volatilities, modular projects offer more flexibility, i.e., if the electricity price path at some point reaches unfavourable prices, the possibility to abandon module construction, with a few units already ordered, exists. Figure 4.9 shows the fraction of scenario paths for which the modular project finishes with different numbers of units ordered. As expected, the fraction of paths for which not all the four units are ordered increases with increasing volatility.

Figure 4.10 gives the expected cashflow from each unit for different volatilities of the electricity prices. In general, the option value of the project increases with an increasing uncertainty. A higher volatility reflects greater future price fluctuations (in either direction) in underlying electricity price levels. This expectation generally results in a higher option premium, especially if the option is exercised optimally. The cashflow from the first unit is smallest because it's the most expensive of the four units. When the cost of the units would be same, as is the case for units 2 to 4, the discount factor plays an important role, making the present value of units that are ordered first larger than the value from units further in the future.

For a firm it is not only important to know the real option value for making an investment decision, but also the *sensitivity*[3] of the value with respect to the param-

[3]Sensitivity analysis in financial options is performed by computing the derivatives with respect to

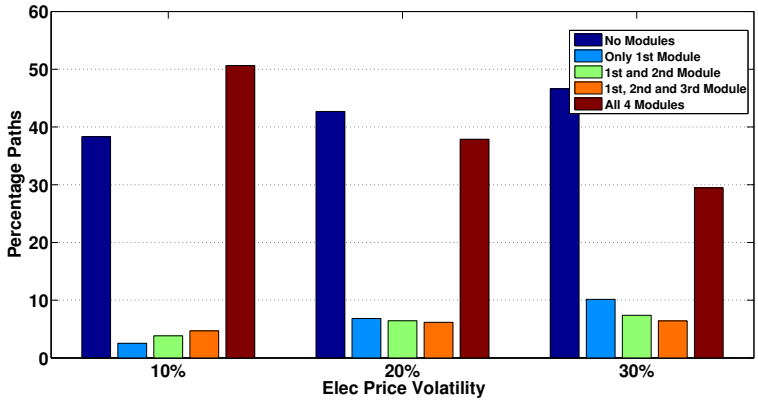

Figure 4.9: Fraction of modules ordered in the end for different scenario paths with increasing volatility. The initial price of electricity is 3 cents/kWh and the decision horizon for the first module is 9 years.

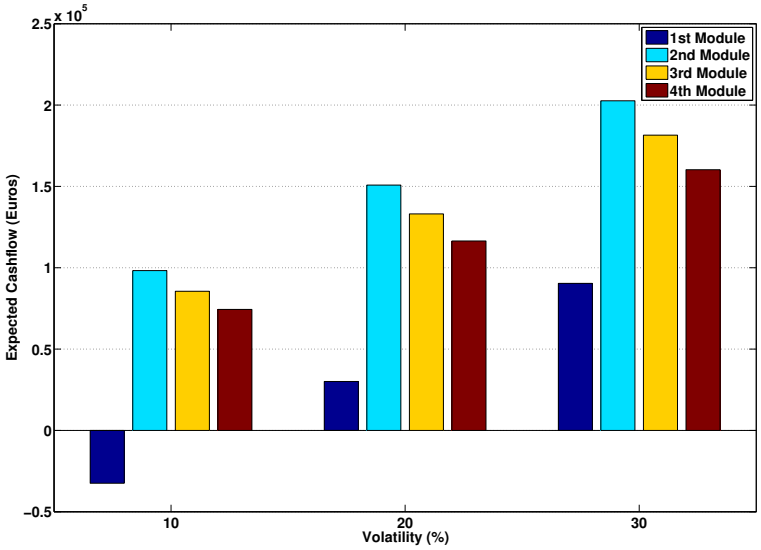

Figure 4.10: Expected cashflow from different units for different volatility values for electricity prices. The initial price of electricity is 3 cents/kWh and the decision horizon for the first module is 9 years.

eters chosen. The *delta* values, i.e. the ratios of the change in the real option value of the reactors to the change in the underlying electricity price, are computed here. High delta values imply that the investment decisions are sensitive to a changing electricity price. Figure 4.11 compares the delta values for the different modules. It can be seen that the delta values for the final module, as expected, converge to one, i.e. when it's optimal to order a new reactor the change in option value is proportional to the change in the electricity price. However, for each prior module the delta values converge to values less than one. For the first module a unit change in the electricity price changes the option price by a factor of 0.8. This makes the modular construction investment option generally more stable, i.e. even if the electricity price drops by 1 the value of investment changes by a factor of 0.8.

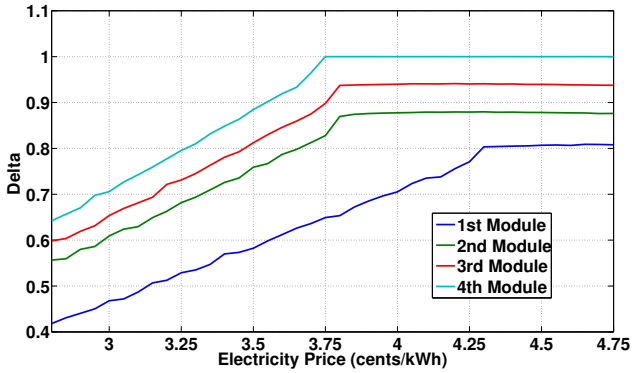

Figure 4.11: The delta values for the four modules when the decision horizon is 12 years and the volatility in electricity price is 20%.

Table 4.2 gives the critical threshold prices for constructing a single large reactor and for the different modules. The SGM results are compared with those computed using the COS method[4] [29] and with Gollier [35]. The COS method is a deterministic method, developed for computational finance applications, which can easily be applied to sequential investment decisions for single "assets". The Monte Carlo and the COS method give identical prices, which is a validation for the MC method, and a clear difference between the results obtained for finite time and those for infinite time horizon decisions in [35] can be seen. From these results it can be concluded that SGM is a good candidate for pricing finite time real option problems.

Modified case: multiple construction, sequential ordering

This section considers the two projects discussed above, except that now one of the constraints in the case of the modular project is relaxed, i.e., a new unit can

various parameters, and these derivatives together are referred to as the *Greeks*.

[4] I would like to thank Marjon Ruijter for the COS method results.

	Final Decision Time (years)	Isolated		Modular			
		LR	Unit 1	Unit 1	Unit2	Unit3	Unit 4
SGM	12	4.56	5.98	4.05	3.44	3.65	3.93
	17	4.60	6.02	4.18	3.50	3.69	3.96
	22	4.62	6.05	4.24	3.53	3.71	3.98
COS	12	4.56	5.98	4.10	3.46	3.65	3.93
	17	4.60	6.03	4.21	3.51	3.69	3.96
	22	4.62	6.05	4.25	3.53	3.71	3.98
Gollier	∞	4.75	6.23	4.29	3.57	3.79	4.10

Table 4.2: Critical threshold electricity prices (cents/KWh) at which new reactors should be ordered, for different decision horizons. There are twenty equally spaced exercise opportunities each year. The volatility of the electricity price is taken as 20 %.

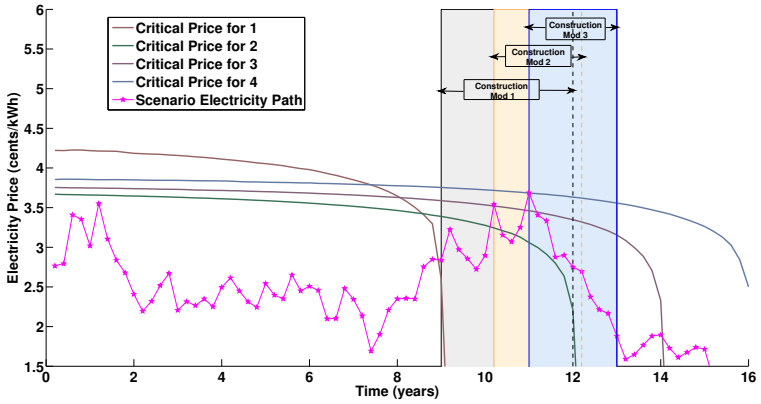

Figure 4.12: Optimal investment policy for ordering sequential modular reactors, with a sample scenario path. Overlap between construction period of two reactors is possible in this case.

be ordered if all previous units have been *ordered* (not constructed). Only a single unit can be ordered at any given time step. It's common practice to have parallel construction of different units in order to achieve cost savings, as it allows rotation of specialized labour between different units [62].

Figure 4.12 shows a scenario path and investment policy for a modular project with the above considerations. It can be seen that in this case overlap in the construction period of different units is possible.

Comparison of two projects and different decision times

Figure 4.13 compares the real option values of the two projects for different decision times. The decision time for the large unit is kept the same as that for the first unit in the modular project. As mentioned before, decisions of generation capacity expansion are based on meeting increasing electricity demands with a certain

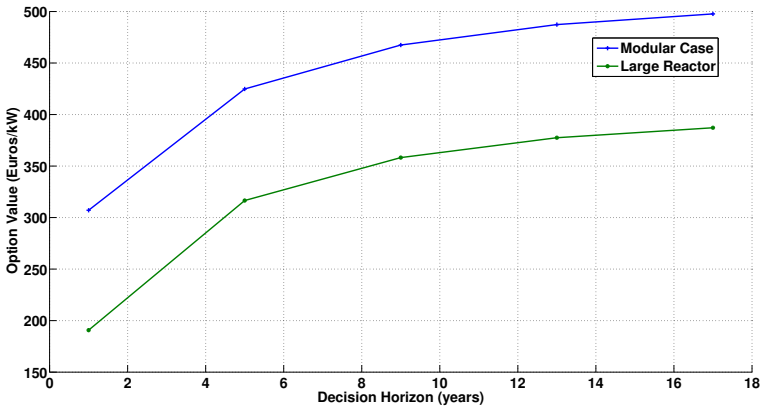

Figure 4.13: Real option value for the large reactor and the modular project for different decision horizons when the initial price of electricity is 3 cents/kWh.

minimum reliability and therefore the decision horizon is chosen to be the same for the first module and the large reactor. In this case the modular project appears more profitable than the single large reactor.

In order to detail the results obtained, the expected cashflow from different units of the modular project and the fraction of modules ordered for different decision times are computed. Figure 4.14 gives the expected cashflow from the four units for different decision times. It can be seen that the expected cashflow grows with decision time. When overlap in the construction periods of different units is allowed, the cashflow from the three similar costing units is almost the same. The reason for this is that most often the three units are ordered around the same time, and so the effect of discounting to present time is almost the same. An important reason for modular projects having higher real option value is that the effective decision horizon for the modular project is significantly longer than that of the large reactor (which is the same as that of the first unit). Another factor which adds up to the profitability of the modular project (when parallel construction is allowed) is that modular units have less construction time, which allows cashflow from the sale of electricity to start before it would start from the large reactor.

Figure 4.15 shows the fraction of different modules constructed by the end of the decision time for the modular project. It is clear that when the constraint of waiting for completion of a unit before ordering a new one is relaxed, that once the first unit is ordered, in most cases it results in all four units being ordered.

Comparison of two projects for different electricity price volatilities

Figure 4.16 compares the real option values of the two projects for different volatility in electricity prices under the relaxed constraint. It can be seen in this case that

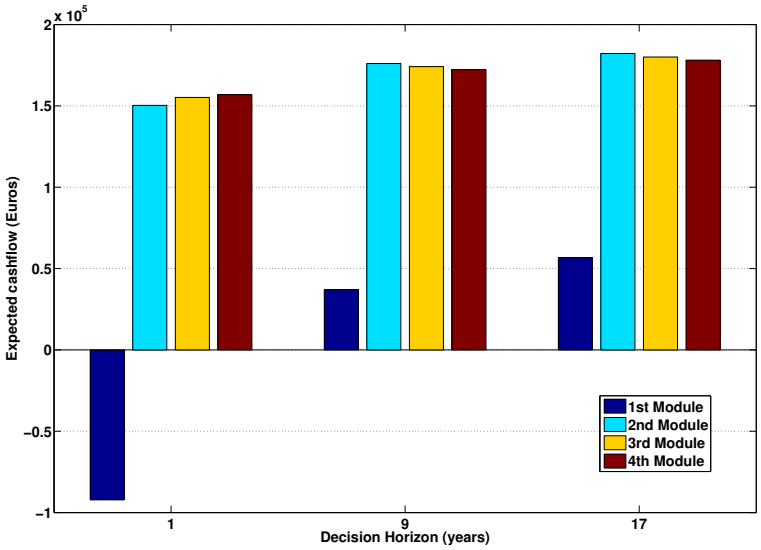

Figure 4.14: Cash flow from different modules with increasing decision time. The initial price of electricity is 3 cents/kWh and volatility in electricity prices is 20%.

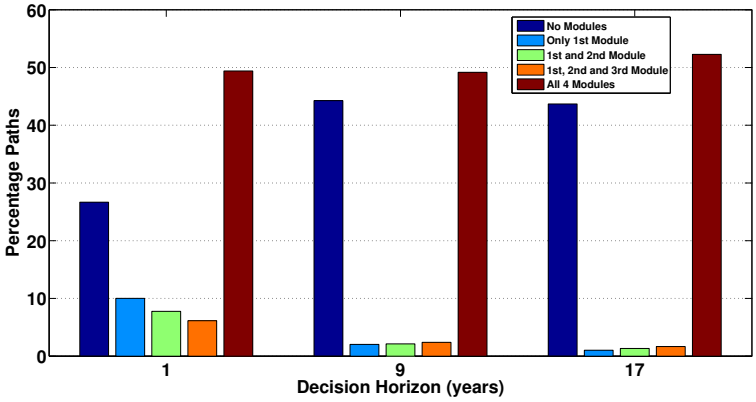

Figure 4.15: Fraction of modules ordered at the end for different scenario paths with increasing decision times. The initial price of electricity is 3 cents/kWh and the volatility in electricity prices is 20%.

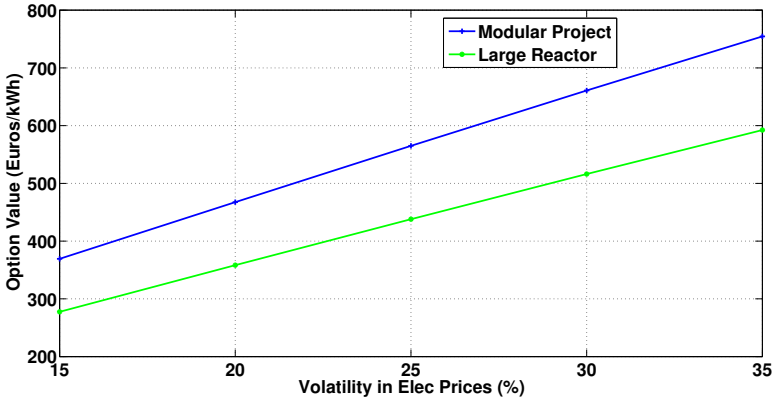

Figure 4.16: Real option value for the large reactor and the modular project for different volatilities when the decision horizon is 9 years and the initial price of electricity is 3 cents/kWh.

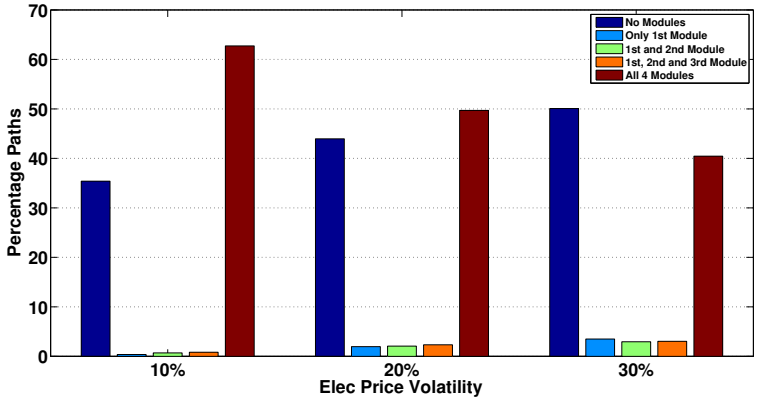

Figure 4.17: Fraction of modules ordered at the end for different scenario paths with increasing volatility. The initial price of electricity is 3 cents/kWh and the decision horizon for the first module is 9 years.

the modular project is always more profitable than a single large reactor.

When the constraint of ordering a new unit is relaxed from waiting until completion of all previous units to waiting until *ordering* of all previous units, most of the units are ordered around the same time, as can be concluded from the discussion above. However, this will not be the case when the uncertainty in the electricity price increases. Figure 4.17 shows the fraction of scenario paths for which different numbers of units are ordered by the end of the decision horizon. It is clear that with increasing uncertainty more often the project ends with fewer units than were planned initially.

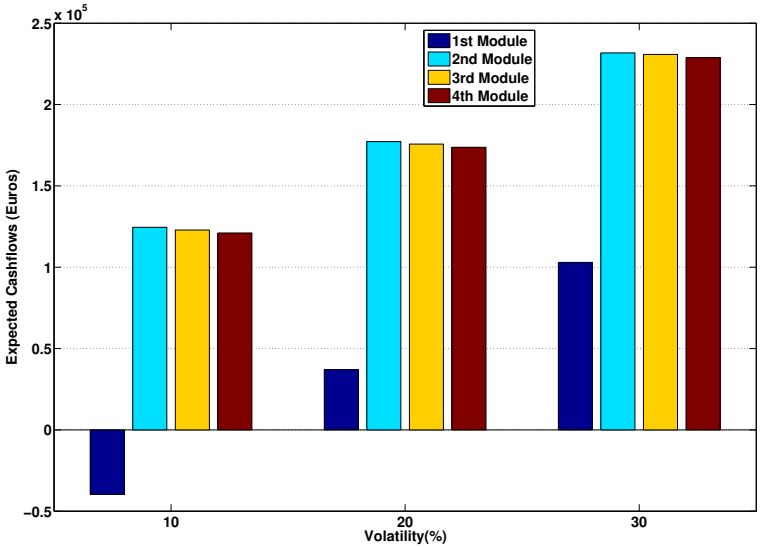

Figure 4.18: Expected cashflow from different units for different volatility of electricity prices. The initial price of electricity is 3 cents/kWh and the decision horizon for the first module is 9 years.

The cashflow for the project from different units however increases with increasing uncertainty in electricity prices, as can be seen in Figure 4.18

4.5 Conclusions

This chapter presented a flexible and accurate valuation method for computing real option values, and critical electricity prices, related to the construction of nuclear power plants. The Stochastic Grid Method from Chapter 2 was used to analyze some scenarios of interest for a utility when choosing nuclear reactors. In particular, the focus of the chapter is on the modular construction and finite decision horizons.

Some of the outcomes can be summarized as:

- SGM is a suitable Monte Carlo method for pricing real options, especially for valuing projects with modularity. The method has been validated against the deterministic COS method for a 1D test case.
- Most investment decision problems are governed by a finite decision time. It has been shown, in various numerical experiments, that decision-making in finite time may result in quite different scenarios compared to infinite time decision problems.

- When the modular project has a restriction on parallel construction of different units then:
  - the real option value of a single large reactor is typically higher when decision horizons are small.
  - For longer decision times modular projects may be more profitable.
  - With an increasing uncertainty in the electricity prices, modelled by a higher volatility in the chosen electricity price model, modular construction may represent a better option. Modular construction may include the possibility of having a few units ordered, when the electricity price reaches unfavourable values. With stable electricity prices the cost effective single large reactor appears to be a better choice.
- When there is the possibility of parallel construction of modular units, then:
  - the option value of the modular project greatly improves, and seems, with our model assumptions, more profitable than a single large reactor for different decision horizons.
  - For different electricity price volatility values, the modular project seems to be the better choice.
  - In many cases, once the first unit of the modular project is ordered it results in all units being ordered.

This chapter serves as a validation of our method against the results obtained in [35]. As future research, a more detailed analysis may include a sophisticated electricity price model, demand and capacity factors, and stochastic construction and operation costs. In the follow-up chapter the impact of features like the construction of twin reactors (parallel construction of modules), effect of learning and rare events on the real option values of various scenarios are also considered.

It's worth noting that the real option values computed here do not include electricity price prediction model uncertainty, which is typically called model risk in finance. It is presumed here that the electricity price follows GBM and the scenario paths used for calculations follow the same model. In reality however, the model presumed by a firm wouldn't exactly replicate the actual electricity price process distribution. Model risk can however be assessed by varying the price dynamics and studying the impact on the real option values. Furthermore, one can always compute sensitivities with respect to the different problem parameters. However, the purpose of the study here was to compare two projects in which case the electricity price prediction model uncertainty effects the valuation of both of them similarly. Note that we have shown results under specific model assumptions.

CHAPTER

5 Construction Strategies and Lifetime Uncertainties for Nuclear Projects: A Real Option Analysis

The contents of this chapter have appeared in [43]. This chapter focuses on the inherent value of flexibility that arises from different construction scenarios of nuclear power plants (NPPs). While the previous chapter focussed on the validation of a real option pricing model for projects involving modular construction in finite decision time horizon; this chapter exploits the real option model developed to analyze more realistic construction scenarios in the nuclear industry. Additionally we use the Stochastic Grid Bundling Method (SGBM), discussed in Chapter 3 to do the underlying computations for the real option problem (unlike the previous chapter where we used SGM for pricing the real options). In Section 5.1 the context of different construction strategies for nuclear power plants and its corresponding mathematical formulation are stated. Section 5.2 deals in detail with the real option analysis of projects under different construction strategies, while Section 5.4 describes the effect of a stochastic life time of operation for nuclear plants. Finally, Section 5.5 gives some concluding remarks.

5.1 Context

Consider a competitive electricity market where the price of electricity follows a stochastic process. A utility needs to make a choice between different projects to meet the same generation capacity expansion. The following construction scenarios are considered:

- The utility is planning a capacity expansion of 1200MWe and needs to make the choice between a single large reactor of 1200 MWe that benefits from the economy of scale or four modules of 300MWe each, that benefit from flexibility, learning, and site sharing costs.
- The utility has a choice between two twin units at the same site or four individual reactors at different sites. Twin units, in order to benefit from site

sharing costs, are constrained to be constructed one after the other. The latter project, although it does not benefit from site sharing costs, has the flexibility to order reactors at favourable times.

The following notations are used: the total number of series units is denoted by n. Unit number i is characterized by discounted averaged cost per KWh equal to θ_i, its construction time is denoted by C_i and the lifetime of its operation by L_i. Construction and lifetime are expressed in years. It is assumed that different modules are constructed in sequence, and the construction of unit $i+1$ can be decided from any time subsequent to the start of the construction of unit i.

A constant risk free interest rate, denoted by r is assumed here.

The value of different construction strategies can be affected by the uncertain life time of operation of nuclear power plants. In Section 5.4 the effect of stochastic life time of operation on the value of NPP is discussed.

5.2 Effects of construction strategies

In this section the real option analysis is used for investment decisions, which arise due to sequential construction of SMRs, for the different scenarios discussed earlier.

Economy of Scale vs Modularity

One of the measures identified to reduce capital costs of nuclear power by the NEA report [62] was increased plant size. The savings arising from the economy of scale when the unit size of power plants increases from the 300 MWe to 1300 MWe range have been studied by experts around the world since the early 1960s. The specific costs ($/kWe) of large nuclear power plants have been quoted within such a broad range that the derivation of scaling factors becomes difficult. In addition to savings arising from increased reactor unit size, cost reductions due to other factors such as construction of several units at the same site, effects of replication and series construction, and learning effects need to be incorporated in the analysis as well. In this test case two projects, one with a single large reactor which benefits from the economy of scale considerations, while the other project consists of a series of four SMRs which benefit from learning and site sharing costs, are considered. Moreover, the modular units benefit from the flexibility to order the reactors at optimal times.

For many years, bigger has been better in the utility industry. The economy of scale arguments have, for some time, and in many cases, reduced the real cost of power

production. The economy of scale can be expressed by the following scaling function, which relates the effect of changing the unit size to the cost of the unit,

$$\frac{TC_1}{TC_0} = \left(\frac{S_1}{S_0}\right)^{\gamma}, \tag{5.1}$$

where TC_0, TC_1 are the total cost for construction of two reactors with size S_0, S_1, respectively, γ is the scaling factor which is usually in the range of 0.4 to 0.7. It is assumed that the two reactors differ only in size, with *other details being equal*.

The effect of a learning curve and the associated cost reduction for nuclear technology has been studied in detail by Zimmerman (1982) [90]. Modular SMRs benefit from learning economies which result from the replicated supply of SMR component by suppliers and from the replicated construction and operation of SMR units by the utilities and their contractors (see Carelli et. al (2010) [19]). Boarin and Ricotti (2011) [11] separate four effects of modular construction:

1. Learning factor: The number of similar plants constructed world-wide will lead to increased experience in construction and therefore in decreased costs;
2. Modularity factor: Modularization assumes capital cost reduction for modular plants, based on the reasonable assumption that the lower the plant size, the higher is the degree of design modularization;
3. Multiple Units factor: The multiple units saving factor shows progressive cost reduction due to fixed cost sharing among multiple units on the same site;
4. Design factor: The design factor takes into account a cost reduction by assumed possible design simplifications for smaller reactors.

The following equation to model the combined impact of multiple units and learning effects on cost savings, as a function of number of units on site is used.

$$K_i = K_0((1-a) + ae^{-bi}), \tag{5.2}$$

where a is the cost-savings factor, which is asymptotically achieved by an increasing number of units, and b gives the rate of on-site cost savings. Factor a would depend on the number of units constructed world wide, the amount of R&D effort put in the technology, etc. Factor b depends on the contractor, the skills of the labour involved, etc. Figure 5.1 shows the price of subsequent units constructed at the same site for varying values of learning rate b. For increasing value of b, the subsequent reactor converges faster to the final cost efficiency gained by learning.

Consider now four modules, each with size 300 MWe, and compare this project with a single unit of size 1200 MWe. The scaling factor for the economy of scale is taken to be $\gamma = 0.65$ [50]. In order to benefit from local learning, a constraint that the construction of a next module can begin only after one year of the start

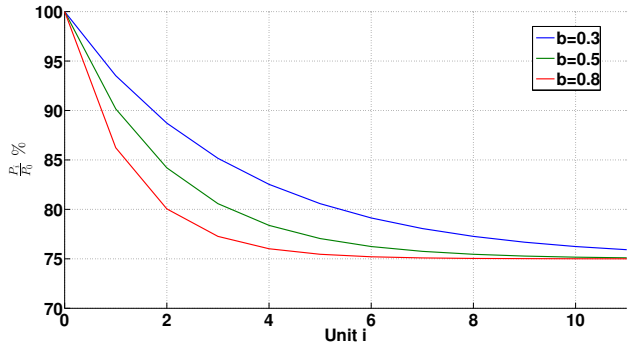

Figure 5.1: Cost savings factor for different local learning rates.

of the construction of the previous module, is added. The rate of local learning is taken as $b = 0.8$ and it is assumed that the cost saving for large numbers of modules would approach a value of 25%. These parameter values correspond to those suggested by Mycoff *et al.* [61], based on various studies in the literature, posit that the combined impact of multiple units and learning effects is a 22% reduction in specific capital costs for the SMR-based power plant. The discounted average costs of a large reactor are taken from Gollier *et al.* (2005) [35] and the corresponding values for modular units are computed based on the discussion above. With these parameters, the construction costs of the modules can be summarized in Table 5.1.

	Construction Time (months)	Discounted Average cost (cents/kWh)	Discounted Average cost ($/kWe)
Modular Units			
Unit 1	36	4.71	4950
Unit 2	24	4.06	4250
Unit 3	24	3.77	3950
Unit 4	24	3.63	3800
Large Reactor			
Unit 1	60	2.9	3000

Table 5.1: Construction time and discounted averaged cost used for the modular units with learning and a single large unit.

Figure 5.2 compares the option value obtained for different volatilities in the electricity price model. It is clear that with increasing uncertainties in the electricity prices the flexible modular project becomes more attractive. However, in a more certain environment a single large reactor seems to be profitable. Table 5.2 reports the option values for the two projects for different decision horizons for ordering the first unit. The gain due to the learning curve and the flexibility of construction, although it improves the option value for the modular units, does not seem to be sufficient to compensate for the economy of scale factor in the case of moderate uncertainty in electricity prices.

Figure 5.3 illustrates the sensitivity of the real option value of the modular project with respect to parameter a, which represents the asymptotic cost reduction that

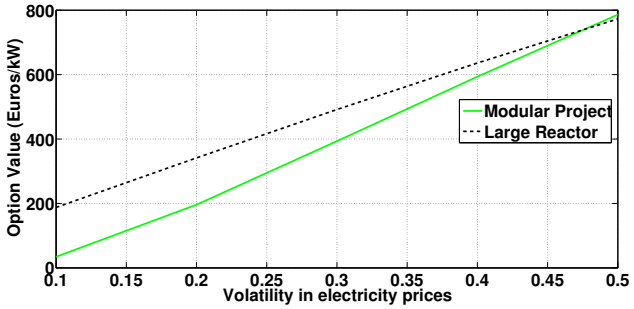

Figure 5.2: The real option value for the two projects for different uncertainties in the electricity prices when the decision horizon for ordering the first unit and the large reactor is equal to 7 years.

Decison Time years	Real Option Value (Euro/kW)
Modular Case	
7	196.17
10	229.72
13	250.15
Large Reactor	
7	341.48
10	364.04
13	376.98

Table 5.2: Option value (Euro/kW) for the modular case ($4 \times 300 MWe$) and for the large reactor (1200 MWe). The volatility for the electricity price is 20% and the initial price of electricity equals 3 cents/kWh.

will be achieved from learning, when the rate of learning (given by parameter b) is kept constant and equal to 0.8. It is seen that the option value of the project grows proportionally to the asymptotic cost saving due to learning.

Figure 5.4 illustrates the sensitivity of the real option value of the modular project for different rates of learning, given by parameter b, when a is kept constant and equal to 0.25. It is seen that for a faster rate of learning the option value of the project is higher as for large b values the cost savings are reflected rapidly in the subsequent units, while for a smaller b value the benefit is reflected only after few units have been constructed.

It can be seen from figures 5.3 and 5.4 that a higher *final cost saving factor* which is reflected by the value of parameter a increases the option value of the project more significantly than a higher on site learning rate as given by the parameter value b. Under our model assumptions a higher final cost saving could be achieved by long term cost reductions resulting from plant upgrades and/or increased R&D efforts to increase the real option value of the nuclear power plants.

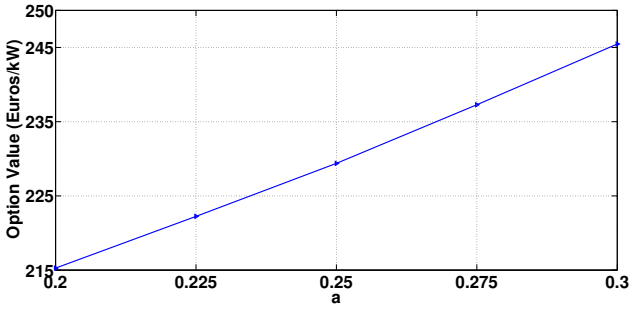

Figure 5.3: The real option value for the modular project when $b = 0.8$, the decision horizon is 10 years and the volatility in the electricity prices is 20%, for different values of parameter a.

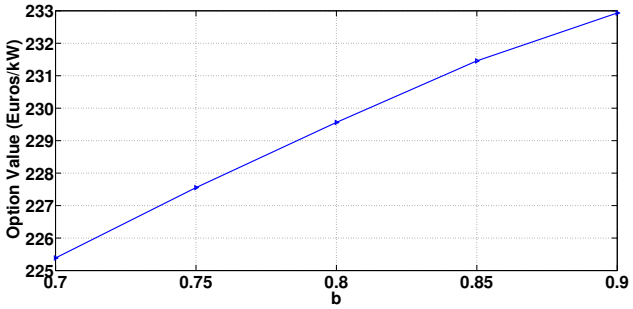

Figure 5.4: The real option value for the modular project when $a = 0.25$, the decision horizon is 10 years and the volatility in the electricity prices is 20%, for different values of parameter b.

Two Twin Units vs Four Single Units

In this section the case describing the sharing of facilities by constructing multiple units at a single site is considered. The parameter values are taken from a real case observed at EdF as described in the NEA report [62]. The averaged costs of a unit reduce with an increasing number of units per site. The case where two pairs of units per site are constructed is compared with a case where four individual units are constructed. All units considered are of the same size, so the economy of scale doesn't play any role in this case.

The aim of the test case is to compare a project with two twin units on a single site with four individual units constructed at different sites. It is assumed that the reactors involved are of the same size (1200 MWe each), and hence the only cost difference comes from the sharing of costs if constructed at a single site. The first reactor, in both cases, is considered to be first-of-a-kind (FOAK), and it is further assumed that the cost of generating electricity for this reactor is 3.5 cents/kWh. The costs of the other units are summarized in Table 5.3. In order to achieve a cost benefit for the reactors constructed at the same site, the reactor units are con-

strained to be constructed in a phased manner. Therefore, the two twin reactors are constructed immediately one after the other, with the construction of the first unit starting when the electricity price crosses the corresponding critical price. The benefits of cheaper subsequent units come at the loss of flexibility to order the subsequent units at optimal electricity prices. The reasons for phased construction include considerable efficiencies and associated savings to be gained from the phased construction and rolling the various craftsmen teams from one unit to the next. In addition, by repetition of construction, there is the *craft labour learning effect* that reduces the time to perform a given task and correspondingly reduces labour cost and schedule. The decision horizon for ordering the first reactor is taken to be 7 years, and it is assumed that once a unit becomes operational it operates at its maximum capacity factor.

On the other hand, when the four units are constructed at separate sites, they do not have the cost benefits of sharing the site specific costs, neither of the productivity effects. However, when constructed individually they benefit from the flexibility to order each unit at its corresponding optimal time.

	Construction Time (months)	Discounted Average cost (cents/KWh)
Two twin units		
Unit 1	60	3.5
Unit 2	48	1.67
Unit 3	48	1.81
Unit 4	48	1.60
Four independent units		
Unit 1	60	3.5
Unit 2 to 4	48	2.25

Table 5.3: Construction times and discounted average costs used for the two twin units at same site and for four units at different sites.

The following assumptions are made when determining the overnight costs:

- Unit 1 bears all of the extra first-of-a-kind (FOAK) costs.
- The cost of engineering specific to each site is assumed to be identical for each site.
- The cost of facilities specific to each site is assumed to be identical for each site.
- The standard cost (excluding the extra FOAK cost) of a unit includes the specific engineering and specific facilities for each unit

If θ_0 is the standard cost (excluding the extra FOAK cost) of a sole unit on a site, the:

- Cost of the first unit is θ = (1+x)θ_0;

- Cost of the following units is θ_0 (for 1 unit per site);
- Cost of the 2nd unit at a site with one pair is y θ_0;
- Cost of the 3rd unit at a site with two pairs is z θ_0;
- Cost of the 4th unit at a site with two pairs is y θ_0 (it is assumed that the cost of the 2nd unit of a pair is independent of the rank of the pair on the site).

The productivity effect depends on the rate of commitment of the units and on the owner's procurement policy. The optimum commitment rate is the one which provides good operational feedback from one construction project to the other, while serving to maintain the apprenticeship effect in the manufacturers' facilities and on the sites. The most favourable procurement policy to obtain the best prices from the suppliers consists in ordering the equipment of all the units under the same contracts [62].

A productivity effect is considered to only occur from the 3rd unit on of a series. If n is the rank of the unit in the series, and θ_n is the cost which results from taking into account the individual unit, it follows that:

$$\theta'_n = \frac{\theta_n}{(1+k)^{n-2}} \quad n \geq 2,$$

where θ'_n represents the cost of a module if there is a productivity gain involved. Using the above formulation, for the case of EdF (x = 55%, y = 74%, z = 82 %, k = 2%) [62] in the case of the two pairs of units per site the relative costs are illustrated in Figure 5.5

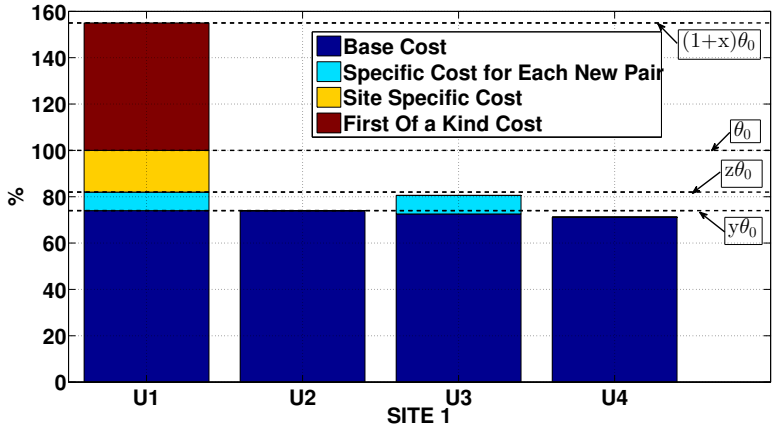

Figure 5.5: Relative cost of the four units when constructed as two twin units on a single site

In the case of two twin units the cash flow due to the sales of electricity once the units become operational is modified from Equation (4.1) into

$$R_1(P_t = x) = \mathbb{E}\left[\left(\int_{t+C_1}^{t+C_1+L_1} e^{-ru}P_u du + \int_{t+(C_1+C_2)}^{t+(C_1+C_2)+L_2} e^{-ru}P_u du + \cdots + \int_{t+(C_1+C_2+C_3+C_4)}^{t+(C_1+C_2+C_3+C_4)+L_4} e^{-ru}P_u du\right) | P_t = x\right]. \quad (5.3)$$

The above Equation (5.3) can be read as the revenue from unit 1, when ordered at time t, starts coming in once its construction is finalized at $t+C_1$, and continues till the end of its lifetime, i.e. until time $t+C_1+L_1$. The construction of unit 2 starts at $t+C_1$, and its revenues start flowing in from $t+C_1+C_2$, till the end of its lifetime at $t+C_1+C_2+L_2$. Equation (4.2) can be modified accordingly. It can be seen that only the first reactor can be ordered at an optimal time, and the construction of the ith reactor is forced immediately after the completion of the $(i-1)$th reactor to achieve the cost reductions shown in Figure 5.5.

For the four independent reactors the revenue is given by Equation (4.1). As significant cost benefits can be achieved after the construction of the FOAK unit, the decision for the construction of subsequent units is made only when the construction of the first unit is done. Also the usual constraint that subsequent modules can be ordered only after the decision on starting the construction of all previous modules has been taken into account. The decision horizon for unit i is $T_i = T_{i-1} +$ *Construction Time of Unit (i-1)*, where T_{i-1} is the decision horizon for the $(i-1)$ th unit.

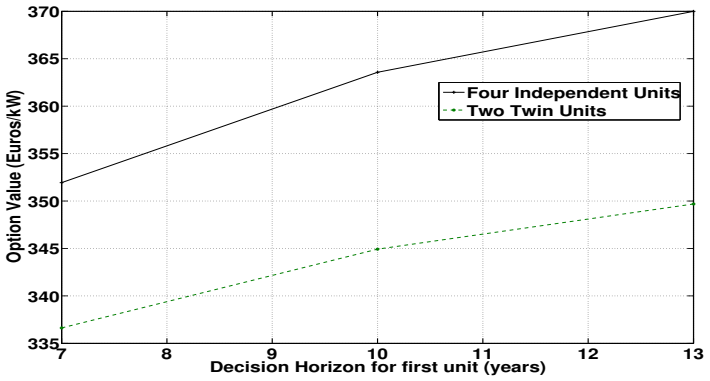

Figure 5.6: Option values (Euro/kW) for the twin units, and the four independent units for different decision horizons. The option values are computed when the initial electricity price is 3 cents/kWh and the volatility in the electricity prices is 20%.

Figure 5.6 compares the option values of the two projects for different decision horizons for starting the construction of the first unit. When a firm has more time

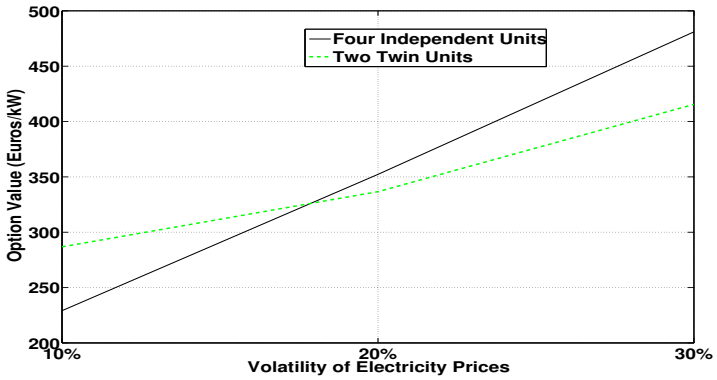

Figure 5.7: Option values (Euro/kW) for the twin units, and the four independent units for different levels of uncertainties in the electricity prices. The option values are computed when the initial electricity price is 3 cents/kWh and decision horizon for the first unit is 7 years.

to decide when to construct the reactor, the option value of the project increases. From these results it seems that for the parameters chosen the project involving four independent units appears to be a better choice, under our model assumptions.

Figure 5.7 compares the two projects for different uncertainties in the electricity prices. The project involving four independent units seems to be less profitable in a more certain environment of electricity prices. The benefit of flexibility in a project becomes more apparent with increasing uncertainties in electricity prices.

It can be seen that although building two twin reactors at a single site significantly reduces the costs of producing electricity for the units, the project loses on the value of flexibility. The four units constructed at different sites, although produce electricity at higher prices, can benefit from the opportunity to construct at more optimal market electricity prices. It should also be noted that the construction of the units in this manner may also benefit since in an event of natural disaster not all units would shut down (see Takashima, Yagi (2011) [78], for more details). Also it is clear that when the uncertainty in electricity price increases, it appears to be advisable to focus on flexibility, by constructing independent reactors, rather than the increased cost efficiencies, by constructing twin units at the same site. On the other hand, when the electricity price is less uncertain not much is gained by the flexibility in the time to order and it appears more profitable to choose cheaper twin units in our model.

5.3 Parameter values for modular construction

Consider a case where an investor needs to decide between two projects, one involving a single large reactor of 1200 MWe and the other consisting of four modules of 300 MWe each. The construction time and costs for the two projects, given in Table 5.4, are taken from the reference case by Gollier [35]. The discount rate is taken as 8% per annum and the predicted growth rate of electricity price is 0% here. The cost of electricity production for the first unit is relatively expensive when compared to the series units, as a large part of the fixed costs for the modular assembly, like the land rights, access by road and railway, site licensing cost, connection to the electricity grid are carried by the first unit.

	Construction Time (months)	Discounted Average cost (cents/KWh)
Large Reactor	60	2.9
Modular Case		
Module 1	36	3.8
Module 2 to 4	24	2.5

Table 5.4: Construction times and discounted averaged costs used for the large reactor and the modular case.

It is assumed that a new unit can be ordered if all previous units have been *ordered* (not constructed). It's common practice to have parallel construction of different units in order to achieve cost savings, as it allows rotation of specialized labour between different units [62].

5.4 Effects of uncertain life time of operation

Uncertain life times of operation should be taken into account when computing the value of investment in an NPP. A detailed analysis would not just take the uncertain life time of operation into account but also uncertain capacity factors during the operation of the reactor. Du and Parsons (2010) [28], perform a detailed analysis on the capacity factor risk in the nuclear power plants. Rothwell(2005) [73] employs a stochastic process for varying capacity factors in his analysis. Here it is assumed (like Gollier [35]) that the nuclear reactors operate at a mean capacity factor of 90% throughout their lifetime, which is a reasonable assumption [86] for modern reactors.

In the analysis it is assumed that the uncertain life time of operation of NPPs can be due to premature permanent shut down on one hand or due to extension of operating licenses and lifetime on the other hand.

Effects of Premature Permanent Shut-down

The term *premature permanent shut down* is used for the case when an operating reactor is permanently shut down before completing its licensed operating life time. Historically, premature permanent shut down of reactors have been observed for direct reasons- like accidents or serious incidents in a reactor (e.g. Three Mile Island 2- 1979, Chernobyl 4- 1986, Fukushima Daiichi 1,2,3,4 - 2011), or it could be indirect by - for example shutting down of reactors due to increased safety measures, economic reasons, changing government policies etc. (e.g. Shoreham, in US 1989).

The arrival time of such an event (elsewhere called *rare events* or *catastrophic events*) has been modelled by a Poisson process, e.g. Clark (1997) [22] for a real options application with a single source for rare events, Schwartz (2003) [76] uses Poisson arrival times to model catastrophic events when investing in R&D. Similarly, here, the arrival time for the cause of premature permanent shut down is modelled as a Poisson process whose arrival frequency, λ, is the expected number of such events a year.

In order to compute the frequency of premature permanent shut down the data available from the IAEA report (2005) and a WNA report [86] is used. In total there have been 133 reactors that have been permanently shut down after they started operating. Out of these, 11 reactors were shut down due to accidents or serious incidents, and 25 have been shut down due to political decisions or due to regulatory impediments without a clear or significant economic or technical justification. The remaining 97 reactors were shut down because they completed their designated lifetime and costs associated with a lifetime extension did not make economic sense for these reactors. The total cumulative life time of operation for the reactors in the world is approximately 14500 years. Therefore, the number of premature permanent shut downs per reactor year is $(25+11)/14500$, which is 0.0025 reactors per year. Thus, the rate of arrival of the cause for premature permanent shut down, is modelled as $\lambda = 0.0025$ events every year, from a statistical point of view.

The inclusion of catastrophic events results in an effective discount rate from r to $r+\lambda$, (see Schwartz (2002) [76] for more details) once the plant gets operational . Figure 5.8 compares the option values for the modular project and the single large reactor, with parameter values as given in Table 5.4. It's clear that the value of the investment option reduces with increasing probability of catastrophic events, however, the modular project seems to be more profitable in the realistic domain of λ values.

Effects of Life Time Extension

Most nuclear power plants originally had a nominal design lifetime of 25 to 40 years, but engineering assessments of many plants have established a longer oper-

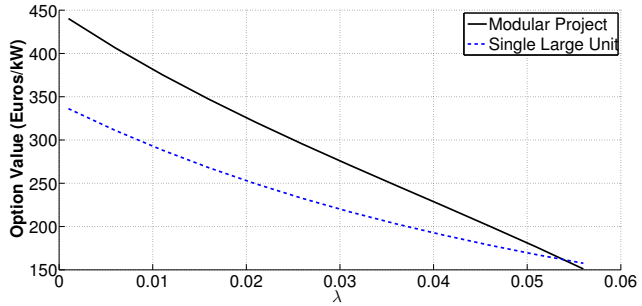

Figure 5.8: The real option price for the large reactor, and the modular project for different values of λ when the volatility of the electricity price is 20% and the decision horizon for the large reactor and the first module equals 7 years.

ation time. In the USA over 60 reactors have been granted licence renewals which extend their operating lives from the original 40 to 60 years, and operators of most others are expected to apply for similar extensions. Such licence extensions at about the 30-year mark justify a significant capital expenditure for replacement of worn equipment and outdated control systems. In 2010 the German government approved a lifetime extension for the country's 17 nuclear power reactors. However, after the Fukushima accident in March 2011, Germany planned a complete phase-out by the year 2022, reverting the previous decision. It is clear therefore that the lifetime of a nuclear power plant can be uncertain.

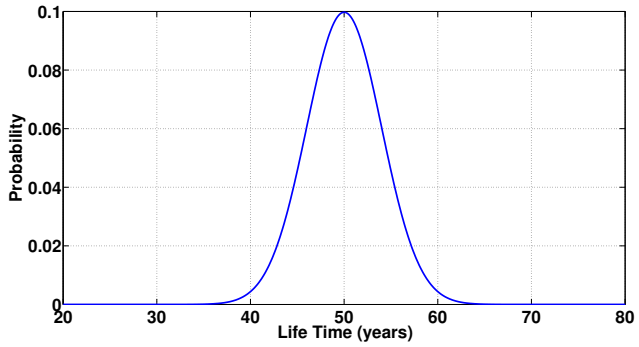

Figure 5.9: The distribution for the lifetime of operation of a nuclear reactor, which was originally licensed for 40 years of operation.

Gen III and Gen IV reactors are mostly designed for a life time of 60 years [50]. However many generation II reactors are being life-extended to 50 or 60 years, and a second life-extension to 80 years may also be economic in many cases [63]. In order to address the uncertain lifetime of operation due to the possibility of lifetime extension, a normal distribution with a mean reactor life of $\mu_l = 50$ years and a variance of $\sigma_l = 4$ is used, which fits well to the discussion above, as can be seen in

Figure 5.9. It should be noted that such a distribution allows for negative life time, however the probability for such lifetime is almost negligible.

In the case where the electricity price follows GBM and the lifetime has a normal distribution, as described above, equations (4.2) and (4.1) can be written as

$$R_i(P_t = x) = e^{-(r-\alpha)C_i} \frac{1 - e^{-((r-\alpha)\mu_l - \frac{\sigma_l^2 (r-\alpha)^2}{2})}}{r-\alpha} x, \tag{5.4}$$

and

$$K_i = e^{-rC_i} \frac{1 - e^{-(r\mu_l - \frac{\sigma_l^2 r^2}{2})}}{r} \theta_i. \tag{5.5}$$

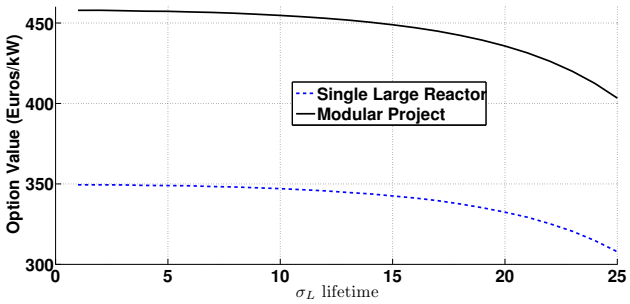

Figure 5.10: Option value vs uncertainty in lifetime of operation

Figure 5.10 compares the real option value for the modular project and the single large reactor with parameter values taken from Table 5.4. The option values are plotted for various σ_l with a mean reactor life of 50 years. It can be seen that with increasing uncertainty in the lifetime the option value reduces in value, although it follows the same trend for both cases considered.

5.5 Conclusion

This chapter presents a real option valuation of different construction strategies of NPPs for finite decision horizon. A few scenarios a utility might be interested in before making a choice of nuclear reactor, are analyzed. The conclusions drawn from the test cases *under the model assumptions* can be summarized as follows:

1. In a finite decision horizon, sequential modular units can be ordered at more competitive electricity prices, compared to a construction of units in isolation.

2. The model shows that the real option value of nuclear power plants increases with implementation of long term cost reductions. Such cost reductions might be achieved e.g. by plant upgrades and/or increased R&D efforts.

3. When twin units are constructed at the same site, significant cost reductions can be achieved. On the other hand, in order to achieve these cost savings, the utility loses the flexibility to order units at optimal market conditions. When the electricity price uncertainty is low it appears that cost savings by a construction of twin modules at the same site would be favourable under our model assumptions, while when electricity prices are more volatile, the flexibility of choice dominates.

4. Uncertain lifetime of operations reduces the option value of both the modular and single large reactor test cases. However, with increasing frequency of premature permanent shutdown, the modular construction works out more profitable than individual large reactors. Uncertainty in life time extension affects the option value of a single large unit and modular units almost similarly in our model.

5. Specific cost of SMRs can be much higher than of a single large unit, because of the economy of scale argument. Some cost reduction is achieved by the learning effect with each new module. However, it appears that cost savings due to learning are not sufficient to make modular SMRs competitive with large units.

CHAPTER

6 Decision-support tool for assessing future nuclear reactor generation portfolios.

The global electricity demand is expected to double to over 30,000 TWh annually by the year 2030 and meeting this demand without substantially exacerbating the risks of climate change requires a solution comprised of a variety of technologies on both the supply and demand side of the energy system (Pacala and Socolow (2004) [65], Holdren (2006) [38] and European Commission (2007)[23]). Nuclear power can play a key role in meeting the projected large absolute increase in energy demand while mitigating the risks of serious climate disruption. The fact that several countries seem keen on building nuclear power stations suggests that their relative costs compared to low-carbon alternatives seem attractive to at least potential investors (Kessides, 2010 [49]). However, there are some concerns related to uncertainties underlying the various costs elements of nuclear power that are reflected in the wide range of cost estimates, cost overruns and schedule delays, for example of Finland's Olkiluoto and France's Flamanville nuclear power plants.

There have been numerous studies on the economics of nuclear power in recent years which use levelized cost[1] of electricity to compare the economics of different generation technologies. The levelized cost methodology used in these studies however does not address the role of risks and uncertainties involved. Methodologies that take into account the large and diverse set of risks characterizing investment in nuclear power maybe a useful alternative. This chapter concentrates on the effect of risks and uncertainties on investment decisions related to the nuclear industry and the use of diversification to mitigate some of these risks. Following Roques *et al.* (2008) [72] and Fortin *et al.* (2007) [30] a two-step approach is used, where first real options optimal investment decisions are taken at the plant level, and then mean-variance portfolio (MVP here after) theory is used to minimize the uncertainties of returns for a portfolio of nuclear reactors.

The seminal literature using MVP techniques in the power sector concentrated on fuel price risk, and focussed on minimizing generation cost, which, under ideal

[1] The levelized cost of a project is equivalent to the constant euro price of electricity that would be required over the life of the plant to cover all operating expenses, interest and repayment obligations on project debt, and taxes plus an acceptable return to equity investors over the economic life of the project.

regulations of a vertically integrated franchise monopoly, should maximise social welfare. Awerbuch and Berger (2003) [4] use MVP to identify the optimal European energy technology mix, considering not only fuel price risk but also Operation and Maintenance (O&M), as well as construction period risks, while Jansen *et al.* (2006) [44] use MVP to explore different scenarios of the electricity system development in the Netherlands. Roques *et al.* (2008) [72] applied the portfolio theory from a private investor perspective to identify optimal portfolios for electricity generators in the UK electricity market, concentrating on profit risk rather than production costs risk. Fortin *et al.* (2007) [30] suggest the use of Conditional Value-at-Risk (CVaR) for portfolio optimization rather than mean-variance portfolio and provide a detailed review of the literature in this area.

Real options analysis (ROA) has been applied to the energy sector planning for years, since the special features of the electricity sector, such as uncertainty, irreversibility and flexibility to postpone investments, make standard investment rules solely relying on the net present value (NPV) not advisable as they ignore the options involved in a sequence of decisions. Using real options it's possible to value the option to delay, expand or abandon a project with uncertainties, when such decisions are made following an optimal policy.

This chapter concentrates on investment in nuclear power plants in a liberalized electricity market, where the energy utility diversifies into different nuclear reactor types as a strategy for reducing exposure to construction costs, fuel and electricity price risks. Mean-variance portfolio (MVP) theory is used to identify the portfolios that maximize the returns for given risk levels. The return distribution of individual nuclear generation types depends on the uncertainties in the costs and revenues of the plant. It is, however, also affected by decisions to continue or abandon a project, that may be taken based on evolution of construction costs and electricity prices. For example, if the construction costs become too high in the future, the management may decide to abandon a project. Using real options the return distribution for each plant assuming the management makes optimal decisions in the future is computed. The return distribution for each plant is then used to compute the mean-variance portfolio.

This chapter uses the Stochastic Grid Bundling Method (SGBM) (see Chapter 3), for computing the return distribution for individual reactors. The simulation also computes the optimal policy to continue or abandon the project in order to maximize its expected cashflows.

The rest of the chapter is structured as follows: Section 6.1 will be concerned with defining the portfolio optimization problem. Section 6.2 gives detailed account of the real options layer used for making optimal decisions at the individual plant level. In Section 6.3 the present model is validated against the results reported in Pindyck (1993) [67]. Section 6.4 illustrates the two steps involved when determining the optimal reactor order fractions through various numerical examples . Under our model assumptions, the sensitivity of reactor order fractions to a different choices of parameter values and constraints on the portfolio are also studied in this

section. The final section will conclude the findings and interpret the general implications.

6.1 Mean Variance Portfolio

While selecting the generating technology, policy makers need to consider not only the cost of the generating technology but also uncertainties in the costs involved. Furthermore, in liberalized energy markets uncertainties are not only limited to the costs of the generating technology but also affect the revenues stream, as utilities are no longer able to pass on their prudently incurred investments costs to consumers. In order to systematically deal with uncertainties in the costs and revenues, like Awerbuch and Berger (2003) [4], Roques *et al.* (2008) [72]the MVP theory[2] is employed to find an optimal mix of generating technologies, that results in the highest expected return for a given level of uncertainty (or standard deviation) of the returns[3].

To compute the optimal reactor order fraction using MVP, the expected return distribution for individual reactors is required. One way of obtaining the return distribution is by simulating several samples of costs (like the fuel prices) and revenues (electricity prices) and then computing the return for each sample. This approach however does *not* address the effect of possible future decisions related to operation of the power plant (for example, abandoning the plant if the expected costs exceed expected revenues at a later date) on the return distribution. In order to include the effect of optimal decisions in the return distribution, first an optimal investment policy for each reactor type is computed. This policy is then applied to simulated paths to determine whether for a particular path there should be an early abandonment. Based on these decisions the costs and revenues for each sample path are computed, which then gives the optimal return distribution. The details for computing an optimal investment policy and the associated return distribution for individual plants are given in Section 6.2.

Suppose an investor has a certain wealth to invest in a set of J reactors. Let the return from operation of reactor i be denoted by random variable R_i, and let w_i represent the proportion of the total investment to allocate in the i-th reactor. The expected return of this portfolio is given by:

$$\mathbb{E}[R_p] = w_1\mathbb{E}[R_1] + \ldots + w_J\mathbb{E}[R_J]. \tag{6.1}$$

[2] MVP is one of the possible ways for portfolio optimization, based on how the risk is expressed, which in the case of MVP is the standard deviation of the returns. Others like Szolgova *et al.* (2011) [77], Fuss *et al.* (2012) [31] use Conditional Value at Risk (CVar) for portfolio optimization.

[3] See Awerbuch and Berger (2003) and Jansen et al. (2006) for a discussion of the assumptions and limitations affecting the application of MVP theory to power generation assets.

The portfolio variance, in turn, is calculated by

$$Var(R_p) = \mathbb{E}\left[\left(\sum_{i=1}^{J} w_i R_i - \mathbb{E}\left(\sum_{i=1}^{J} w_i R_i\right)\right)^2\right]. \tag{6.2}$$

So,

$$Var(R_p) = \sum_{i=1}^{J}\sum_{j=1}^{J} \mathbb{E}\left[(R_i - \mathbb{E}[R_i])(R_j - \mathbb{E}[R_j])\right] w_i w_j. \tag{6.3}$$

Representing each entry i, j of the covariance matrix Q by

$$q_{ij} = \mathbb{E}\left[(R_i - \mathbb{E}[R_i])(R_j - \mathbb{E}[R_j])\right], \tag{6.4}$$

one has

$$Var(R_p) = \mathbf{w}^\top Q \mathbf{w},$$

where $\mathbf{w} = (w_1, \ldots, w_J)^\top$.

As w_i represents the weight of reactor i, the weights are required to satisfy an additional constraint:

$$\sum_{i=1}^{I} w_i = 1.$$

As a portfolio of nuclear reactors is dealt with, additional conditions on the weights, like that they cannot be negative, need to be applied. Additionally, weights of individual reactors might be constrained by an upper and lower bound, for example, if the utility decides that the new portfolio should not excessively deviate from the existing one. In general, it can be stated that:

$$L_i \leq w_i \leq U_i,\ i = 1, \ldots, J,$$

for given lower L_i and upper U_i bounds on the weights.

MVP theory does not prescribe a single optimal portfolio combination, but rather a *range* of efficient choices for each level of return, which form a *Pareto efficient frontier* composed of non dominated points. This means that a rational investor should use an external criterion to choose a portfolio out of the set at hand. Investors will choose a risk-return combination based on their preferences and risk aversion. By solving the mean-variance optimization problem a portfolio for given risk tolerance, λ, of the investor, of minimum variance is identified amongst all that provide a return equal to $R_{\min}$, or, in other words, minimize the risk for a given level of return. The formulation can be written as:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{\lambda}\mathbf{w}^\top Q\mathbf{w}, \\ \text{subject to:} \quad & \mathbb{E}[R_p] = R_{\min}, \\ & \sum_{i=1}^{J} w_i = 1, \\ & L_i \le w_i \le U_i,\ i = 1,\dots,J. \end{aligned} \tag{6.5}$$

Equation (6.5) is a convex quadratic programming problem for which the first-order necessary conditions are sufficient for optimality. The classical Markowitz mean-variance model can be seen as a way of solving the bi-objective problem, which consists of simultaneously minimizing the portfolio risk (variance) and maximizing the portfolio return (profit), i.e.

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{\lambda}\mathbf{w}^\top Q\mathbf{w}, \\ \max_{\mathbf{w}} \quad & \mathbb{E}[R_p], \\ \text{subject to:} \quad & \sum_{i=1}^{J} w_i = 1, \\ & L_i \le w_i \le U_i, i = 1,\dots,J. \end{aligned} \tag{6.6}$$

The solution of Equation (6.6) is non-dominated, efficient or Pareto optimal for Equation (6.5). Efficient portfolios are thus the ones which have the minimum variance among all that provide a certain expected return or, in other words, those that have maximal expected return among all upto a certain variance.

6.2 Plant level optimization using real options

The real option valuation of nuclear power plants should take into account the major uncertainties that affect the decision making process associated with them. Of the several risks involved in the life cycle of nuclear power plants (see Kessides (2010) [49] for a comprehensive review), the following have been identified as significant from the perspective of economic risks and are taken into account in our model.

- *The construction or capital costs, and the speed to build:* The length of the pre-construction period and the time it takes to construct the plant are highly uncertain as there are several factors that make forecasting nuclear plant construction costs difficult. As pointed out by Kessides (2010) [49] one of the reasons for this is that new nuclear plants require a significant amount of on-site engineering, which accounts for a major portion of the total construction

cost (Thomas, 2005 [80]). It is generally difficult to manage and control the costs of large projects involving complex on-site engineering. While major equipment items (turbine generators, the steam generators, and the reactor vessel) can be purchased on turnkey terms, it would difficult for the entire nuclear plant to be sold on turnkey terms precisely because of the lack of confidence on the part of vendors that they can control all aspects of the total construction costs. Additionally, governmental licensing and certification procedures can add up significantly to construction costs and delays.

- *The O&M and fuel costs:* The O&M component includes expenses related to health and environmental protection and accumulation of funds for spent-fuel management and for eventual plant decommissioning. It also includes the cost for insurance coverage against accidents. Thus, several potential externalities are internalized in O&M costs.
- *The price of electricity:* Electricity prices are highly uncertain and vary significantly not just between different seasons but also during a single day. Thus, the revenues generated by a power plant are uncertain and an important parameter for making optimal decisions.

Modelling uncertain construction costs

Construction or capital costs constitute almost 60% of the total costs associated with nuclear power plants and are the major source of uncertainty when it comes to a comprehensive cost-benefit analysis of nuclear power. An economic assessment that reflects on the uncertainty in construction costs by employing probabilistic scenario analysis can help making economic decisions related to NPPs. To capture the uncertainties associated with the construction costs and their effect on the decision making process, the model proposed by Pindyck (1993) [67] for irreversible investment decisions when projects take time to complete and are subject to uncertainties over the cost of completion, is followed.

Expenditure of nuclear power plants are sunk costs that cannot be recovered should the investment turn out, *ex post*, to have been an unfavourable one, i.e. the firm cannot disinvest and recover the money spent. Cost uncertainties have implications for irreversible investment decisions. The uncertainties in construction costs of nuclear power plants can be classified into two different types. The first, as Pindyck (1993) [67] states, is *technical uncertainty*, that relates to the technical difficulties associated with the completion of the nuclear power plant, i.e. if the costs of raw materials, labour etc. are fixed then the uncertainty reflects how much time, effort and material will ultimately be required. Technical uncertainties involved in the construction of the plant can be resolved only by undertaking the project which unfolds the actual costs and construction time as the project proceeds.

The second type of uncertainty that affects the construction costs is external or independent of what the firm does and is called *input cost uncertainty*. Input cost uncertainty arises when the prices of labour, land, materials needed to build the

plant fluctuate unpredictably, or when there are unpredictable changes in government regulations (for example a change in the required quantities of construction inputs or certification time). As prices and government regulations change irrespective of whether or not the construction of a plant has already begun, input costs uncertainties affect the expected plant costs.

Consider the expected cost of completion of a nuclear power plant to be a random variable K, then, following Pindyck (1993) [67], the stochastic differential equation (SDE) governing the dynamics of K_t can be written as:

$$dK_t = -I\,dt + \beta(IK_t)^{\frac{1}{2}}\,dW_\beta + \gamma K_t\,dW_\gamma, \tag{6.7}$$

where I is the rate of investment. When the construction of a nuclear power plant has begun the expected change in K_t over an interval dt is $-I\,dt$, but the realized change can be greater or less than this due to the random fluctuations in the cost to completion of the project. The term $\beta(IK_t)^{\frac{1}{2}}\,dW_\beta$ constitutes a part of the fluctuation in the project cost due to the technical uncertainty, where the noise is introduced by the Wiener process W_β and the amplitude of the noise depends on the remaining expected costs of the project and the rate of investment I, and β. When the firm is not investing, i.e., I is zero the project cost is not influenced by technical uncertainties. The term $\gamma K_t\,dW_\gamma$ constitutes the part of the fluctuation in the project costs due to input cost uncertainty. As discussed before, this uncertainty affects the cost of the plant irrespective of I, i.e. whether the firm is investing or not. Higher values of parameters β and γ, result in greater uncertainties in realized construction costs of the power plant. The time for completion of the power plant is a stochastic variable $\tilde{T}$ and is the time when K_t falls to zero. W_β and W_γ are uncorrelated Wiener processes, with W_β being also uncorrelated to the economy and the stock market, while W_γ may be correlated with the market.

It is assumed that the firm invests in the project at a constant rate (i.e. I is constant), a fact also observed in practice as shown in Table 6.1, where the fraction of the overnight costs[4] for the construction of a power plant in different countries incurred each year is almost equal.

Modelling uncertain O&M , fuel and electricity prices

During a nuclear power operation period, the *generating costs* consist of operational and maintenance cost, back-end and front-end fuel cycle costs. Following Rothwell (2006) [73] and Zhu (2012) [89] the uncertain generation costs are modelled by Geometric Brownian Motion (GBM). The dynamics of the generation costs are described by the following SDE:

$$dC_t = \mu_c^* C_t\,dt + \sigma_c C_t\,dW_C, \tag{6.8}$$

[4] Overnight cost is the cost of a construction project if no interest was incurred during construction, as if the project was completed "overnight."

Year	CAN	USA	FIN	NLD	CHE	JPN	ROU
-8							16.5
-7							12.5
-6		10			3	5	12.5
-5	8	20	10	20	19	15	12.5
-4	22	20	22	20	19.5	20	12.5
-3	29	20	28	20	19.5	20	16.5
-2	21	20	20	20	19.5	18.5	12.5
-1	12.5	10	20	20	19.5	21.5	4.5
1	7.5						

Table 6.1: Expense schedule for nuclear power plant construction from country to country expressed as percentage of total overnight construction cost per year. Source: OECD (2005), CAN: Canada, FIN: Finland, NLD: The Netherlands, CHE: Switzerland, ROU: Roumania. Year stands for number of years before the plant becomes operational. [64].

where C_t is the instantaneous cost of generation in € per kWh, μ_c^* is a risk adjusted drift[5] of the generation costs and σ_c is the volatility of the generation costs. W_C is a Wiener process which may be correlated to the market.

Modelling *electricity spot prices* is difficult primarily due to factors like:

- Lack of effective storage, which implies that electricity needs to be continuously generated and consumed.
- The consumption of electricity is often localized due to constraints of the grid connectivity.
- The prices show other features like daily, weekly and seasonal effects, that vary from place to place.

As decisions for setting up power plants look at *long term evolution* of electricity prices, like Gollier *et al.* (2005) [35], the GBM is used as the electricity price process. However, it should be noted that within the present modelling approach other price processes can easily be included. The dynamics of electricity prices in our model are now described by

$$dP_t = \mu_p^* P_t dt + \sigma_p P_t dW_P, \tag{6.9}$$

where P_t is the instantaneous cost of electricity in € per kWh, μ_p^* is the risk adjusted drift of electricity price process and σ_p gives the volatility of electricity prices.

[5] if μ_c is the true drift of generation cost then the risk adjusted drift is $\mu_c^* = \mu_c - \eta$, assuming that the Intertemporal Capital Asset Pricing model of Merton (1973) [59] holds, the risk premium η is equal to the β^* of the successful project times the risk premium of market portfolio: $\eta = \beta^*(r_m - r_f)$.

Value of the power plant after it becomes operational

When the construction of a power plant is finished, i.e. $K_t = 0$, the value of the project depends only on the net cashflow to be generated from the operation of the power plant. Let $h_t(P_t, C_t)$ be the value of the power plant, once it becomes operational, at time t when the instantaneous cost of electricity is P_t € per kWh and the combined O&M and fuel cycle costs are C_t € per kWh. Let $\tilde{t}_S$ denote the time when the plant starts its operation, i.e. $\tilde{t}_S$ is the first instance when $K_t = 0$. Then, the time when it will be decommissioned, $\tilde{t}_f$, is equal to,

$$\tilde{t}_f = L + \tilde{t}_S,$$

where L is the designed lifetime of operation for the power plant and $\tilde{t}_S \leq t \leq \tilde{t}_f$. The expected discounted stream of future differences in cash flows at time t, under the risk neutral measure $\mathbb{P}$, from the remaining operation of the power plant, assuming the plant is decommissioned only after completing its designed lifetime, is then a function of its current state, P_t, C_t, and is equal to:

$$\begin{aligned} h_t(P_t, C_t) &= \mathbb{E}\left[\int_t^{\max(\tilde{t}_f, t)} e^{-r_f \tau}\,(P_\tau - C_\tau)\,d\tau | P_t, C_t\right] \\ &= e^{-(r_f - \mu_p^*)t} P_t \frac{1 - e^{-(r_f - \mu_p^*)(t_f - t)^+}}{r_f - \mu_p^*} \\ &\quad - e^{-(r_f - \mu_c^*)t} C_t \frac{1 - e^{-(r_f - \mu_c^*)(t_f - t)^+}}{r_f - \mu_c^*}, \end{aligned} \tag{6.10}$$

where r_f is the risk free discount rate and $(t_f - t)^+$ is used to denote $\max(t_f - t, 0)$.

Real option value of the power plant

The option value of the power plant before it becomes operational depends on the electricity price, P_t, combined fuel cycle and O&M costs, C_t, that would be incurred if the plant becomes operational and on the expected cost of completion, K_t of the power plant. The option value, $V_t(P_t, C_t, K_t)$, of the plant can be computed using Ito's lemma to obtain the differential equation for dV :

$$\begin{aligned} dV &= \frac{\partial V}{\partial t}dt + \frac{\partial V}{\partial P}dP + \frac{\partial V}{\partial C}dC + \frac{\partial V}{\partial K}dK \\ &\quad + \frac{1}{2}\frac{\partial^2 V}{\partial^2 P}dP^2 + \frac{1}{2}\frac{\partial^2 V}{\partial^2 C}dC^2 + \frac{1}{2}\frac{\partial^2 V}{\partial^2 K}dK^2 \\ &\quad \frac{1}{2}\frac{\partial^2 V}{\partial P \partial C}dPdC + \frac{1}{2}\frac{\partial^2 V}{\partial P \partial K}dPdK + \frac{1}{2}\frac{\partial^2 V}{\partial K \partial C}dKdC, \end{aligned}$$

and substituting equations (6.7), (6.8), (6.9) into the corresponding Bellman equation for optimality (see Pindyck (1993) [67]) with the final condition :

$$V_{\tilde{t}_S}(P_{\tilde{t}_S}, C_{\tilde{t}_S}, K_{\tilde{t}_S}) = \max(h_{\tilde{t}_S}(P_{\tilde{t}_S}, C_{\tilde{t}_S}), 0). \tag{6.11}$$

Here $h_{\tilde{t}_S}(P_{\tilde{t}_S}, C_{\tilde{t}_S})$ is given by Equation (6.10).

Solving the partial differential equation so obtained can be cumbersome due to the free boundary condition, as the date at which the power plant starts its operation, $\tilde{t}_S$, is a random variable. The problem considered here has a dimensionality of three, but in practice it can be even higher, which makes the use of finite difference based methods for solving the above PDE cumbersome. Like Schwartz (2004) [76], a simulation-based approach is used here to solve the optimal investment decision problem.

Computing the real option value using simulation

It is assumed that a complete probability space $(\Omega, \mathscr{F}, \mathbb{P})$ and finite time horizon $[0, T]$, with Ω the set of all possible realizations of a stochastic economy between 0 and T. The information structure in this economy is represented by an augmented filtration $\mathscr{F}_t : t \in [0, T]$, and $\mathbb{P}$ is the probability measure on elements of $\mathscr{F}$. Further it is assumed that the state of economy is represented by an $\mathscr{F}_t$-adapted Markovian process (P_t, C_t, K_t), i.e. the electricity price rate, the generation cost rate and the expected cost of completion of the power plant, respectively, at time t. The state space is generated at discrete time steps and for simplicity the time horizon is divided into M equal parts, with $t \in [t_0 = 0, \ldots, t_m, \ldots, t_M = T]$. The length of each time step is equal to

$$\Delta t = \frac{T}{M}.$$

The simulation begins by generating N stochastic paths for the remaining expected construction cost K_t, generation cost C_t and electricity price rate P_t. The vector $P_{t_m}(n), C_{t_m}(n), K_{t_m}(n)$, where $n \in [1, \ldots, N]$ and $m \in [0, \ldots, M]$, defines a unique state at time step t_m. The random cost of completion paths are simulated using the following discrete approximation to Equation (6.7).

$$K_{t_{m+1}}(n) = K_{t_m}(n) - I\Delta t + \beta (IK_{t_m}(n))^{\frac{1}{2}} (\Delta t)^{\frac{1}{2}} X_\beta + \gamma K_{t_m(n)} (\Delta t)^{\frac{1}{2}} X_\gamma, \tag{6.12}$$

where X_β, X_γ are uncorrelated standard normal variates. Time point $\tilde{t}_S(n)$ is the first time step at which $K_t(n)$ reaches a value less than or equal to zero and $K_t(n)$ is set to zero for all $t \geq \tilde{t}_S(n)$. Figure 6.1 shows a few of the scenario paths obtained using Equation (6.12), and Figure 6.2 gives an example of the distribution of the

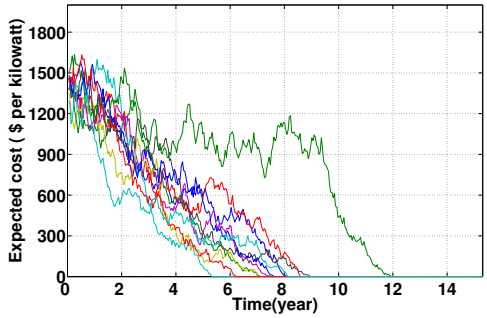

Figure 6.1: Sample paths for expected cost of completion at different time steps

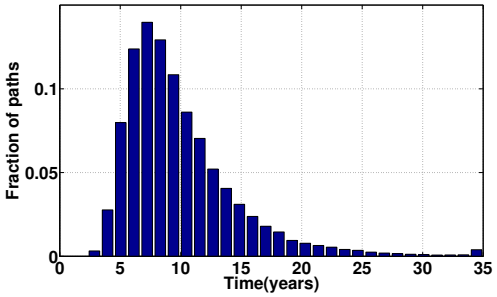

Figure 6.2: Distribution of construction time when construction costs are uncertain.

total construction time. The generation cost rate C_t and the electricity price rate P_t paths are simulated as:

$$C_{t_{m+1}}(n) = C_{t_m}(n) e^{\left(\mu_c^* - \frac{1}{2}\sigma_c^2\right)\Delta t + \sigma_c \sqrt{(\Delta t)} X_C}, \tag{6.13}$$

$$P_{t_{m+1}}(n) = P_{t_m}(n) e^{\left(\mu_p^* - \frac{1}{2}\sigma_p^2\right)\Delta t + \sigma_p \sqrt{(\Delta t)} X_P}, \tag{6.14}$$

where X_γ, X_C and X_P are standard normal variates that can be correlated.

Time horizon T is taken sufficiently long, so that the construction of the plant is almost surely finalized before T, i.e. $\tilde{t}_S < T$ with very high probability.

The real option value problem, like its financial counterpart the Bermudan option, is solved backwards in time, starting from the final time step, $t_M = T$. For those paths where the construction of the plant is finalized the option value at any time

step is given by Equation (6.10). Particularly, the option value at the time point at which the plant becomes operational is given by:

$$V_{\tilde{t}_S}(P_{\tilde{t}_S}(n), C_{\tilde{t}_S}(n), 0) = e^{-(r_f-\mu_p^*)\tilde{t}_S} P_{\tilde{t}_S}(n) \frac{1-e^{-(r_f-\mu_p^*)L}}{r_f-\mu_p^*} \tag{6.15}$$

$$-e^{-(r_f-\mu_c^*)\tilde{t}_S} C_{\tilde{t}_S}(n) \frac{1-e^{-(r_f-\mu_c^*)L}}{r_f-\mu_c^*}, \tag{6.16}$$

where $n \in [1,\ldots,N]$ and L is the designed lifetime of the plant.

For those paths where investment is still ongoing the optimal decision to continue the investment is based on the *continuation value* $Q_{t_m}(P_{t_m}, C_{t_m}, K_{t_m})$, which is given by:

$$Q_{t_m} := e^{-r_f \Delta t} \mathbb{E}\left[V_{t_{m+1}} | P_{t_m}, C_{t_m}, K_{t_m}\right], \tag{6.17}$$

where the simplified notations Q_{t_m} and $V_{t_{m+1}}$ are used for $Q_{t_m}(P_{t_m}, C_{t_m}, K_{t_m})$, and $V_{t_{m+1}}(P_{t_{m+1}}, C_{t_{m+1}}, K_{t_{m+1}})$, respectively. It is optimal for the firm to continue with the investment, when the construction is not yet finalized, i.e. if $Q_{t_m}(n) \geq I\Delta t$, and abandon it otherwise. More intuitively, irrespective of how much the firm has already spent on the construction of the power plant, the optimal decision at a given state point is just based on whether the net future expected revenues are greater than zero. The option value at a state described by path n, at time step t_m, is then:

$$V_{t_m}(n) = \max(Q_{t_m}(n) - I\Delta t, 0). \tag{6.18}$$

Once the option value has been computed for all paths at t_m, the above process (6.17,6.18) is followed recursively moving backwards in time until one reaches the starting time t_0. The main challenge here is to efficiently compute the continuation value given by Equation (6.17), for which the *Stochastic Grid Bundling Method* (SGBM) is used, details of which are discussed in Chapter 3.

The policy for continuing or abandoning the construction of the plant obtained above is used to compute the real option value, i.e. the expected discounted cashflow, and the distribution of the net cashflow obtained following the optimal policy. The mean and the distribution of the optimal cashflow are required as inputs for the portfolio optimization step described in Section 6.1. To compute them another set of N_l paths[6] is generated and the policy computed above is applied to continue or abandon the construction of the plant. If the n-th path enters the critical zone, i.e. reaches a state $(P_t(n), C_t(n), K_t(n))$ where it is optimal to abandon,

[6] Fresh paths are generated as using the same set of paths that were used to obtain the optimal policy may result in an option value which is biased high, due to perfect foresight (or over-fitting).

the plant is abandoned for that path and revenues for the path are set to zero, i.e. Revenue(n) = 0. The costs incurred until the plant was abandoned are discounted to time t_0 to :

$$\text{Cost}(n) = \sum_{t=t_0}^{\tilde{t}_a(n)} e^{-r_f t} I \Delta t,$$

where $\tilde{t}_a$ is the first time the path enters the abandonment region. For those paths whose construction is successfully completed (i.e. the paths which never enter the abandonment region), the revenues as seen at time t_0 are:

$$\text{Revenue}(n) = e^{-r_f \tilde{t}_S} V_{\tilde{t}_S}(P_{\tilde{t}_S}(n), C_{\tilde{t}_S}(n), 0), \tag{6.19}$$

and the costs of construction of the plant, discounted to time t_0, are:

$$\text{Cost}(n) = \sum_{t=t_0}^{\tilde{t}_S} e^{-r_f t} I \Delta t, \tag{6.20}$$

where $\tilde{t}_S(n)$ is the time when the plant starts its operation along the n-th scenario path. The real option price or the net expected cash-flow following the optimal policy of the power plant is then given by

$$V_{t_0}(P_{t_0}, C_{t_0}, K_{t_0}) = \frac{1}{N_l} \sum_{n=1}^{N_l} (\text{Revenue}(n) - \text{Cost}(n)). \tag{6.21}$$

The option price so obtained is a *lower bound*[7] of the true price as the policy used is generally sub-optimal due to numerical errors involved.

6.3 Validation: A Case from Pindyck

Pindyck (1993) [67] examined the decision to start or continue building of a nuclear power plant. To apply the model the estimates of the expectation and variance of the cost of building a kilowatt of nuclear generating capacity are used. The variance is decomposed into two parts to obtain estimates for technical uncertainty and input cost uncertainty. The survey of individual nuclear power plant costs published by the Tennessee Valley Authority (1977 to 1985) was used, which provided data on expected cost of a kilowatt of generating capacity on a plant-by-plant basis. A cross-section regression analysis over time was employed to estimate the expected costs and variance of a power plant. The variance of the costs and their decomposition were estimated from time-series and cross-sectional variations of the data,

[7] Lower bound implies that if the same Monte Carlo simulation is performed several times, with different initial seeds, the mean of V_{t_0} so obtained would be lower than V_{t_0}.

Initial expected cost K_0	\$ 1435 per kilowatt
Investment rate I	\$ 144 per annum
Discount factor r	0.045
Life Time of reactor	40 years
Revenue	\$ 2000 per kilowatt or 1.23 cents per kWh

Table 6.2: Parameter set used for validation case

	Pindyck		SGBM		LSM	
β	V_{t_0}	K_0^*	V_{t_0}	K_0^*	V_{t_0}	K_0^*
0	121	1550	120.64	1550.5	120.64	1550.5
0.24	131	1609	128.89	1582	128.75	1612
0.59	215	1881	211.46	1798	210.36	1887

Table 6.3: The real option value and critical expected construction cost for different levels of technical uncertainties. K_0^*, is the critical expected construction cost at time t_0, above which the project should not be undertaken.

using the fact that the variance of cost due to technical uncertainty is independent of time, whereas the variance due to input cost fluctuations grows with time. Based on these estimates the technical uncertainty parameter β in (Pindyck (1993) [67]) is found to vary from 0.24 to 0.59, while γ in (Pindyck (1993) [67]) varies between 0.07 to 0.2. In this analysis an instant revenue as soon as the construction is finalized was considered.

As a first numerical validation experiment, like Pindyck (1993) [67] the parameter set given in Table 6.2[8] is used. Table 6.3 compares the values reported by Pindyck with those obtained using the simulation method SGBM as well as the *least squares method* (LSM) (see Longstaff and Schwartz (2002)[53], Schwartz (2005)[76] for details on LSM), for different levels of technical uncertainty. It can be seen that without uncertainties in the construction costs the closed-form solution and results from simulations are almost identical, where a minor difference is due to the discretization of Equation (6.7). When technical uncertainty, β, is non-zero the real option values from simulation are slightly lower than the closed form values from Pindyck (1993) [67], as simulation results are biased low. The option values obtained using SGBM are slightly higher than those obtained using the *least squares method* for the same set of paths, which implies that in the discrete time version the critical costs for abandonment, K_0^*, obtained using SGBM are somewhat more accurate.

The role of real options in computing the net expected cashflow and its distribution when a firm is flexible to take decisions during the course of construction and operation of the reactor, needs to be emphasized. If the underlying stochastic fac-

[8] Note that prices are in USD here in accordance to the reference values from the literature.

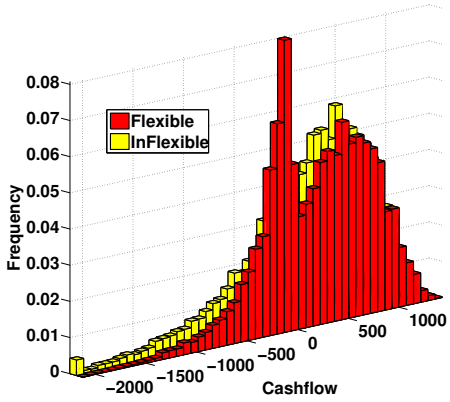

Figure 6.3: Distribution of net cashflow when the firm has flexible and inflexible decision opportunities to abandon the project in the future. The policy, when early abandonment is possible, is computed using SGBM. Same set of scenario paths is used for the two cases.

tors like expected cost of completion turn unfavourable in the future the firm uses its discretion to abandon the project[9] in such a way that the net expected cashflow is maximized. Figure 6.3 compares the cashflow distribution when (a) the firm doesn't have the flexibility to change its decision in the future and continues with the construction of the reactor irrespective of whether the scenario is favourable or not, (b) the firm has the flexibility to change its decision and continues or abandons the project following the policy computed using SGBM. It can be seen that the option to abandon the project under unfavourable price scenarios reduces the possibility of extreme losses. Table 6.4 compares the expectation and standard deviation of the net cashflow for the above two cases.

Figure 6.4 shows the fraction of scenario paths for which the project is abandoned at different time steps when the policy from SGBM is followed for the above case. It's more likely for a project to be abandoned in its early phases than in later stages. As the project commences the remaining expected construction costs (due to the ongoing investment) and also the remaining expected time to finish the construction reduce while the anticipated revenues increase (as the revenues are expected to start flowing in relatively sooner, which implies they are discounted less), which reduces the chance of the project being abandoned.

[9] It is assumed that the firm behaves rationally throughout the life cycle of construction and operation of a nuclear power plant, although there is some empirical evidence which suggests that management might act otherwise when sunk costs are involved, for example see "Throwing good money after bad ? : Nuclear power plant investment decisions and the relevance of sunk costs"by Bondt and Makhija (1988) [24].

	Inflexible Case	Flexible case (SGBM)
Expected net cashflow ($/kWe)	186	221
Standard deviation	600	500

Table 6.4: The expected value and standard deviation of the net cashflow corresponding to the distribution in Figure 6.3. The reactor parameters are taken from Table 6.2 and $[\beta, \gamma]$ values are [0.59 0.07], respectively.

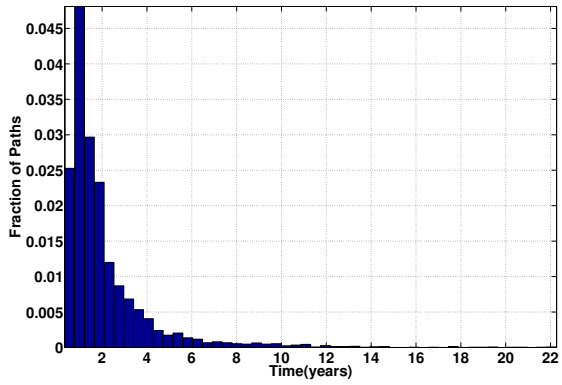

Figure 6.4: Fraction of paths abandoned at different time steps, when the policy from SGBM is followed, corresponding to the case considered in Table 6.4.

6.4 Numerical Examples

This section illustrates by various examples the two steps involved in deciding the optimal mix of NPPs for a power utility (or country or otherwise). A more realistic case, where not only the costs are uncertain but also the market price of electricity, is considered. The real option value, optimal decision rules to start, continue or abandon the construction of a reactor, distribution of costs and cashflows obtained following the optimal policy for different reactors, are analyzed. Finally, based on the expected net cashflow and its distribution, the optimal reactor order fraction for the different reactors considered is found.

Choice of nuclear power plants

This section discusses the economics of different nuclear reactors, considered here, for determining an optimal portfolio in energy generation planning. Here, not only

the expected costs of completion of the reactors are uncertain, but also the source of revenues, i.e. the electricity prices. The optimal decisions do not only depend upon the expected costs of the reactor but also on the present market price of electricity. Under our model assumptions, the construction of an unfinished reactor continues as long as the expected cost of completion is below some critical value and the electricity prices are above the corresponding threshold electricity price. For a given expected construction cost if the present electricity price (per annum) falls below a threshold the expected net cashflow would be negative and hence the construction of the plant is discontinued in our model. Similarly for a given electricity price if the expected cost of completion increases above a threshold price the construction of the plant will be abandoned in our model.

For the analysis the following types of reactors are considered for the portfolio.

- *Generic Gen III type Light Water Reactor (LWR):* The light water reactor (LWR) is a type of thermal reactor that uses water as its coolant and a neutron moderator and solid compound of fissile elements as its fuel. Thermal reactors are the most common type of nuclear reactor, and light water reactors are the most common type of thermal reactor.

- *Fast Reactors (FR):* Fast reactors or fast neutron reactors are a category of nuclear reactors in which the fission chain reaction is sustained by fast neutrons. They are considered an attractive option because of their potential to reduce actinide wastes, particularly plutonium and minor actinides which eliminate much of the long-term radioactivity from the spent fuel. Fast reactors with closed fuel cycle allow a significantly improved usage of natural uranium. The Sodium Cooled Fast Reactor (SFR), Lead Cooled Fast Reactor (LFR) and Gas Cooled Fast Reactor (GFR) are examples of fast reactors featured in the Generation IV roadmap (2002) [33].

- *High Temperature Reactor (HTR):* Also featured in the Generation IV roadmap, HTRs are graphite-moderated nuclear reactors with a once-through uranium fuel cycle. The high temperatures enable applications such as an emission-free process heat or hydrogen production, which effectively increase the efficiency of the reactor by as much as 20% (Generation IV roadmap (2002)) [33].

- *Super Critical Water Reactor (SCWR):* Featured in the Generation IV roadmap, SCWRs resemble light water reactors (LWRs) but operate at higher pressure and temperature, with a direct once-through cycle like a boiling water reactor (BWR), with the water always in a single fluid state like the pressurized water reactor (PWR). The SCWR is an advanced nuclear system because of its high thermal efficiency of 45% vs. 33% for current LWRs, and simple design (Generation IV roadmap (2002)) [33].

The size, efficiency and capacity factors of the reactors considered, taken from Roelofs *et al.* (2011) [69], are given in Table 6.5.

	Power (Thermal) MW	Power (Electric) MW	Efficiency (%)	Capacity Factor (%)	Life Time (years)
Gen III	4500	1600	35.5	90	60
FR	3600	1500	42	85	60
HTR	500	(200 +100)	(40 + 20)	90	60
SCWR	2300	1000	43.5	90	60

Table 6.5: The specification of the reactors considered.

Notice that HTRs have an efficiency of 40% + 20%, as not only would the reactor produce 200 MW of electricity, but also 100 MW of process heat. This is incorporated in the model by assuming that the cost of electricity is 2.32 times the process heat costs, as in Gandrik (2012) [32], which results in a revenue for this reactor equal to $1.21 \times P_t$.

The reference values for the expected construction costs, fuel cycle costs, operation and maintenance costs and also the confidence interval or standard deviation of these costs are taken from van Heek *et al.* (2012) [85] and Roelofs *et al.* (2011) [69]. These are engineering cost estimates as there isn't sufficient experience to estimate these values from historical data. Table 6.6 reports the expected construction costs and fuel, operation and maintenance costs as derived from the values in van Heek *et al.* (2012) [85] and Roelofs *et al.* (2011) [69]. In the case of the HTR additionally the benefits of modular construction (increased standardisation and faster learning curves) are included which is different from van Heek *et al.* (2012) [85]. The analysis of Boarin & Ricotti (2011) [11] is followed, where four effects of modular construction are distinguished:

1. *Learning factor:* The number of similar plants constructed world-wide will lead to increased experience in construction and therefore in decreased costs;
2. *Modularity factor:* The modularization factor assumes a capital cost reduction for modular plants, based on the reasonable assumption that the smaller the plant size, the higher the degree of design modularization;
3. *Multiple units factor:* The multiple units saving factor shows a progressive cost reduction due to fixed cost sharing among multiple units at the same site;
4. *Design factor:* The design factor takes into account a cost reduction by assumed possible design simplifications for smaller-sized reactors.

Figure 6.5 shows the curve constructed when all these separate effects are combined. A fitted curve that gives the modular construction factor is then given by,

$$mcf = \min\left(0.195\ln\left(\frac{\text{Power}_{mod}}{100}\right) + 0.63 - 10^{-4} \times \text{Power}_{ref}, 100\%\right), \tag{6.22}$$

where Power_{ref} is 1100 MWe and Power_{mod} is on the x-axis of Figure 6.5. Following Equation (6.22) based on the assessment of Boarin *et al.* (2011) [11], the modularity

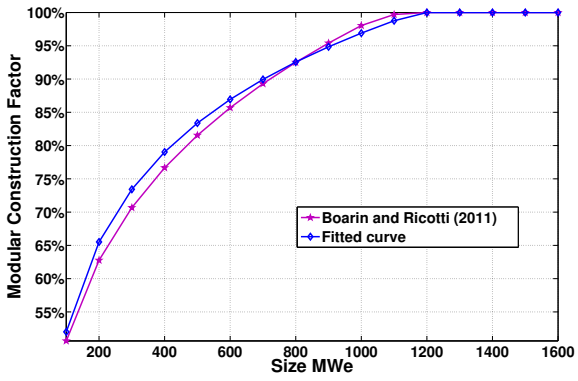

Figure 6.5: Modular construction factor as a function of size of the reactor. The reference power plant size is 1100MWe.

Reactor	Expected construction cost €/kWe	Expected Fuel and O&M cost €/kWe/a
Generic Gen III LWR	2900 (320)	140 (35)
FR	4600 (580)	185 (35)
HTR	3600 (750)	165 (35)
SCWR	3400 (400)	140 (33)

Table 6.6: Expected construction, fuel and O&M costs for different reactors considered. The values in brackets are standard deviations of these costs.

construction factor would be 65.5% for a Power$_{mod}$ = 200 MWe HTR, which brings down the expected costs of construction of HTRs from 6100 to 3600 €/kWe.

As the values reported in Table 6.6 are "engineering estimates" the uncertainty in these values can primarily be attributed to technical uncertainty. When only technical uncertainties are involved the variance of the expected cost of construction is given by (see Pindyck (1993) [67]) :

$$\text{Var}(K) = \left(\frac{\beta^2}{2-\beta^2}\right) K^2;$$

a relation used here to compute the corresponding value of β for different reactors in the portfolio. The Brownian motions driving the input cost uncertainties (see Section 6.2) for different reactors can be correlated to each other (and the economy), as raw material required and government regulations are similar for differ-

Generic GenIII	K_0 = 2900 (€/kWe), γ=0.07, β = 0.15, expected construction time = 5 years, C_0 = 1.36 (cents/kWh), σ_C = 0.25;
Generic Fast Reactor	K_0 = 4600 (€/kWe), γ =0.07, β = 0.18, expected construction time = 7 years, C_0 = 1.95 (cents/kWh), σ_C = 0.19;
HTR	K_0 = 3600 (€/kWe), γ=0.07, β = 0.17, expected construction time = 4 years, C_0 = 1.70 (cents/kWh), σ_C = 0.22;
SCWR	K_0 = 3400 (€/kWe), γ = 0.07, β = 0.16, expected construction time = 5 years, C_0 = 1.43 (cents/kWh), σ_C = 0.24;

Table 6.7: Initial expected cost of completion, input cost uncertainty parameter γ, technical uncertainty parameter β, expected construction time, present value of combined O&M and fuel charges C_0 and the corresponding volatility for different reactors. For all cases considered it is assumed that the correlation coefficient ρ between W_P and W_C to be 0.5 and the growth rate in O&M costs, μ_C^* to be 0. The rate of investment I for each reactor is taken as their initial expected construction costs divided by their expected construction times.

ent reactors, while technical uncertainties for different reactors are assumed to be uncorrelated. Table 6.7 summarizes the parameter choices related to Table 6.6.

The real option value of the reactors and the distribution of the net cashflow under optimal policy for construction and operation of the reactors, depends on, amongst others, the expected growth rate for electricity prices (μ_p^*), uncertainty in electricity prices (σ_p), and the discount rate used (r)[10]. Table 6.8 gives values considered for these parameters. For the base case, values corresponding to the row 'Medium' in Table 6.8 are taken, and the initial price of electricity P_{t_0} is set to 8.5 cents/kWh.

[10] As the Brownian motions dW_γ, dW_C, dW_P may be correlated with the market, the risk-free interest rate for discounting cannot be used, especially if spanning is not possible. Instead different discount rates which represent different levels of risk premiums added to the risk-free rate, are considered.

	Growth rate μ_p^* (% per annum)	Uncertainty σ_p (% per annum)	Discount rate r (% per annum)
Low	0	10	6
Medium	3	20	8
High	5	30	10

Table 6.8: Values of electricity price growth rate μ_p^*, uncertainty in electricity prices, σ_p and discount rate r considered in various examples.

Real option value analysis

Real option analysis is used here to determine the optimal policy to start, continue or abandon the construction of a project, so that the net expected discounted cash-flow is maximized. As stochastic construction costs K_t, combined O&M and fuel cycle costs C_t, and cost of electricity P_t, are considered, the optimal decision will depend on these three state variables. It will be optimal to abandon the project, if :

- the expected cost of completion is too high,
- the O&M and fuel cycle costs are too high,
- the electricity prices are too low.

Figure 6.6 shows the early abandonment region at an intermediate time step of the simulation. Here the $x-$ axis represents the expected costs of completion of the reactor and the $y-$ axis represents the cost of electricity minus fuel and O&M costs. The red coloured grid points represent the states at which the construction of the reactor should be abandoned, while green colour represents the ones for which the construction should continue.

Table 6.9 reports the critical price of electricity above which each of these reactors should be ordered and their real option values when the initial price of electricity equals $P_0 = 8.5$ cents/kWh. Reactor specific parameters are taken from Table 6.7. The same set of simulated electricity paths should be used for different reactors.

Under the present model assumptions and parameter choices, it is observed that the HTRs, despite their high expected capital costs, appear economically most attractive, primarily due to their higher efficiencies. The Gen III LWRs have the lowest critical electricity price above which they can be ordered, while the fast reactors seem economically least viable in our model settings.

Optimal portfolio analysis

If a firm has to choose amongst the above reactors, solely based on their capital costs (Table 6.7), then their portfolio would contain only Generic Gen III type LWRs,

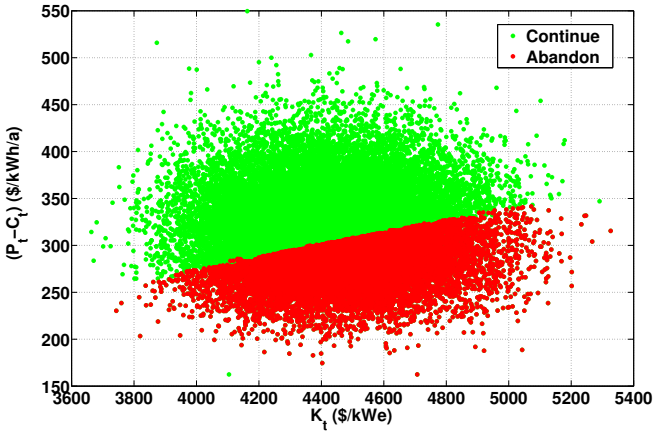

Figure 6.6: Simulated K_t, P_t, C_t paths at a particular time step. Green represents the states where it is optimal to continue and red where it is optimal to abandon the project.

Type	Critical electricity price P_0^*	Option value $P_0 = 8.5$ cents/kWh
Gen III LWR	4.0	3100
FR	6.25	875
HTR	4.5	3500
SCWR	4.7	2650

Table 6.9: Critical price of electricity $P_{t_0}^*$ in (euro cents/kWh) above which the reactors should be ordered and their option values (in €/kWe) when the initial price of electricity is 8.5 euro-cents/kWh. The reactor parameters are taken from Tables 6.7 and 6.8.

something also observed in practice. However, such a portfolio excludes the role of uncertainties of cashflows for these reactors. Application of MVP theory takes into account not only the expected returns but also the uncertainties or risks associated with these returns.

An efficient frontier gives the optimal reactor order fraction for a portfolio designed to meet a given expected return while minimizing the uncertainties of these returns. In order to determine the efficient frontier the expected returns and the covariance matrix of the returns from the reactors considered are required. The distribution of returns for each reactor optimally constructed is sampled by computing the returns along each simulated path.

The following constraints on the portfolio are considered:

- *Budget constraint:* Under a budget constraint, the optimal reactor order fraction for every euro spent is computed. Returns corresponding to a euro spent

on a reactor are given by,

$$R_i(n) = \frac{\text{Revenue}_i(n) - \text{Cost}_i(n)}{\text{Cost}_i(n)}, \tag{6.23}$$

and the constraint for the portfolio optimization problem is then:

$$\sum_{i=1}^{J} w_i = 1,$$

$n = 1,\ldots,N$ being the path index and $i = 1,\ldots,J$ again indicate the different reactors considered. The weights correspond to the fraction of money invested in different reactors, which is then used to compute the reactor order fraction (per kWe) by taking into account the expected construction costs as reported in Table 6.7.

- *Capacity constraint:* Under a capacity constraint, the optimal reactor order fraction for every kWe of capacity ordered is computed. Returns corresponding to a kWe ordered are given by,

  $$R_i(n) = \text{Revenue}_i(n) - \text{Cost}_i(n),$$

  and the constraint for the portfolio optimization problem is:

  $$\sum_{i=1}^{J} w_i = 1,$$

  $n = 1,\ldots,N$ being the path index, and $i = 1,\ldots,J$ indicate the different reactors considered. The constraint implies here that reactor order fractions should add up to a kWe.

For both constraints, the weights are additionally bounded as,

$$0 \leq w_i \leq 1,\ i = 1,\ldots,J,$$

which comes naturally from the fact that short selling is not possible here, and thus the weights cannot be negative.

The quadratic programming problem expressed by Equation (6.5) is solved using the optimization toolbox of MATLAB, which solves general problems of the kind:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2}\mathbf{w}^\top Q w + f'\mathbf{w}, \\ \text{such that:} \quad & A\mathbf{w} \leq a, \\ & B\mathbf{w} = b, \\ & \mathbf{L} \leq \mathbf{w} \leq \mathbf{U}, \end{aligned} \tag{6.24}$$

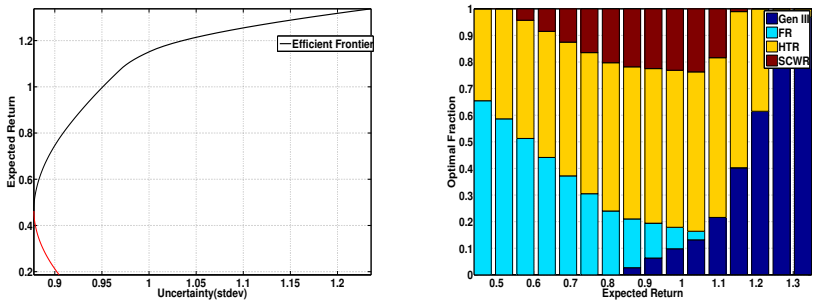

(a) (b)

Figure 6.7: (a) Efficient frontier for the base case when portfolios have the budget constraint and (b) reactor order fraction corresponding to points on the efficient frontier.

	Gen III	FR	HTR	SCWR
Expected Return	1.3376	0.1863	1.0645	0.9744
Stdev Return	1.2356	0.9046	0.9926	1.0527

Table 6.10: The expected return and its standard deviation per euro spent for the base case.

using the command w = quadprog(Q,f,A,a,B,b,L,U).

Figure 6.7 displays the efficient frontier and the corresponding optimal reactor order fraction when the portfolio has the *budget constraint*. The mean and standard deviation of the simulated returns for the individual reactors are reported in Table 6.10. Under our model assumptions and choice of parameter values, the GenIII LWRs have the highest expected returns (based on Equation (6.23), while the FRs have the lowest returns per euro spent. However, the uncertainty of returns for FRs is lower than that for GenIII LWRs. An investor who wants to minimize the uncertainty of returns and is willing to take a lower expected return in order to do so, will choose a portfolio with more Gen IV type reactors. An investor who wants higher returns and is indifferent to the uncertainty of returns, will hold a portfolio with more Gen III type reactors.

Figure 6.8 shows the efficient frontier and optimal reactor order fraction corresponding to points on the optimal frontier, when the portfolios have the *capacity constraint*. Expected returns and their standard deviations per kWe of reactor ordered are reported in Table 6.11. It is observed that unlike the case with the budget constraint, where portfolios with high returns were dominated by Gen III LWRs, here portfolios with higher expected returns are dominated by *both* HTR and GenIII LWRs. This difference can be explained as the returns in Equation (6.23) are scaled by the individual reactor costs.

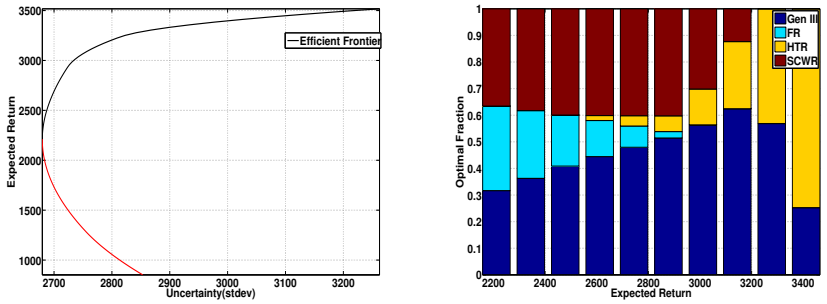

(a) (b)

Figure 6.8: (a) Efficient frontier for the base case when there is a capacity constraint and (b) the reactor order fractions corresponding to points on the efficient frontier.

	Gen III	FR	HTR	SCWR
Expected Return (€/kWe)	3100	880	3500	2625
Stdev Return	2825	2870	3225	2850

Table 6.11: The expected returns and their standard deviations per kWe of reactor ordered for the base case.

Portfolio sensitivity

In addition to the constraints on the portfolio, the choice of parameter values affects the structure of the optimal portfolio. The optimal portfolio for varying parameter values is studied here, which gives an intuition about the portfolio's sensitivity with respect to these parameters. In particular, the following cases are considered:

- Different discount rates r, with other parameters constant.
- Varying electricity price growth rates μ_p^*, with other parameters constant.
- Varying uncertainties in electricity prices σ_p, with other parameters constant.

From here on, only portfolios that have capacity constraints are considered.

Varying discount rates

The reference case considered a discount rate of 8% per annum. Here the portfolios sensitivity to varying discount rates is examined. A change in discount rate affects the expected revenues, costs and the optimal investment strategy, which in turn affects the returns. This makes the discount rate an important parameter while computing the efficient frontier and corresponding optimal reactor order fractions.

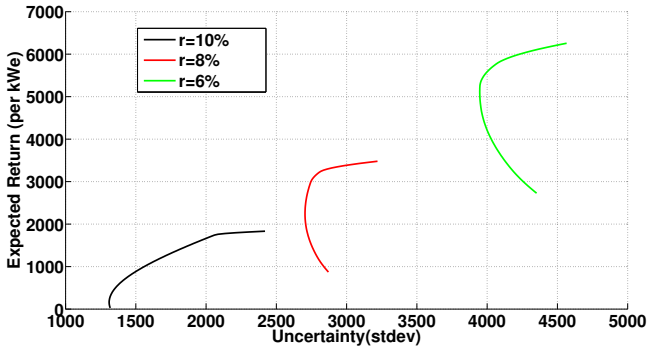

Figure 6.9: Efficient frontiers for varying discount rates. Parameter values are taken from Tables 6.7 and 6.8.

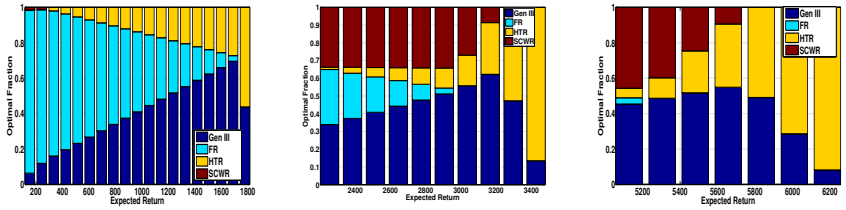

(a) (b) (c)

Figure 6.10: Optimal reactor order fractions when (a) $r = 10\%$, (b) $r = 8\%$, and (c) $r = 6\%$. Parameter values are taken from Tables 6.7 and 6.8.

Figure 6.9 shows the efficient frontier for low, medium and high discount rates, with corresponding values taken from Table 6.8. Lowering the discount rate can help realize higher expected returns, although at increased uncertainty (variance) in returns. Although both the expected returns and the variance of returns increases, however, the increase in the expected returns is more significant than increase in the variance of returns. Therefore, reactors with higher expected returns would then be more favoured in the mean-variance portfolio.

The optimal reactor order fractions corresponding to the points on the efficient frontier are shown in Figure 6.10. Under the model assumptions and parameter choices, it is seen that lowering discount rates results in a portfolio dominated by reactors having greater expected returns, while higher discount rates result in a portfolio where reactors with lower uncertainties dominate.

Varying electricity price growth rates

Long term growth rates of electricity prices are difficult to predict. A sensitivity analysis of the optimal portfolio with respect to different electricity price growth

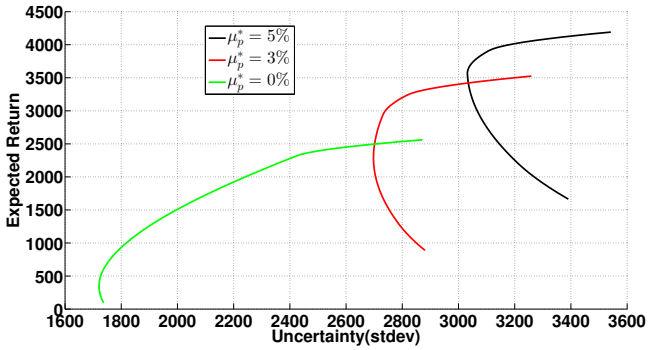

Figure 6.11: Efficient frontier for varying electricity price growth rate, where the reactor specific parameters are taken from Table 6.7, and economic parameters from Table 6.8.

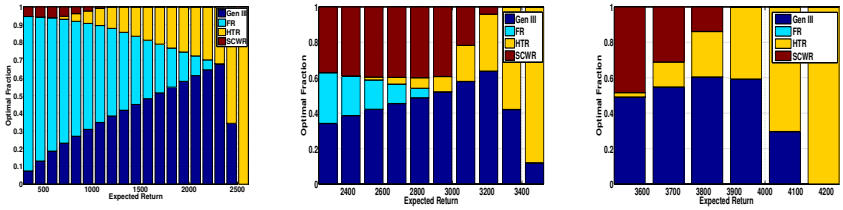

(a) (b) (c)

Figure 6.12: Optimal reactor order fractions when (a) $\mu_p^* = 0\%$, (b) $\mu_p^* = 3\%$. (c) $\mu_p^* = 5\%$. Parameter values are taken from Tables 6.7and Table 6.8.

rates is then essential. An optimal portfolio analysis for low, medium and high growth rate scenarios for electricity prices is performed.

Figure 6.11 shows the efficient frontiers corresponding to different electricity price growth rates. A higher growth rate in electricity prices results in portfolios which can achieve greater expected returns.

The optimal reactor order fractions corresponding to the points on the efficient frontiers for different electricity price growth rates are shown in Figure 6.12. Under the model assumptions, it is observed that a higher expected growth rate in electricity prices leads to portfolios that are dominated by reactors with higher expected returns (HTR and GenIII), while for low growth rate scenarios optimal portfolios can have reactors with lower returns (like FRs).

Varying uncertainty in electricity prices

Uncertainty in electricity prices affects the expected return and its distribution for different reactors. The mean-variance portfolios for low, medium and high uncer-

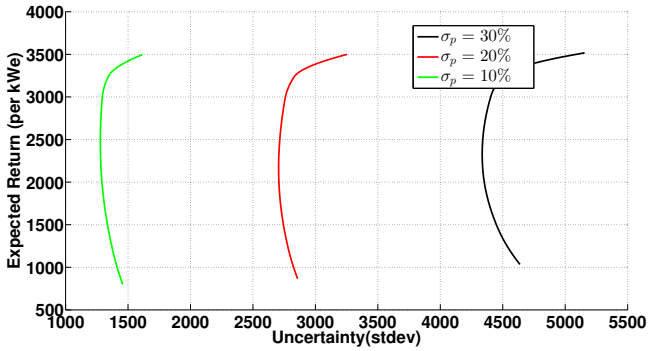

Figure 6.13: Efficient frontier for varying uncertainty in electricity prices, where parameter values are taken from Tables 6.7 and 6.8.

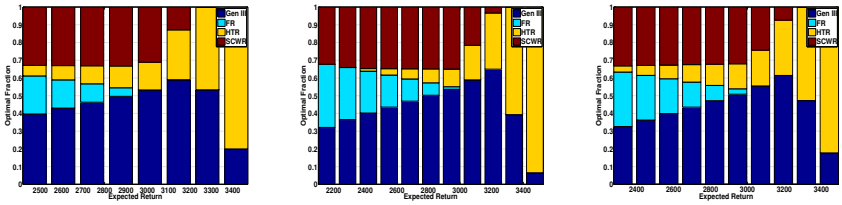

(a) (b) (c)

Figure 6.14: Optimal reactor order fractions when (a) σ_p = 10%, (b) σ_p = 20%, and (c) σ_p = 30%. Parameter values are taken from Tables 6.7 and 6.8.

tainty in electricity prices, with the corresponding values for σ_p taken from Table 6.8 are studied. Figure 6.13 plots the efficient frontiers for the three different scenarios considered. With increasing uncertainty in electricity prices, the uncertainty in returns of the optimal portfolio increases for a given level of expected returns.

The optimal reactor order fractions corresponding to the points on the efficient frontiers for the three scenarios considered are presented in Figure 6.14. Under our model assumptions, the reactor order fractions seem less sensitive to uncertainty in electricity prices, when compared to their sensitivity to discount rates or electricity price growth rates.

6.5 Conclusion

While the future of nuclear power depends on resolving the issues of safety of operations, safe management of radioactive wastes and measures to prevent proliferation (MIT, 2003) [26], in a deregulated electricity market, the economics of NPPs will a very most important determinant of nuclear energy's role in the future global

energy mix. A decision-support tool, which takes into account major factors and their uncertainties for studying the economics of individual reactors as well as for a portfolio of reactors has been presented here.

Specifically, real option analysis and portfolio optimization have been used to study optimal reactor order fractions within the nuclear sector. A two-step approach is proposed, where first optimal decisions are taken at the plant level, and then the resulting distributions of returns for each reactor-type are used as inputs to a portfolio optimization problem solved using MVP theory. The main contributions on the methodological side can be stated as:

- The method adequately accounts for uncertain reactor construction costs and schedule, and reflects their effect on the return distribution for different reactors.
- An optimal policy for continuing the construction or abandoning of the project is computed taking into account the uncertainties in construction costs, electricity prices and O&M and fuel cycle costs involved.
- A detailed study on optimal portfolios based on MVP theory is conducted.
- The effect of different constraints on portfolio diversification is studied.
- The sensitivity of the optimal portfolio with respect to electricity price growth rates, uncertainty in electricity prices and discount rates is studied.

It should be emphasized here that, although careful attention has been paid to choose realistic parameter values for the reactors considered, however, the main focus is to illustrate a methodology that accounts for the various economic uncertainties related to nuclear power plants. Under our model assumptions, it has been shown that certain scenarios lead to portfolios that are dominated by Generation IV type reactors, while others result in conventional Gen III type LWRs being the dominant ones. Following the methodology described here can be useful when decisions related to reactor order fraction need to be made.

As possible future direction of research, the portfolio optimization step should in addition to the variance of the returns also consider other risk measures, such as, value at risk and conditional value at risk, as in [77]. The resulting portfolios will then not only minimize the variance of the returns but will also avoid reactors which are likely to be abandoned in the future.

CHAPTER

7 Conclusions and Outlook

This thesis we dealt with models for making investment decisions under uncertainty and used them to study different scenarios for investment decisions in the nuclear industry. The first half of the thesis was concerned with the development of algorithms for valuing and taking optimal early exercise decisions. As options with early exercise features are common in the financial industry, such as American and Bermudan options, the algorithm developed were motivated by these financial options. The latter half of the thesis adapted the pricing methods developed for financial options to value real options and take optimal decisions related to investment in nuclear power plants.

Chapter 2 introduced the stochastic grid method SGM, for pricing and exercising Bermudan options. SGM employed dynamic programming and least squares regression for option pricing. One of the main achievements of the algorithm was its ability to reduce a multi-dimensional problem to a problem that deals with single-dimensional features. The method had a few advantages, but it suffered from some disadvantages such as; the method was computationally expensive when sub-simulations were required, the early exercise policy obtained using SGM was inferior to those obtained using the least squares method (LSM) for high-dimensional problems, and therefore the option values from SGM were less accurate, SGM for high-dimensional problems required the use of peripheral paths to improve the early exercise policy but the approach to generate the peripheral paths was not well defined and the convergence was guaranteed only for the single-dimensional problem.

In order to deal with the drawbacks of SGM, Chapter 3 introduced the *Stochastic Grid Bundling Method*, SGBM. SGBM is a hybrid of regression-based and bundling-based Monte Carlo methods and appears computationally at least as attractive as existing methods. SGBM overcame the drawbacks of SGM by utilizing bundling to improve the approximated conditional expectations involved in the valuation. The chapter provided a proof for convergence of the method for high-dimensional problems, and demonstrated through numerical examples the rate of convergence with respect to the number of bundles chosen in the method. One of the main advantages of SGBM is that for similar computational time a higher accuracy, not just at the final time step, but also at intermediate time steps is obtained when

compared to the least squares method of Longstaff and Schwartz [53]. This makes SGBM a suitable candidate for computing upper bounds using duality-based methods. Another advantage of SGBM is that it can be used for fast approximation of the option price sensitivities.

Investment in small and medium size reactors which can benefit from modular construction has been an attractive option in the nuclear industry. Chapter 4 presented a method for computing real option values for the construction of nuclear power plants when the project can benefit from the flexibility offered by modular construction in finite decision time horizon. SGM from Chapter 2 was used as the underlying pricing method. It was found that decision-making in finite time, which is often the case, may result in quite different scenarios when compared to infinite time decision horizon problems. Chapter 4 showed that under model assumptions in comparison to the modular project the real option value of a single large reactor is typically higher when decision horizons are small, while for longer decision times modular projects may be more profitable. An interesting result in Chapter 4 was that with an increasing uncertainty in the electricity prices modular construction may represent a better option as they include the possibility of having a few units ordered, when the electricity price reaches unfavourable values. With stable electricity prices the cost effective single large reactor appeared under the assumptions to be a better choice.

The real option model developed in Chapter 4 was further exploited in Chapter 5 to value and compare different construction strategies of NPPs for finite decision horizon. An finding of the chapter was that the real option value of nuclear power plants increases with implementation of long term cost reductions. Such cost reductions might be achieved e.g. by plant upgrades and/or increased R&D efforts. It is known that when twin units are constructed at the same site, significant cost reductions can be achieved. On the other hand, in order to achieve these cost savings, the utility loses the flexibility to order units at optimal market conditions. Chapter 5 shows that when the electricity price uncertainty is low the cost savings by a construction of twin modules at the same site would be more profitable under the model assumptions, while for volatile electricity prices, the flexibility of choice dominates. In this chapter a basic model to include the effect of uncertain lifetime of operations and premature permanent shutdown of nuclear reactors on the real option value of projects was also included. It was found that some cost reduction was achieved by the learning effect with each new module in a modular construction. However, under the model assumptions considered in Chapter 5 it appeared that cost savings due to learning are not sufficient to make modular SMRs competitive with large units.

Finally Chapter 6 presented a decision-support tool which took into account uncertainties in major economic factors for individual reactors to provide, under model assumptions, an optimal portfolio of reactors. Real option analysis together with mean-variance portfolio optimization was used to compute optimal reactor order fractions. The decision tool used a two-step approach; where first optimal decisions were taken at the plant level, and then the resulting returns distributions for

each reactor-type were used as input for a mean-variance portfolio optimization problem. The tool could take into account uncertain reactor construction costs and schedule, and compute their effect on the return distribution for different reactors. The chapter also presented the optimal portfolio and its sensitivity with respect to electricity price growth rates, uncertainty in electricity prices and discount rates. It was observed that under model assumptions certain scenarios led to portfolios that were dominated by Generation IV type reactors, while others resulted in conventional Gen III type LWRs being the dominant ones.

7.1 Future Outlook

This section discusses possible future directions for research, both on the method development side that was discussed in Part 1 of the thesis, as well as on the application side that was dealt with in Part 2 of the thesis.

The proof for convergence of SGBM was provided in Chapter 3. Although it was shown that the method would converge to the true price as the number of paths and bundles approach infinity, the rate of convergence of the method with increasing number of bundles was only demonstrated through numerical examples. A detailed numerical analysis however can provide the optimal number of bundles that should be chosen for a given number of total paths. The intuitive reason behind the existence of an optimal number of bundles is that, while on one hand increasing the number of bundles results in reduced error as result of more accurate sampling of the conditional distribution, on the other hand with an increasing number of bundles the number of paths within each bundle reduces when the total number of paths is fixed. This results in larger standard error for the Monte Carlo estimates within the bundle, as this error is inversely proportional to the square root of number of samples. If the number of bundles is more than the optimal number of bundles, the reduced error from accurate sampling of the conditional distribution would be lost by increased standard errors from too few paths per bundle.

At present the value of an option with early exercise feature is computed by finding a policy that maximizes the expected cashflows. An interesting direction for future research could be to examine policies that not only maximize the expected future cashflows but also, like a mean-variance portfolio, minimize the variance of the future cashflows. Such policies could be more stable as deviating slightly from the optimal policy wouldn't dramatically affect the expected future cashflows. Dynamic portfolio allocation problems have gained popularity in recent times and it would be interesting to compute them using SGBM, wherein constant portfolio weights are assigned to the states that belong to the same bundle.

A comprehensive analysis of different nuclear power plants cannot exclude the role of exogenous factors, such as the cost of carbon emissions, cost of natural gas etc, something which hasn't been dealt within the present work. Also in order to better evaluate nuclear power plants it is important to consider their role in long term nuclear waste management. Some reactors although at present might appear attrac-

tive, could become increasingly unattractive in absence of long term waste management solutions. For example the value of LWRs could significantly improve if the waste from it could be reprocessed, while it could become unattractive if it has to bear the economic burden for the development of expensive long term nuclear waste storage facilities. Such complex interactions between nuclear fuel cycle systems of different reactors need to be addressed for more informed decision making process in nuclear industry.

In Chapter 6 to compute the optimal reactor order fractions we used the mean-variance portfolio theory. It can happen that the optimal portfolio so computed may contain reactors that are very likely to be abandoned at a future date when following the computed policy for reactor abandonment. It would be ideal to construct the optimal portfolio such that the fraction of scenario paths for which all the reactors become operational lies in some high percentile range. This could be achieved by using alternate portfolio optimization techniques, such as those based on the conditional value at risk (CVaR).

7.2 Integrating decision-support tool with DANESS

Argonne National Laboratory (ANL) developed an integrated nuclear energy system model called Dynamic Analysis of Nuclear Energy System Strategies (DANESS). DANESS allows modelling the full mass-flow chain of time varying mixes of nuclear reactor plants and associated fuel cycle options (Van den Durpel et al. (2008) [84]). It has been used to simulate fuel cycles from uranium mining, reprocessing, to geological disposal. For any modelled combination of reactor types and fuel cycles, DANESS projects electricity production cost, fuel mass flows, and waste quantities as a function of time, spanning periods from decades up to centuries. Roelofs and Estorff (2013) [70] use DANESS to predict the top down workforce demand corresponding to different nuclear energy scenarios.

In DANESS new reactor types — characterized by techno-economic parameters, representing their fuel consumption and overall effectiveness — are introduced based on the requirement to fulfill a certain scenario dependent nuclear energy demand.

Given a nuclear energy demand scenario, DANESS orders reactors and associated facilities at discrete time steps so as to meet the time varying energy demand as closely as possible. DANESS determines ΔE, the energy generation shortage, by using a prediction of the energy demand, an assumed planning horizon and the installed reactor capacities. While determining ΔE, reactors which will be taken out of operation as well as reactors which are already under construction are also considered. DANESS covers small changes in ΔE by exploiting the flexibility in load factors of the different reactors. Adjustment in load factors reflects the possibility of a utility to decide to have a longer maintenance stop when the energy demand is low enough. Finally, if ΔE is large enough, reactors are constructed based on the specified reactor order distribution target.

DANESS orders from different available choices of reactor types to meet this energy shortage. The shares of different reactor types to be ordered are usually determined using some economic criteria. A criterion for the selection could be to order the reactor with the lowest levelised costs. However, this could lead to only a single reactor type to be ordered all the time, something that is not observed in practice. Another approach is to use the Reactor Order Decision Model (RODM) of Roelofs *et al.* (2011) [69] to determine the reactor order distribution to be used in DANESS. Briefly, RODM involves the following steps:

1. The levelised fuel cycle costs are determined for each reactor type and associated fuel cycle at every time step by the DANESS code. These levelised fuel cycle costs take into account all front-end and back-end costs.

2. The levelised costs are time averaged for each available reactor type i. Evidently, reactor types which are not available for construction, e.g. because they are not yet fully developed, are not taken into account. The time averaged costs are defined as C_i.

3. The initial reactor order distribution target is determined as $\frac{F_i}{\sum_i F_i}$, in which $F_i = \frac{\max(C_1,...C_i,...)}{C_i}$. In scenarios where the energy production cannot match the energy demand because of fuel limitations for a certain reactor type and associated fuel cycle, the target is adjusted towards a reactor type which does not suffer from fuel availability limitations.

The RODM model has it advantages and drawbacks. The advantages are that it is easy to implement and understand. Above that, it reflects the market conditions in a way that not all utilities select only the reactor with lowest levelised costs, and thus we see a mixture of reactors in the market. However, it can be argued that the scheme used in RODM is not based on strong economic or physical principles.

Alternatively, the reactor order fraction determined by the MVP based decision tool, as described in Chapter 6, can be used as the reactor order distribution in DANESS. The market share of reactors then ordered would still be diversified, however, the underlying rules for diversification would be based on established economic principles. We therefore recommend that in the near future, for analysis using DANESS, both the RODM and the MVP decision-tool should be used to determine the reactor order distribution. Conclusions should be drawn by comparing the future reactor park development results obtained from DANESS, in terms of revenues, waste management etc. for these two models.

7.3 Integrating Real options with DANESS to evaluate scenarios

Modelling the worldwide nuclear reactor park can be challenging as the nuclear fuel cycle includes feedback loops representing physical feedbacks within the system. Moreover, and most prominently, socio-political feedbacks in the decision-

making on the various available deployment pathways for nuclear energy (Van den Durpel et al. (2008) [84]).

DANESS can be used to model nuclear energy systems composed of various reactor types and fuel cycle facility technologies, as it allows for mutual exchange of actinide mass flows and keeps a time-varying track of various inventories involved. Also it helps in analysing how efficiently a facility would be utilized during its lifetime for some given nuclear energy demand scenario.

Given a nuclear energy deployment pathway DANESS can accurately compute the physical mass flows and other involved system dynamics. However, the nuclear energy deployment pathways that generally span time periods over several decades; can be highly uncertain, due to amongst others, socio-political and technical reasons. In addition there can be several deployment strategies for which a given nuclear energy demand scenario can be met efficiently. A challenging task then is to identify based on certain objectives an optimal deployment strategy.

One of the objectives to identify the optimal deployment strategy can be the cost-benefit analysis of the strategy. In the absence of any uncertainties this would be equivalent to identifying the cheapest deployment strategy. However, as usually is the case with economics of nuclear energy systems, the costs and benefits are highly uncertain. Deployment strategies then should ideally be valued using the Real Options based methodologies, as it involves valuing flexible decisions. Under uncertain scenarios, it is likely to switch from one planned deployment strategy to another, depending upon how the future unfolds. Any nuclear deployment strategy from which it is easier to switch to other possible strategies in the future would have an inherent value of flexibility. Integrating Real Options based valuation techniques for decision making process involved in scenario studies using DANESS can therefore lead to valuable results.

Curriculum vitae

Shashi Jain (1983) obtained his master degree in Mechanical Engineering specializing in Energy Technology at the Indian Institute of Technology, Madras, India in 2007. Before starting his PhD research, from August 2007 till March 2010, he worked as a software engineer for Intuit Inc. in Bangalore, India. In 2010 he started his PhD research to develop quantitative models for investments in nuclear power plants. He performed his research at Centrum for Wiskunde and Informatica (CWI), Amsterdam, Nuclear Research Group (NRG), Petten and TU Delft. During the period of his PhD research he spent 3 months as a visiting researcher at the Tata Institute of Fundamental Research (TIFR), Mumbai, India and 2 months as an intern for Areva, Paris, France.

List of publications

- **Chapter 2**
  S. Jain and C.W. Oosterlee, (2012). Pricing high-dimensional Bermudan options using the stochastic grid method. *International Journal of Computer Mathematics*, 89(9), 1186-1211.

- **Chapter 3**
  S. Jain and C.W. Oosterlee. The Stochastic Grid Bundling Method: Efficient Pricing of Bermudan Options and their Greeks.*submitted manuscript*.

- **Chapter 4**
  S. Jain, F. Roelofs and C.W. Oosterlee, (2013). Valuing modular nuclear power plants in finite time decision horizon. *Energy Economics*, 36(C), 625-636.

- **Chapter 5**
  S. Jain, F. Roelofs and C.W. Oosterlee, (2013). Construction strategies and lifetime uncertainties for nuclear projects: A real option analysis. *Nuclear Engineering and Design*, 265, 319-329.

- **Chapter 6**
  S. Jain, F. Roelofs and C.W. Oosterlee. Decision-support tool for assessing future nuclear reactor generation portfolios. *submitted manuscript*.

Bibliography

[1] N. Aleksandrov and B. Hambly. A dual approach to multiple exercise option problems under constraints. *Mathematical Methods of Operations Research*, 71(3):503–533, 2010.

[2] L. Andersen and M. Broadie. A Primal-Dual simulation algorithm for pricing multi-dimensional American options. *Management Science*, 50:1222–1234, 2004.

[3] K. J. Arrow and A. C. Fisher. Environmental preservation, uncertainty, and irreversibility. *The Quarterly Journal of Economics*, 88(2):312–319, 1974.

[4] S. Awerbuch and M. Berger. Energy security and diversity in the EU: A mean-variance portfolio approach. *IEA Research Paper*, 2003.

[5] V. Bally, G. Pagès, and J. Printems. A Quantization Tree method for pricing and hedging multi-dimensional American options. *Mathematical Finance*, 15(1):119–168, 2005.

[6] M. T. Barlow. A diffusion model for electricity prices. *Mathematical Finance*, 12(4):287–298, 2002.

[7] J. Barraquand and D. Martineau. Numerical valuation of high dimensional multivariate American securities. *Journal of Financial and Quantitative Analysis*, 30(03):383–405, 1995.

[8] D. Belomestny, G. Milstein, and J. Schoenmakers. Sensitivities for Bermudan options by regression methods. *Decisions in Economics and Finance*, 33(2):117–138, 2010.

[9] C. Bender. Dual pricing of multi-exercise options under volume constraints. *Finance and Stochastics*, 15(1):1–26, 2011.

[10] C. Bender, A. Kolodko, and J. Schoenmakers. Policy iteration for american options: overview. *Monte Carlo Methods and Applications*, 12(5):347, 2006.

[11] S. Boarin, G. Locatelli, M. Mancini, and M. E. Ricotti. Financial case studies on small-and medium-size modular reactors. *Nuclear technology*, 178(2):218–232, 2012.

[12] P. Boyle, M. Broadie, and P. Glasserman. Monte carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21:1267–1321, 1997.

[13] P. P. Boyle and Y. Tse. An algorithm for computing values of options on the maximum or minimum of several assets. *Journal of Financial and Quantitative Analysis*, 25(2):215–227, 1990.

[14] M. Brennan and E. Schwartz. Evaluating natural resource investments. *Journal of Business*, 58:135–157, 1985.

[15] M. Broadie and M. Cao. Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance*, 8(8):845–861, 2008.

[16] M. Broadie and J. Detemple. American option valuation: New bounds, approximations, and a comparison of existing methods. *Review of Financial Studies*, 9:1211–1250, 1996.

[17] M. Broadie and P. Glasserman. A Stochastic Mesh Method for pricing high-dimensional American options. *Journal of Computational Finance*, 7:35–72, 2004.

[18] L. Capriotti and M. Giles. Fast correlation Greeks by adjoint algorithmic differentiation. *arXiv preprint arXiv:1004.1855*, 2010.

[19] M. Carelli, P. Garrone, G. Locatelli, M. Mancini, C. Mycoff, P. Trucco, and M. Ricotti. Economic features of integral, modular, small-to-medium size reactors. *Progress in Nuclear Energy*, 52:403–414, 2010.

[20] J. F. Carriere. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: mathematics and Economics*, 19(1):19–30, 1996.

[21] N. Chiara, M. J. Garvin, and J. Vecer. Valuing simple multiple-exercise real options in infrastructure projects. *Journal of infrastructure systems*, 13(2):97–104, 2007.

[22] C. E. Clark. The greatest of a finite set of random variables. *Operations Research*, 9(2):145–162, 1961.

[23] E. Commission. The sustainable nuclear energy technology platform. *Special Report*, 2007.

[24] W. F. De Bondt and A. K. Makhija. Throwing good money after bad?: Nuclear power plant investment decisions and the relevance of sunk costs. *Journal of Economic Behavior & Organization*, 10(2):173–199, 1988.

[25] Department of Trade and Industry (DTI). *The future of nuclear power in a low carbon economy.* Consultation Document, United Kingdom, 2007.

[26] J. Deutch, E. Moniz, S. Ansolabehere, M. Driscoll, P. Gray, J. Holdren, P. Joskow, R. Lester, and N. Todreas. The future of nuclear power. *an MIT Interdisciplinary Study, http://web. mit. edu/nuclearpower*, 2003.

[27] A. K. Dixit. *Investment under uncertainty*. Princeton university press, 1994.

[28] Y. Du and J. Parson. Capacity factor risk at nuclear power plants. *Center for Energy and Environmental Policy Research*, 2010. Working Paper 10-016.

[29] F. Fang and C. W. Oosterlee. Pricing early-exercise and discrete barrier options by Fourier-cosine series expansions. *Numerische Mathematik*, 114(1):27–62, 2009.

[30] I. Fortin, S. Fuss, J. Hlouskova, N. Khabarov, M. Obersteiner, and J. Szolgayova. An integrated CVaR and real options approach to investments in the energy sector. Technical report, Institute for Advanced Studies, 2007.

[31] S. Fuss, J. Szolgayová, N. Khabarov, and M. Obersteiner. Renewables and climate change mitigation: Irreversible energy investment under uncertainty and portfolio effects. *Energy Policy*, 40(0):59 – 68, 2012.

[32] A. M. Gandrik. HTGR application economic model usersâĂŹ manual. Technical report, Idaho National Laboratory (INL), 2012.

[33] R. Gen IV. US DOE Nuclear energy research advisory committee and the generation IV international forum. A technology roadmap for generation IV nuclear energy systems. gif002-00, december 2002, 2002.

[34] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer, 2004.

[35] C. Gollier, D. Proult, F. Thais, and G. Walgenwitz. Choice of nuclear power investments under price uncertainty: valuing modularity. *Energy Economics*, 27(4):667–685, 2005.

[36] M. B. Haugh and L. Kogan. Pricing American options: A duality approach. *Operations Research*, 52(2):258–270, 2004.

[37] C. Henry. Investment decisions under uncertainty: The irreversibility effect. *American Economic Review*, 64:1006–1012, 1974.

[38] J. P. Holdren. The energy innovation imperative: Addressing oil dependence, climate change, and other 21st century energy challenges. *Innovations*, 1(2):3–23, 2006.

[39] IEA. World energy outlook 2011, 2011.

[40] S. Jain and C. W. Oosterlee. Pricing high-dimensional Bermudan options using the stochastic grid method. *International Journal of Computer Mathematics*, 89(9):1186–1211, 2012.

[41] S. Jain and C. W. Oosterlee. The Stochastic Grid Bundling Method: Efficient pricing of Bermudan options and the Greeks. *Available at SSRN 2293942*, 2012.

[42] S. Jain, F. Roelofs, and C. W. Oosterlee. Valuing modular nuclear power plants in finite time decision horizon. *Energy Economics*, 36:625–636, 2012.

[43] S. Jain, F. Roelofs, and C. W. Oosterlee. Construction strategies and lifetime uncertainties for nuclear projects: A real option analysis. *Nuclear Engineering and Design*, 265(0):319 – 329, 2013.

[44] J. Jansen, L. Beurskens, and X. Van Tilburg. Application of portfolio analysis to the Dutch generating mix. *Energy research Center at the Netherlands (ECN) report C-05-100*, 2006.

[45] X. Jin, H. H. Tan, and J. Sun. A State-Space Partitioning Method for pricing high-dimensional American-style options. *Mathematical Finance*, 17(3):399–426, 2007.

[46] K. H. F. Kan, G. Frank, V. Mozgin, and M. Reesor. Optimized least-squares monte carlo for measuring counterparty credit exposure of american-style options. *Mathematics-in-Industry Case Studies*, 2, 2010.

[47] R. Kaniel, S. Tompaidis, and A. Zemlianov. Efficient computation of hedging parameters for discretely exercisable options. *Operations research*, 56(4):811–826, 2008.

[48] S. Kendall, A. Stuart, and J. Keith Ord. *Kendall's Advanced Theory of Statistics. Vol. 1*. Charles Griffin & Compagny Limited, 1987.

[49] I. N. Kessides. Nuclear power: Understanding the economic risks and uncertainties. *Energy Policy*, 38(8):3849 – 3864, 2010.

[50] I. N. Kessides. The future of the nuclear industry reconsidered: Risks, uncertainties,and continued promise. *Energy Policy*, 48:185 – 208, 2012.

[51] S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[52] G. Locatelli, M. Mancini, and N. Todeschini. Generation IV nuclear reactors: Current status and future prospects. *Energy Policy*, 61(0):1503 – 1520, 2013.

[53] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial studies*, 14(1):113–147, 2001.

[54] r. Lord, F. Fang, F. Bervoets, and C. Oosterlee. A fast and accurate fft-based method for pricing early-exercise options under lévy processes. *SIAM Journal of Scientific Computing*, 30(4):1678–1705, 2008.

[55] J. J. Lucia and E. S. Schwartz. Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of Derivatives Research*, 5(1):5–50, 2002.

[56] H. Markowitz. Portfolio selection*. *The Journal of Finance*, 7(1):77–91, 1952.

[57] R. McDonald and D. Siegel. The value of waiting to invest. *Quarterly Journal of Economics*, 101:707–727, 1986.

[58] N. Meinshausen and B. M. Hambly. Monte carlo methods for the valuation of multiple-exercise options. *Mathematical Finance*, 14(4):557–583, 2004.

[59] R. C. Merton. An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, pages 867–887, 1973.

[60] W. Milne. On the degree of convergence of the gram-charlier series. *Transactions of the American Mathematical Society*, 31:422âĂŞ 443, 1929.

[61] C. Mycoff, M. Carelli, B. Petrovic, and K. Miller. Strategies to demonstrate competitiveness of SMRs in world markets. *IAEA*, 2010.

[62] NEA. *Reduction of Capital Costs of Nuclear Power Plants.* OECD/NEA, 2000.

[63] NEI. No reason why npps cannot live beyond 60. *Available at: www.neimagazine.com/story.asp?sectionCode=132&storyCode=2057693.*, 2010.

[64] OECD. Projected costs of generating electricity. Technical report, OECD, 2005.

[65] S. Pacala and R. Socolow. Stabilization wedges: Solving the climate problem for the next 50 years with current technologies. *Science*, 305(5686):968–972, 2004.

[66] D. Pilipovic. *Energy risk: Valuing and managing energy derivatives*, volume 300. McGraw-Hill New York, 1998.

[67] R. S. Pindyck. Investments of uncertain cost. *Journal of Financial Economics*, 34(1):53–76, 1993.

[68] N. S. Rasmussen. Control variates for Monte Carlo valuation of American options. *Journal of Computational Finance*, 9(1):83, 2005.

[69] F. Roelofs, J. Hart, and A. Van Heek. European new build and fuel cycles in the 21st century. *Nuclear Engineering and Design*, 241(6):2307–2317, 2011.

[70] F. Roelofs and U. von Estorff. Top down work force demand extrapolation from nuclear energy scenarios. *JRC Scientific and Policy Reports*, 2013.

[71] L. C. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.

[72] F. A. Roques, D. M. Newbery, and W. J. Nuttall. Fuel mix diversification incentives in liberalized electricity markets: A mean-variance portfolio theory approach. *Energy Economics*, 30(4):1831–1849, 2008.

[73] G. Rothwell. A real options approach to evaluating new nuclear power plants. *The Energy Journal*, 0(1):87–54, 2006.

[74] M. Ruijter and C. Oosterlee. Two-dimensional Fourier cosine series expansion method for pricing financial options. *SIAM Journal on Scientific Computing*, 34(5):B642–B671, 2012.

[75] J. Schoenmakers. A pure martingale dual for multiple stopping. *Finance and Stochastics*, 16(2):319–334, 2012.

[76] E. S. Schwartz. Patents and R&D as real options. *Economic Notes*, 33(1):23–54, 2004.

[77] J. Szolgayová, S. Fuss, N. Khabarov, and M. Obersteiner. A dynamic CVaR-portfolio approach using real options: an application to energy investments. *European Transactions on Electrical Power*, 21(6):1825–1841, 2011.

[78] R. Takashima and K. Yagi. Optimal investment timing and location under uncertainty. *13th International Conference on Real Options*, 2009.

[79] The Economist. When the wind blows. *Available at: http://www.economist.com/news/europe/21585029-hopes-fears-and-worries-europes-quest-renewable-energy-when-wind-blows*, 2013.

[80] S. Thomas. The economics of nuclear power: analysis of recent studies. *Available at: www.psiru.org/reports/2005-09-E-Nuclear.pdf*, 2005.

[81] J. A. Tilley. Valuing American options in a path simulation model. *Transactions of the Society of Actuaries*, 45(83):104, 1993.

[82] L. Trigeorgis. *Real Options: Managerial Flexibility and Strategy in Resource Allocation*. Cambridge MA: MIT Press, 1996.

[83] J. N. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *Neural Networks, IEEE Transactions on*, 12(4):694–703, 2001.

[84] L. Van Den Durpel, A. Yacout, D. Wade, and T. Taiwo. DANESS v4. 0: An integrated nuclear energy system assessment code. *PHYSOR*, 2008.

[85] A. van Heek, F. Roelofs, and A. Ehlert. Cost estimation with G4-Econs for generation IV reactor designs. In *GIF Symposium Proceedings*, page 29, 2012.

[86] J. Varley. *World Nuclear Industry Handbook*. Nuclear Engineering International., 2012.

[87] Y. Wang and R. Caflisch. Pricing and hedging American-style options: a simple simulation-based approach. *Journal of Computational Finance*, 13(4):95, 2010.

[88] World Nuclear Association. Plans for new reactors worldwide. *Available at: http://world-nuclear.org/info/Current-and-Future-Generation/Plans-For-New-Reactors-Worldwide*, 2013.

[89] L. Zhu. A simulation based real options approach for the investment evaluation of nuclear power. *Comput. Ind. Eng.*, 63(3):585–593, 2012.

[90] M. Zimmerman. Learning effects and the commercialization of new energy technologies: The case of nuclear power. *The Bell Journal of Economics*, 13(2):297–310, 1982.