Analysis of Employee Stock Options and Guaranteed Withdrawal Benefits for Life

by

Premal Shah
B. Tech. & M. Tech., Electrical Engineering (2003),
Indian Institute of Technology (IIT) Bombay, Mumbai

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2008

Author ....
Sloan School of Management
July 03, 2008

Certified by ....
Dimitris J. Bertsimas
Boeing Professor of Operations Research
Thesis Supervisor

Accepted by ....
Cynthia Barnhart
Professor
Co-director, Operations Research Center

Analysis of Employee Stock Options and Guaranteed Withdrawal Benefits for Life

by

Premal Shah

B. Tech. & M. Tech., Electrical Engineering (2003),

Indian Institute of Technology (IIT) Bombay, Mumbai

Submitted to the Sloan School of Management
on July 03, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Operations Research

Abstract

In this thesis we study three problems related to financial modeling.

First, we study the problem of pricing Employee Stock Options (ESOs) from the point of view of the issuing company. Since an employee cannot trade or effectively hedge ESOs, she exercises them to maximize a subjective criterion of value. Modeling this exercise behavior is key to pricing ESOs. We argue that ESO exercises should not be modeled on a one by one basis, as is commonly done, but at a portfolio level because exercises related to different ESOs that an employee holds would be coupled. Using utility based models we also show that such coupled exercise behavior leads to lower average ESO costs for the commonly used utility functions such as power and exponential utilities. Unfortunately, utility based models do not lead to tractable solutions for finding costs associated with ESOs. We propose a new risk management based approach to model exercise behavior based on mean-variance portfolio maximization. The resulting exercise behavior is both intuitive and leads to a computationally tractable model for finding ESO exercises and pricing ESOs as a portfolio. We also study a special variant of this risk-management based exercise model, which leads to a decoupling of the ESO exercises and then obtain analytical bounds on the implied cost of an ESO for the employer in this case.

Next, we study Guaranteed Withdrawal Benefits (GWB) for life, a recent and popular product that many insurance companies have offered for retirement planning. The GWB feature promises to the investor increasing withdrawals over her lifetime and is an exotic option that bears financial and mortality related risks for the insurance company. We first analyze a continuous time version of this product in a Black Scholes economy with simplifying assumptions on population mortality and obtain an analytical solution for the product value. This simple analysis reveals the high sensitivity the product bears to several risk factors. We then further investigate the pricing of GWB in a more realistic setting using different asset pricing models, including those that allow the interest rates and the volatility of returns to be stochastic. Our analysis reveals that 1) GWB has insufficient price discrimination and is susceptible to adverse selection and 2) valuations can vary substantially depending on which class of models is used for accounting. We believe that the ambiguity in value and the presence of significant risks, which can be challenging to hedge, should create concerns to the GWB underwriters, their clients as well as the regulators.

Finally, many problems in finance are Sequential Decision Problems (SDPs) under uncertainty. We find that SDP formulations using commonly used financial metrics or acceptability criteria can lead to dynamically inconsistent strategies. We study the link between objective functions used in SDPs, dynamic consistency and dynamic programming. We then propose ways to create dynamically consistent formulations.

Thesis Supervisor: Dimitris J. Bertsimas
Title: Boeing Professor of Operations Research

Acknowledgments

First, I would like to express my sincere gratitude towards *Prof. Dimitris Bertsimas* for his guidance and support in the work related to this thesis and beyond. From day one of my association with him, Prof. Bertsimas was never short of interesting research ideas and insights into problems. His emphasis on the practicability of research has kept me true. Without him this work would have been but an exercise in futility.

Dimitris has been a great source of inspiration, encouragement and counsel to me outside research and academics as well. His versatility in terms of dealing with a broad range of research areas in both theory and applications, his repertoire of interests and activities and his drive are stupendous. I consider myself very fortunate to have been able to tap into the same, especially given the several sources competing for his time!

I would also like to acknowledge the members of my thesis committee - *Prof. Jiang Wang,* for valuable inputs and suggestions to improve this thesis and *Prof. David Gamarnik,* for his encouragement and interest. Discussions with David about any problem have always refined my understanding significantly. I am also thankful to David, Dimitris and *Prof. John Tsitsiklis* for giving me the opportunity to serve as a teaching assistant with them. Teaching a course with one of them at charge has been a surprisingly good learning experience as well!

I would also like to express my gratitude towards the Operations Research Center administrators - *Paulette Mosley, Laura Rose and Andrew Carvalho* - their efficiency and co-operation have made things so much easier.

The years at ORC were sweet and memorable thanks to the fellow ORCers - *Ilan Lobel, Yunqing (Kelly) Ye, Vinh Doan, Ruben Lobel, Doug Fearing, Dan Iancu, Carine Simon, Gareth, Juliane, Dung, Mike, Rajiv, Pranava, Shobhit, Carol, Pavithra, Theo, Andy, Yann, Elodie, Katy, Rags, Pallav, Dmitriy, Dave, Nikos, Margret, Lincoln, Karima, Mallory, Shubham, Thibault, Martin, Shashi, Sun Wei, Kostas, Nelson, ...* . Thank you for the countless wonderful moments inside and outside the ORC. You all will be sorely missed!

Thanks is also due to all the room-mates and their extensions - *Sudhir Chiluveru, Saurabh Tejwani, Hoda, Jean-Paul, Alec, Kurt, Raf, Mukul, Riccardo, Anatoly* who made leaving in the dorms so much more homely.

I am also grateful to MIT for the opportunity to meet and make so many wonderful friends - *Rajiv Menjoge, Aditya Undurti, Theta Aye, Vikas Sharma, Kranthi*

Vistakula, Neha Soni, Mithila Azad, Mira Chokshi. Thanks to them and the rest of the kajra-re and wandering gopis gangs - *Jen Fang, Sukant, Preet, Tulika, Silpa, Natasha, Tara(2!), Nick, Silpa, Saba(2!), Kartik, Chief, ...* for all the fun times!

The time spent here, in the cold of Boston, also helped me to thaw some old friendships - *Nandan Bhat, Ajay Raghavan, Suchit Kaul and Prafulla Krishna* - it was great to have rediscovered you all!

I am also thankful to Sangam and the '05 executive council fellow members - *Vivek Jaiswal, Saurabh, Biraja* for helping to re-create here Diwali, Holi and all the other things that I missed from home. I would also like to thank the members of the Edgerton Board '07 - *Yvonne, Nicholas, Kelly, Biz, Grace, Becky, Ron,...* and the Edgerton House Masters - *David and Pamela Mindell;* working with you all was a great experience.

The last several months at MIT have been different and not just because of this thesis. Thanks to *Tara Sainath* for helping with proof-reading this thesis and a million other things and all the care and affection I could have asked for.

Finally and most importantly, I would like to thank my parents, my sister *Payal* and my relatives in India - their love, support and care have been a source of great strength for me. Nothing would have been possible without this foundation of my life.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Much of quantitative finance is based on models of market variates - how prices of securities traded in the markets evolve, how market participants behave and the interaction between the two.

In a landmark paper, Black and Scholes [18] proposed a model to price a stock-option based on the price of the underlying stock. They provided not only a way to unambiguously price an option, but also a method to "hedge" out and in theory, eliminate the risk of holding or underwriting this option. This spawned an entire new field of financial engineering and a fresh body of work based on the concept of risk-neutral pricing and hedging was created and is still being pursued very actively. Simultaneously, in industry, a plethora of complex financial derivatives are being created, marketed and sold to institutions and households.

The assumptions underlying the risk neutral pricing theory are that the markets are complete and that the market players do not face constraints in buying and selling the various instruments traded in the market. In this ideal setting, all derivative securities are redundant and can not only be priced unambiguously but also hedged perfectly from the prices of other traded instruments. The quantitative link between various traded instruments in the market is established by using a parametric model for price process of the underlying(s) on which the derivative instruments are written. This method works fairly well for pricing standard instruments that are heavily traded. But as the derivative instrument to be priced becomes more intricate in its dependence on the underlying(s) and the market for it more constrained, the financial engineering methodology becomes less precise and more subjective. The model that is used to link together prices of different securities then itself becomes a key factor in valuation. Due care must then be taken in creating models to price such an instrument as well as interpreting results from the so called no-arbitrage pricing models.

In this thesis, we examine two problems in pricing derivatives in "incomplete" markets and show how even reasonable models can sometimes leave out significant determinants of value.

Models for behavior of market participants such as investment managers or consumers are in general decision and control problems. Again here, we observe in a dynamic setting, an ad-hoc formulation of an investment manager's problem that employs the commonly used financial metrics can lead to a model where the manager takes conflicting decisions over time. We also examine in this thesis, general properties of a dynamic decision framework for "consistency" in decisions.

1.2 Contributions

Modeling Exercise Behavior for pricing Employee Stock Options

We consider the problem of pricing Employee Stock Options (ESOs) from the point of view of the companies that issue them. Although, off late, ESOs have been losing popularity due to scandals and less favorable accounting regulations, they still constitute a sizeable chunk of many companies' compensation costs. Since employees are constrained in both trading as well hedging ESOs, the standard risk-neutral pricing framework cannot be directly applied to ESOs. Employees, would exercise ESOs to maximize a personal measure of value or utility. Pricing ESOs then encompasses twin problems - we first need to model employee exercises, which can be sub-optimal from a risk neutral perspective, while being subjectively optimal; and then compute the risk-neutral cost of the ESOs under these exercise policies. The choice of model for exercise behavior, will have a big impact on the valuation of ESO costs.

Most ESO exercises are driven by the need of a risk-averse employee to limit the uncertainty of an option payoff that she cannot hedge. The basic financial tenet of diversification would suggest that as all ESOs are exposed to common risk sources, they have diminishing marginal value to the employee. It is then plausible that incremental option grants will, in general, be exercised differently. This, in turn, would lead to a different incremental cost for an ESO grant to the employer. The models that have been proposed in literature, typically price ESOs on an individual basis and in isolation and would not take this effect into consideration. We propose a new approach - to explicitly take into account the employee's need for diversification with respect to the entire portfolio of ESOs that she holds while modeling her exercise behavior. Thus we propose to price ESOs at a portfolio level rather than at individual level.

We first augment the conventional utility based framework that leads to an endogenous exercise model for ESOs and show that in general bundling together of ESOs affects exercise behavior and tends to cause an employee to exercise her ESOs earlier on average. This makes her forgo a larger part of the option value of the portfolio, thereby reducing its cost for the employer. Also, an immediate consequence of taking these portfolio effects into account is that the cost of an option portfolio is no longer linear or equal to the sum of its parts. Further, issuance of new options can have a retrospective effect on the cost of already issued options.

We then use the concept of risk management and ideas in portfolio optimization to motivate a model, where the employee exercises options so as to optimize a risk-adjusted value of the entire portfolio at each time step. This causes the employee to exercise options in decreasing order of a barrier function that can be interpreted as a pseudo Sharpe-ratio for the option. The advantage of this model is that it leads to a computationally simple framework to both price the ESO portfolio and also allocate its costs amongst its components. For a special risk-management based exercise model we show that option exercises decouple and one can think of applying the "pseudo Sharpe-ratio" criterion to options on a one-by-one basis. In this case, we recover a linear pricing rule for ESOs and also derive tight analytical bounds on the cost of an ESO.

Pricing Guaranteed Withdrawal Benefits for Life

Complex financial derivatives are often embedded in retail investment products. We consider one such recent and extremely popular innovation in the private pension product space - the Guaranteed Withdrawal Benefits (GWB) for life. The GWB for life option, which is usually available as an add-on to a Variable Annuity (VA) investment fund, guarantees an investor a non-decreasing stream of payments in her retirement until death, with her funds always staying invested in the VA. In return, the investor pays a small fee indexed to the quantum of the guarantee, every year. While prima-facie, GWB for life appears to be just another, somewhat exotic, financial option, pricing it poses many challenges. Due to its exposure to population longevities and dependence on investor behavior over time and complex dependence on various financial market factors, the complete markets hypothesis does not hold for the GWB.

We first undertake an analysis of GWB for life in a simplified setting using the Black-Scholes model for asset prices. This allows us to get an almost closed form solution for the value of GWB. We use this solution to draw insights and investigate the impact of potential risk factors and find almost all of them to be quite significant.

We then price GWB in a more realistic setting using models that allow interest rates and equity-market volatilities to be stochastic. We find that accounting for these additional risks can alter valuations significantly. In addition, GWB has considerable exposure to realized investor population longevity. These facts suggest that hedging GWB is likely to be only partially successful in practice. We also find that the typical GWB for life offering with its uniform pricing across fund classes and investor ages is susceptible to adverse selection in its customer profile and needs price discrimination.

Dynamic Consistency For Sequential Decision Problems

Portfolio Optimization and Risk Management are standard problems in finance. In a dynamic setting, these can be viewed as instances of a much broader class of problems - Sequential Decision Problem (SDP)s. We study SDPs in context of an important normative criterion for a good SDP model - that it should lead to dynamically consistent planning. A lack of dynamic consistency would mean that the decision maker would make plans, while being fully aware that she will not carry them through. While SDPs arise in several decision and control settings, we show that tendency to be dynamically inconsistent is particularly severe for financial applications. This is because SDPs based on many of the performance metrics such as Sharpe-ratio, variance adjusted mean, Value at Risk etc as well as acceptability criteria based on so called dynamic risk measures turn out to be dynamically inconsistent.

We explore the connection between dynamic consistency and the algorithmic notion of dynamic programming or Bellman's principle and find them to be closely related, though not identical. We show that most dynamically consistent strategies can be considered to be arising from SDPs that have a sum decomposable representation across time and event space. We then examine how "conflicts" due to dynamic inconsistency can be resolved for specific applications. We also propose a new class of dynamically consistent performance metrics that are essentially expectations with respect to probability measures distorted in a specific way.

1.3 Thesis Overview

The thesis is organized as follows. In Chapter 2, we introduce the problem of ESO pricing for the company that issues them. We examine how the presence of other ESOs in an employee's portfolio can affect the exercise decisions concerning ESOs and thereby their cost to the company. In Chapter 3, we then propose risk management based models for exercise behavior, and through them, a tractable way to price ESO portfolios. In Chapter 4, we then turn to the problem of pricing GWB for life. In this chapter, we propose an analytical framework to price GWB for life for a continuous time counterpart of the product. In Chapter 5, we price GWB for life using realistic models and investigate the impact of modeling interest rates and volatilities as stochastic processes on pricing. In Chapter 6, we look into the issue of dynamic consistency for SDPs, especially in the context of some standard problems in finance. We summarize the findings and some interesting research directions in Chapter 7.

Chapter 2

Pricing Employee Stock Options

2.1 Introduction

Employee Stock Options (ESOs) are commonly used by corporations as an effective, but often controversial, form of compensation for mid and high level employees. An ESO is typically an American Call Option that the employee can exercise between two pre-specified dates, the earlier date called the *Vesting Date* and the later the *Expiration Date*. By its very structure, an ESO acts like a performance - linked compensation for higher echelon executives. Moreover, the employing company would realize the cost of this pay only in the event of its stock performing well.

The popularity of ESOs as a means of compensation for employers can be attributed to three important reasons:

- They directly link pay to performance and serve to align, at least partially, the management's interests with those of the shareholders, thereby mitigating some of the "agency problems"[1].

- Long-term options create an incentive for the employee to stay with the company and thus options with vesting schedules can be used to retain talent.

- Most ESOs are at the money (ATM) call options. Until recently, under Financial Accounting Standards Board (FASB)'s alternate accounting provisions, companies could expense stock option grants to employees at their intrinsic value, i.e., zero costs for the ATM options. Thus this form of compensation,

[1]Agency problems arise because in general the goals and objectives of a company's management do not coincide with those of its share-holders.

at grant, would not add to company's expenses or show up in its profit and loss/earnings statements.

On the negative side, ESOs issued to executives can also create a conflict of interest between management and shareholders, as the asymmetric option payoff can incentivize managers to undertake projects with unduly high risk. More important, perhaps, is the fact that ESOs actually amount to a significant liability on a company's balance-sheet that often goes under-expensed. For example, when ESOs issued by technology companies were exercised by the employees during the dot-com boom, there were payoffs amounting to tens and sometimes hundreds of millions of dollars. These were effectively a transfer of value from the shareholders to the employees, mainly executives.

2.1.1 Motivation

Empirical surveys show high-level executives receive a bulk of their compensation as stock options, a trend which has recently started to somewhat reverse because of scandals and controversies. For example, according to the data compiled by Henderson [67], in 2002, 58% of the net CEO pay in the US and 24% in the UK was options related. In terms of balance-sheet liabilities, Hall and Murphy, [63] report that in 1992, firms in the Standard & Poor's 500 granted their employees options worth a total of 1.1 billion at the time of grant. This figure reached 119 billion in 2000 before dropping down to 71 Billion in 2002, still a sizeable figure. Because ESO related costs can amount to a substantial fraction of the firms's balance-sheet, evaluating this cost is important for investors and regulators.

Pricing ESOs is however made difficult because of the fact that they are not tradeable and hence do not have a directly observable market price. We also cannot price these options using standard models such as the Black-Scholes framework [18], because the option bearer faces constraints that would not allow her to hedge these options effectively[2]. This coupled with risk-averseness causes a typical employee to exercise an ESO in a way that would substantially reduce its cost below its Black-Scholes or risk-neutral price.

More specifically, the employing company realizes the ESO cost if and when the employee exercises the option. Unlike a regular call option which typically gets ex-

[2]To hedge a call option, the employee should be able to short the employing company's stock. This is usually prohibited by regulatory bodies. Also if the employee were able to hedge out the options, then many of the objectives of issuing the options as a means of incentive and retentive compensation will be lost.

ercised at or close to expiry[3], an ESO is typically exercised much earlier, see Hall and Murphy [63]. ESOs follow subjective exercise patterns that are difficult to predict because the employee exercise ESOs to realize a measure of personal value. The clauses related to vesting and forfeiture of options (in the event of employee quitting or being terminated) further complicate the problem of pricing ESOs.

Unfortunately, but perhaps also unsurprisingly, there is no consensus in literature or in practice about what the fair cost of ESOs is. After some accounting controversies and several debates, the FASB issued a revision to Statement 123 that deals with accounting principles for stock related compensation in 2004 [54]. This made it mandatory for publicly traded corporations to expense ESOs at their "fair value" (or levels more representative of the cost incurred than the intrinsic value accounting), effective 2005. The European counterpart, International Accounting Standards Board (IASB), had laid down similar stipulations earlier through [74]. However FASB, IASB and other regulatory bodies have only laid down broad guidelines when it comes to methods and models to estimate a "fair cost" of ESOs, going only so far as indicating preferences for some models - such as the lattice model. For example, the Securities and Exchange Commission (SEC) Bulletin guidelines [108], state that the accounting practice used to price an ESO must be based on sound financial economic theory and be generally accepted in the field but stops short of laying down a specific accounting rule, see Cvitanic and Zapatero [47]. While it is broadly agreed that the true cost of an ESO lies between its intrinsic value and the Black-Scholes value, there continues to be an active debate about what exactly the "fair cost" of an ESO is.

In this chapter, and Chapter 3, we seek to develop a framework to model ESO exercises so as to estimate the cost of the outstanding ESOs on a company's balance-sheet. Our focus in this chapter will be to understand the functional nature of ESO costs. In Chapter 3, we seek practical methods to expense ESOs that are driven by economic reasoning and at the same time are simple to implement in practice.

While not considered in this thesis, an interesting problem related to costing ESOs is estimating the value of ESOs to the employee. This has been studied well e.g. Lambert, Larcker and Verrechia in [83], Ingersoll in [73]. Understanding how an employee would value the options grant is useful in designing compensation packages to incentivize desired management or employee actions. How an employee would value an option grant though will not be the same as the cost it represents to the company. In this context, it is worth mentioning a generally accepted conclusion that

[3] It is well-known that an American Call Option on the stock of a company that does not pay dividends is optimally exercised at expiry.

the option's value to the employee is less than the cost of issue to the employer as has been discussed by Hall and Murphy in [62] and [63]. The difference, a "deadweight loss", may also be seen as a price that the company pays to solve its agency problems and retain talent, see Kadam [77].

2.1.2 Related Work

ESOs have been an active topic of research and debate and there is a vast body of work that deals with pricing of ESOs and other issues related to them. As the topic touches so many fields, the contributions also come from diverse areas including Accounting, Econometrics, Asset Pricing and Mathematical Finance. Chance [32] provides a detailed analysis of the issues related to ESOs from many different perspectives as well as a sound critique of the approaches that have been proposed to address these issues in the literature. Hall and Murphy [63] analyze historical trends in issuance of ESOs by corporations to executives and lower level employees and the possible attractions and pitfalls of using them as incentive - pay. Another paper by the same authors, [62], provides a good understanding of the role of ESOs in incentivizing executives and how risk-averseness and other idiosyncratic investor characteristics might affect exercise behavior using a stylized model. The authors propose a simple utility based framework and give numerical examples to illustrate the effects of risk-averseness and trading restrictions on employee's exercise behavior and cost of ESOs to the issuers.

Huddart and Lang [71] present an empirical analysis of how employees tend to exercise their ESOs using over 10 years of data. Bettis, Bizjak and Lemmon [15] provide an analysis of exercise behavior and incentive effects of ESOs using an empirically calibrated utility model.

The employee's decision-process remains fundamental to pricing an ESO[4]. Hence, even though we do not seek to value ESOs from an employee's perspective, we still need to have a model of the employee's exercise behavior.

In general, for modeling exercise behavior, two broad approaches have been used, as observed in Carr and Linetsky [31]. The first approach treats ESO exercise as an endogenous process and models it as a decision triggering from typically a utility optimization consideration. There are several factors that can potentially influence

[4]However, approaches to pricing, that circumvent this have also been proposed. For example, Bulow and Shoven [27] suggest an alternate way of accounting for ESO costs, in which the ESOs are expensed as rolling options of quarterly maturity until they are exercised or lapse. In another strikingly different approach to pricing, Core and Guay [43] suggest an empirically calibrated model that uses data available in a company's proxy statement to price ESOs.

exercise decisions:

- Employee's Risk Averseness - Employees are typically over-exposed to their employers. The option-bearer might want to offload some of this by exposure by cashing the ESOs.
- Tax Implications.
- Career related moves can cause early ESO exercises or forfeitures.
- Liquidity crunches may force an early exercise of the option.
- ESO terms and company policies - Companies tend to reset strikes of options in the event the stock price goes significantly below the strike and also issue new ESO grants (termed as "reload") on exercises. Sometimes these features are explicitly embedded in the ESOs and can impact exercise related decisions.

To retain tractability, endogenous exercise models often retain focus on one or few of the several possible factors that can influence exercise decisions. For example, for Constant Relative Risk Aversion (CRRA) utilities, Ingersoll [73] derives the subjective and objective value of ESOs when the employee is constrained to hold a certain fixed proportion of her wealth in the stock of the employing company. Detemple and Sundaresan [49] analyze the value of a non-tradeable option using dynamic programming methods under a binomial stock price model for CRRA utilities - that is directly applicable to pricing an ESO. Using a simple 2-period binomial model Kulatilaka and Marcus [82] show how liquidity constraints and other idiosyncratic factors related to an employee can influence the exercise behavior. The authors remark that FASB recommended methods miss out on or inaccurately estimate the effect of such factors.

The second approach is to model exercise behavior as an exogenous process. A justification for this is provided by Carpenter [29] who showed that empirically calibrated utility based models do a no better job of predicting exercises when compared to a model that uses random exogenous exercises and forfeits. This motivated researchers to look at intensity based models where exercise is modeled as an independent random process. An example is the model in Carr and Linetsy [31], where exercise occurs as an arrival in a Poisson process whose intensity is albeit modulated by the stock price. This model gives an analytical expression for the ESO cost. In a similar spirit, Hull and White [72], while pointing out drawbacks of the methods proposed by FASB to expense options, suggest an "Enhanced FASB 123" method to expense options. Their approach is to use a binomial stock price process and an employee behavior model in which exercise is triggered whenever the stock price hits a certain multiple of the strike. Cvitanic and Zapatero [47] use a similar framework, but in continuous time, which is solvable analytically. They also employ a fictitious barrier based exercise policy for the employee, in this case the employee would exercise her option when the stock price hits a barrier that decreases exponentially with time, and also allow for exercises due to employee exiting the company.

Sircar and Xiong [109] propose an elaborate framework that takes into account reload (wherein exercise of options leads to a new grant) and reset (underwater options have their strikes reset) features of options, and gives analytical formulae for the option price under the assumptions of no expiry and no hedging constraints. Dybvig and Lowenstein in [52], Hemmer, Matsunaga and Shelvin in [66] and Acharya, John and Sundaram in [3] also consider the impact of reload features on option prices. Bodie, Ruffino and Treussard [21] propose a broad framework that an employee can use to weigh ESO benefits while making career related decisions.

2.1.3 Contributions

Most of the proposed ESO costing methods implicitly assume that employees exercise ESOs in an "all or none" fashion. While this assumption is appropriate for traded options, the possibility of partial exercise must be considered while valuing ESOs, as it is reasonable to expect employees to exercise options in batches to distribute the risk over time. Notable exceptions that consider the effect of partial exercises of options are Jain and Subramnaian [75] and Grasselli [61] which provide an analysis of how allowing for partial exercises can impact values and cost of ESOs using primarily a two-period binomial model. Grasselli [61] also considers the possibility of partial hedging using correlated instruments on option prices. Recent work by Leung and Sircar [86] and Rogers and Scheinkman [104] have also considered these effects and solved for optimal exercises numerically using a utility based framework.

What we propose here is to take this reasoning a step further. ESOs are granted in lots and batches and most employees at any given point of time will, in fact, have a basket of unexercised ESOs with varying strikes, expiries and vesting dates. Since, most researchers agree that risk-averseness and over-exposure to the employing company's fortunes drives early exercise of the ESOs, the degree of this exposure, accumulated primarily through the employee's own ESO portfolio, will weigh on her decision to exercise an option. Moreover, unlike many of the other quantities alluded to in ESO pricing models, unexercised ESO grants to an employee constitute information that should be readily available to the company and hence easy to use. We therefore argue that unlike the usual approach taken in literature so far to price ESOs, and unlike FASB recommended methods, ESOs perhaps need to be priced not one by one but as an entire portfolio of options held by a particular employee. Thus, ESO exercises can not only be "partial' but also "coupled". This would also mean that ESO costs in general will not be linear. For example, the cost of a lot of ESOs does not increase linearly with its size. More generally speaking, the cost of an ESO portfolio will not be the same as the sum of its parts.

In this chapter, we examine the case for portfolio pricing of options by trying to study the qualitative implications of a portfolio approach on ESO costing. We use a standard expected utility maximizing framework as a basis for the employee's decision process. We begin with the case where employees have several ESOs with the same terms. We find that even in this simple case, exercise policies and hence the implied cost of the grant for the employing company can, in general, vary arbitrarily depending on the nature of employee's utility function and the stock price dynamics. However, for commonly used utility functions such as the class of Constant Relative Risk Aversion (CRRA) and Constant Absolute Risk Aversion (CARA) utilities, for any stock price dynamics diversification needs would cause the employee to exercise a proportionally larger component of her grant earlier. As a result, the cost of the portfolio in these cases increases sub-linearly with the grant size for options with similar terms. We find this effect interesting because it suggests that ESOs not only have diminishing marginal utility for employees, as one would expect, but in some sense also diminishing marginal costs (or more generally, diminishing average costs) for employers issuing them. We then seek to extrapolate these findings to the case where the employee has multiple types of ESOs in their portfolio. Surprisingly, even for CARA utilities, cost of an option portfolio can turn out to be super-additive, i.e., more than the sum of it parts for some stock price dynamics. However, with an additional but reasonable assumption on stock price dynamics that can be linked to diversification, the cost of the portfolio with multiple types of ESOs can be shown to be sub-additive or less than the sum of its parts. Our analysis thus establishes that a one-by-one costing of ESOs is likely over-estimating the cost of ESOs.

2.1.4 Chapter Layout

In Section 2.2, we briefly describe the model used in this chapter. In Section 2.3, we discuss the simple two period case and conditions on utility functions that will make the average cost of an option grant decreasing in the size of the grant. Next, in Section 2.4, we consider the multi-period case for a single option type and show that the average cost under both CARA and CRRA utilities will be decreasing for arbitrary stock price dynamics. In Section 2.5, we then consider the case of ESO portfolios with multiple option types and show that there exists a partial order of exercise between options of different types. In Section 2.6, we examine in detail the problem of expensing multi-type multi-period option portfolios, particularly for CARA utilities. Section 2.7 provides a summary of the results obtained.

2.2 Model

We work with a discrete time model, where the employee treats her ESO grants as investment rather than consumption instruments. Further, the employee will never require to exercise her ESOs for liquidity reasons. We do not consider the impact of the employee quitting or being fired on option exercises or costs. We also do not take into account the "reset" and "reload" tendencies/policies that companies sometimes have in more exotic ESO grants as we would like to retain focus on the key goal of this chapter i.e., the impact of a portfolio approach on exercise behavior and the implied option cost.

The employee has a concave utility function $U(\cdot)$ and a planning horizon T and seeks to maximize the expected utility of her wealth position W_T at T.

In our model, the employee may have N different type of options in her portfolio $\mathbf{P}$. The type i option is characterized by a strike price, denoted by K_i, an expiry T_i and a vesting date V_i. Also, the number of unexercised options of type i in $\mathbf{P}$ at time t are denoted by $\alpha_{i,t}$, with $\alpha_i \triangleq \alpha_{i,0}$. Restricted stock grants can be treated within this framework as options with strike 0. We also impose the natural restrictions that the employee can neither trade nor hedge against these options. We allow for partial exercises and for mathematical convenience "fractional exercises" to avoid integrality constraints.

We assume that all non-ESO wealth is invested in a well diversified portfolio, whose returns are independent of the employing company's stock. Any proceeds from option exercise are also likewise invested. An assumption that helps us simplify the analysis considerably is that the employee continues to measures her wealth at time T in units of wealth indexed to some reference time, by discounting time t cash-flows a subjective discounting factor β_t. β_t can be interpreted as an "opportunity cost of cash" for the employee[5].

This approach is in similar spirit but slightly more general than the one used in several utility based models for exercising ESOs, notably Kadam [77], Huddart and Lang [70], Kulatilaka and Marcus [82] and more recently, the model discussed in Rogers and Scheinkman [104]. In principle, upon exercise, the employee has the freedom to invest the proceeds along with other non-option wealth in the markets, and hence must jointly solve the problem of investing non-option wealth and exercising ESOs. Such an approach has been taken for example in Leung and Sircar [86], Grasselli [61]. Allowing for this additional flexibility usually requires making some other restrictive assumptions in order to maintain analytical tractability. Both Leung and Sircar [86] and Grasselli [61] assume that the employee has Constant Absolute Risk Aversion (CARA) type utilities and also assume simple dynamics for stock price processes. By using a model where the employee always uses her wealth at time 0 as a numeraire and subjectively discounts future cash-flows, we can decouple the exercise decision from the non-option wealth investment decision and simplify the analysis considerably. We, in fact, assume no particular form of dynamics on the stock price process except that it is a Markov process, to keep notation less cumbersome.

Under the model described above then, the employee's decision problem can be described by the following optimization problem. β_t denotes the time t discount factor for the employee, based on her opportunity cost of cash and W, the current non-option wealth in the reference time units.

$$\begin{aligned}
\max V &= \mathbb{E}[U(W_T)] \ ; \\
\text{s.t. } W_T &= W + \sum_{i=1}^{N} \sum_{t=s}^{T} x_{i,t} \beta_t (S_t - K_i)^+ \ , \\
& x_{i,t} \text{ is } \mathcal{F}_t - \text{measurable. } , \\
& \sum_{t=0}^{T} x_{i,t} \ = \ \alpha_i \ldots 1 \leq i \leq N \ , \\
& x_{i,t} \ = \ 0 \text{ if } t < V_i \text{ or } t > T_i \ .
\end{aligned} \tag{2.1}$$

The exercise problem in (2.1) is a dynamic programming problem.

The quantity $x_{i,t}$ denotes the number of options of type i to be exercised at time t. The expectations in (2.1) are with respect to the employee's belief about the

[5]The analysis presented carries through even when the subjective discount factor β_t is taken to be stochastic. We only require that β_t is almost surely decreasing with time. For simplicity, we will however consider β_t to be non-stochastic.

employing company's stock price process. We will assume that there exists a unique preferred solution to the problem in (2.1). To make this precise, if there are multiple solutions to the problem, then the employee chooses for implementation at any given time an exercise policy that has the smaller value of the ordered set $\{x_{i,t} : 1 \leq i \leq N\}$, when comparisons are made in the lexicographic order. We denote the exercise policy so obtained by $\mathbf{x}^*$ and the optimal exercise for option i at time t by $x^*_{i,t}$.

We assume, there is also a unique risk-neutral measure, $\mathbb{Q}$, that prices securities, and is absolutely continuous with respect to the real or believed stock price process. Let $D_t = \exp(-\int_0^t r_s ds)$ denote the time t discount factor based on risk-free interest rates r_s. Then the cost of the grant to the employer is given by:

$$C(P) \;=\; \mathbb{E}^{\mathbb{Q}}\left[\sum_{n=1}^{N}\sum_{t=0}^{T} x^*_{i,t} D_t (S_t - K_i)^+\right] . \tag{2.2}$$

We also assume that the stock does not pay dividends. The quantity $D_t(S_0 - K)^+$ is a $\mathbb{Q}$ sub-martingale by Jensen's inequality.

Most people, typically expect stocks to appreciate on average over time, i.e., believe the stock price process to be a sub-martingale. By Jensen's inequality, the expected value of an option's payoff which is a convex function of the stock price, is increasing in the time of exercise. However, the employee's utility function is concave and hence delaying exercise need not increase the expected utility of the payoffs received for the employee. Thus, there is a trade-off between exercising immediately and waiting, and the exercise decision will be impacted by several factors including the nature of utility function, the time to expiry, assumed dynamics of stock prices, the level of current non-option wealth and the number of unexercised options.

Our goal in this chapter is to analyze qualitatively the nature of an employee's exercise policy, as governed by (2.1) and thereby get some comparative statics on the cost of issuance to the employer, as given by (2.2). We start with a simple case, where $N = 1$, i.e., the employee has several options of a single *type* in the portfolio.

2.3 Nature of ESO Cost Functions

In this section, we show that even for concave utility functions, the cost of issuing ESOs can become convex.

Let us take the simplest instance of (2.1), with the number of types $N = 1$ and expiry time $T = 2$.

We will assume that the options in question have already vested. The corresponding exercise decision problem is then,

$$\begin{aligned} \max V &= \mathbb{E}[U(W + x\beta_0(S_0 - K)^+ + \beta_1(\alpha - x)(S_1 - K)^+)] \ ; \\ \text{s.t.} & \quad 0 \le x \le \alpha \ . \end{aligned} \tag{2.3}$$

As there is only one type of option involved, we have dropped the sub-scripts referencing the option type. The quantity x denotes the only decision variable in this problem - the number of options to be exercised at time 0. If the optimal value of this decision variable is x^*, then the implied cost to the employing company is given by

$$C(\alpha, W_0) = x^*(S_0 - K)^+ + D_1(S_1 - K)^+ \ .$$

We find that, even for this simple case, the employee exercise policy or the implied cost for the employers cannot be generalized. The value of the option grant to the employee, as measured in terms of her optimal attainable utility will always be concave in the option grant, so long as the employees utility function is concave. As a function of the grant-size, the cost can however become super-linear.

We now give a simple example where the cost of the option grant becomes convex.

Example 2.1. *The employee has a utility function*

$$U(y) = \min\left(y, \frac{1}{3}y + 40\right)$$

and an initial wealth $W = 50$. The employee's utility function is concave (though not differentiable). Assume the employee also has α ESOs with strike 90 expiring in one period, i.e., $T = 1$. To keep things simple, we set both the opportunity cost of cash as well as risk-free interest rates to zero. Suppose the current stock price $S_0 = 100$ and the employee believes that at $T = 1$, S_1 will be either 120 or 80, with equal probability. The probability of these movements in the risk neutral measure can then also be verified to be $\frac{1}{2}$. This implies that the fair value of an American call option expiring at $T = 1$ is 15. It is relatively straightforward to verify that the optimal no.

of options, x_0^*, that the employee should exercise at time 0 *is given by*

$$x_0^* \;=\; \min(1, \alpha).$$

Thus the resulting cost $C(\alpha, W)$ to the company is:

$$C(\alpha, W) \;=\; 10 \cdot \min(1, \alpha) + 15 \cdot (\alpha - 1)^+ \tag{2.4}$$

(2.4) shows that the cost of this grant to the employing company is convex in the size of the grant in this case.

In Lemma 2.1 that follows, we characterize a class of utility functions for which the cost function becomes sub-linear in grant-size. This shows that mere concavity of the utility function, which is sufficient to infer diminishing marginal value to the employee, is however not sufficient to draw any conclusions about the implied cost functions.

Lemma 2.1. *Consider the two-period, single ESO type, exercise problem in (2.3) and the corresponding cost function (2.4). We assume that the employee's utility function $U(\cdot)$ is twice continuously differentiable. Then, the average cost of an ESO i.e., $\frac{C(\alpha,W)}{\alpha}$ is decreasing in α, for all values of $W > 0$, irrespective of the believed dynamics of S_1, if the following condition on the utility function $U(\cdot)$ is satisfied:*

$\frac{-y \cdot U''(x+y)}{U'(x+y)}$ *is increasing in y for all $x > 0, y > 0$.*

Proof. We fix the non-option wealth W and treat it as a constant for this proof. We first rewrite the problem (2.3) in terms of a different control variable $y \triangleq \frac{x}{\alpha}$. Also for ease of notation we let $P_0 \triangleq \beta_0(S_0 - K)^+$, $P_1 \triangleq \beta_1(S_1 - K)^+$ and $\Delta \triangleq P_1 - P_0$. Then the employee's problem is

$$\begin{aligned} \max_{0 \le y \le 1} V &= \mathbb{E}[U(W + \alpha y P_0 + \alpha(1-y)P_1)] \\ &= \mathbb{E}[U(W + \alpha(P_1 - y\Delta))]. \end{aligned} \tag{2.5}$$

Let us first assume that the optimal y in (2.5), say y^*, satisfies $0 < y^* < 1$. Then (2.5) is a concave maximization problem. First order optimality conditions require

$$\mathbb{E}[U'(W + \alpha(P_1 - y^*\Delta)) \cdot \Delta] \;=\; 0\ . \tag{2.6}$$

Differentiating (2.6) w.r.t. α, we get

$$\begin{aligned}
\mathbb{E}\left[U''\left(W+\alpha(P_1-y^*\Delta)\right)\Delta\left(P_1-y^*\Delta+\alpha\frac{\partial y^*}{\partial\alpha}\Delta\right)\right] &= 0 ; \\
\text{i.e., } \alpha\mathbb{E}\left[-U''(W+\alpha(P_1-y^*\Delta))\Delta^2\right]\frac{\partial y^*}{\partial\alpha} & \\
= \mathbb{E}[-U''(W+\alpha(P_1-y^*\Delta))(P_1-y^*\Delta)\Delta] . & \qquad (2.7)
\end{aligned}$$

Now, if the condition specified in the lemma is satisfied, as $W > 0$ and $P_1 - y^*\Delta = y^*P_0 + (1-y^*)P_1 \geq 0$, we must have

$$\begin{aligned}
\frac{-U''(W+\alpha(P_1-y^*\Delta))\alpha(P_1-y^*\Delta)}{U'(W+\alpha(P_1-y^*\Delta))} &\geq \alpha\gamma(W+\alpha P_0); \quad \text{if } \Delta > 0 , \\
\frac{-U''(W+\alpha(P_1-y^*\Delta))\alpha(P_1-y^*\Delta)}{U'(W+\alpha(P_1-y^*\Delta))} &\leq \alpha\gamma(W+\alpha P_0); \quad \text{if } \Delta < 0 ;
\end{aligned}$$

where

$$\gamma(W+\alpha P_0) \triangleq \frac{-U''(W+\alpha P_0)P_0}{U'(W+\alpha(P_1-y^*\Delta))} .$$

Hence, we must have

$$-U''(W+\alpha(P_1-y^*\Delta))(P_1-y^*\Delta)\Delta \geq \gamma(W+\alpha P_0)U'(W+\alpha(P_1-y^*\Delta))\Delta .$$

Using this fact in (2.7), we get

$$\begin{aligned}
\alpha\mathbb{E}\left[-U''(W+\alpha(P_1-y^*\Delta))\Delta^2\right]\frac{\partial y^*}{\partial\alpha} &\geq \gamma(W+\alpha P_0)\mathbb{E}[-U'(W+\alpha(P_1-y^*\Delta))\Delta] \\
&= 0 .
\end{aligned}$$

Now as $U(\cdot)$ is concave, $U''(\cdot)\Delta^2 < 0$. This in turn implies

$$\frac{\partial y^*}{\partial\alpha} \geq 0 .$$

But,

$$\frac{C(\alpha,W)}{\alpha} = C_1 - y^*(C_1 - C_0) ,$$

where C_1 is the fair value of the European call option with strike K maturing at $T = 1$ and $C_0 = (S_0 - K)^+$. Since $C_1 \geq C_0$, for non-dividend paying stocks, it follows that the average cost of a grant, i.e., $\frac{C(\alpha,W)}{\alpha}$ is decreasing in α in the region where $0 < y^* < 1$. Now consider an $\hat{\alpha}$, such that $y^* = 0$. Since y^* must be continuous as an implicit function of α, if $y^* > 0$ for some α, then we must have $y^* = 0$ for all $\alpha < \hat{\alpha}$. Similarly, if $y^* = 1 \Rightarrow \frac{C(\bar{\alpha},W)}{\bar{\alpha}} = C_0$, for some $\bar{\alpha}$, then $y^* = 1$ for any $\alpha > \bar{\alpha}$. As $C_0 \leq C_1$, this completes the proof. □

Remark 2.1. *The condition in Lemma 2.1 is satisfied by the two most commonly used classes of utility functions. For CARA or exponential type utility functions, where* $U(x) = -e^{-cx}$ *with* $c > 0$*,*

$$\frac{-U''(W+y)y}{U'(W+y)} = cy \ ,$$

which is indeed increasing in y*. For CRRA, or power utility functions, where* $U(x) = \frac{x^{1-a}-1}{1-a}$*,* $a \geq 1$*, we have*

$$\frac{-U''(W+y)y}{U'(W+y)} = a\frac{y}{W+y} = a\left(1 - \frac{W}{W+y}\right) \ ,$$

which is again increasing in y*.*

Remark 2.2. *Note that Lemma 2.1, shows that the average cost of an ESO grant is decreasing in the size of the grant for certain utility functions. This is a weaker condition than to say that the cost of the ESO grant is concave. The latter implies decreasing marginal cost and consequently subsumes decreasing average costs.*

In the next section, we examine the nature of ESO cost functions for CARA and CRRA utilities for the case where the employee's ESO portfolio consists of options of a single type, but expiring after several periods.

2.4 The multi-period problem for one type of options and nature of cost function

In this section, we seek to characterize the ESO cost function, as in Lemma 2.1 for the multi-period case. We continue to assume that the employee has only one type of options, i.e., $N = 1$. Dropping the sub-scripts corresponding to the option type, the multi-period version of the single ESO type portfolio can be written as

$$\begin{aligned}\max V &= \mathbb{E}\left[U\left(W + \sum_{t=0}^{T} \beta_t x_t (S_t - K)^+\right)\right] \; ; \\ \text{s.t.} \quad & \sum_{t=0}^{T} x_t = \alpha \; , \\ & x_t \geq 0 \; . \end{aligned} \tag{2.8}$$

As before, we assume that there is a unique preferred optimal exercise policy[6] denoted by $\mathbf{x}^*$. The corresponding cost function (2.4) is given by

$$C(\alpha, W_0) = \sum_{t=0}^{T} D_t x_t^* (S_t - K)^+ \; .$$

We next consider the CARA and CRRA utility cases separately and show that as in the two-period case, the average ESO cost is decreasing in the size of the grant even when the expiry is several periods away.

Exponential or CARA Utilities

The function $U(\cdot)$ in (2.8) in this case is given by $U(y) = -e^{-cy}$, for some $c > 0$. We first show that the optimal exercise policy for this class of utility functions assumes a relatively simple form. In fact, it is independent of the non-option wealth level W.

Lemma 2.2. *For CARA utilities, the optimal exercise policy x_0^* is independent of the employee's non-option wealth level W and has the form $x_0^* = (\alpha - \eta^*)^+$, where η^* is a quantity that is independent of α, and W.*

Proof. Suppose an exercise policy, $\mathbf{x}^*$ maximizes (2.8) for some α and $W > 0$ for the CARA utility $U(y) = -\exp(-cy)$. This happens if and only if $\mathbf{x}^*$ is a solution to the

[6]In case of multiple competing optimal policies, the employee chooses the one that requires exercising the fewest number of options immediately. This policy should lead to a conservative estimate of costs to the employer.

problem

$$\begin{aligned}
\min V &= \mathbb{E}\left[\exp\left(-c\sum_{t=0}^{T}\beta_t x_t(S_t-K)^+\right)\right] \; ; \\
\text{s.t.} & \quad \sum_{t=0}^{T} x_t = \alpha \; , \\
& \quad x_t \geq 0 \; .
\end{aligned}$$

which has no dependence on W. It then follows that the optimal exercise policy is independent of W.

Since we are only concerned with the dependence of exercise policy on α and W, we fix the other parameters, i.e., current stock price S_0 and the discount factor β_0 that should be applied to any payout received in current period to convert it to its reference time equivalent and treat them as constants. The optimal number of options to be exercised in the current period can be characterized as

$$x_0^*(\alpha, W) = \inf\{x \mid 0 \leq x \leq \alpha \text{ and } x_0^*(\alpha - x, W + \beta_0(S_0 - K)^+) = 0\} \; .$$

Using the independence of optimal policy from non-option wealth, then we have

$$x_0^*(\alpha, W) = \inf\{x \mid 0 \leq x \leq \alpha \text{ and } x_0^*(\alpha - x, 0) = 0\} \; . \tag{2.9}$$

Now, we define the set A_0 and η^* as follows:

$$\begin{aligned}
A_0 &\triangleq \{\alpha \geq 0 \mid x_0^*(\alpha, 0) = 0\} \\
\eta^* &\triangleq \sup A_0 = \sup\{\alpha > 0 \mid x_0^*(\alpha, 0) = 0\}
\end{aligned} \tag{2.10}$$

Suppose A_0 is empty, then we simply set $\eta^* = \infty$ and the lemma holds, as $x_0^*(\alpha, W) = 0$ for all α, W. The proof is also trivial if $x_0^*(\alpha, W) = \alpha$ for all α, in which case we set $\eta^* = 0$. Hence, we consider the case that A_0 is not empty, i.e., $x_0^*(\alpha, W) = x_0^*(\alpha, 0) > 0$ for some $\alpha > 0$. We claim that in this case

- A_0 is bounded above and η^* is finite

- If $\alpha < \eta^*$ then $x_0^*(\alpha, 0) = 0$

If either of the above is not true, then as $x_0^*(\alpha, 0)$ must be continuous in α, there must exist $\underline{\alpha}$, $\bar{\alpha}$ such that $\underline{\alpha} < \bar{\alpha}$ and

$$\begin{aligned} x_0^*(\underline{\alpha}, 0) &= 0 \text{ and} \\ x_0^*(\bar{\alpha}, 0) &= 0 \\ \text{but } x_0^*(\alpha, 0) &> 0 \text{ if } \underline{\alpha} < \alpha < \bar{\alpha} \ . \end{aligned}$$

Let $x_0^*(\bar{\alpha} - \delta, 0) = y$ for some $\delta : \bar{\alpha} - \underline{\alpha} > \delta > 0$. Then, using (2.9),

$$y = \bar{\alpha} - \underline{\alpha} - \delta \ .$$

But, this would mean

$$\lim_{\alpha \to \bar{\alpha}} x_0^*(\alpha, 0) = \bar{\alpha} - \underline{\alpha} > 0 \ ,$$

while $x_0^*(\bar{\alpha}, 0) = 0$. The optimal exercise policy is then discontinuous in grant size, which leads to a contradiction.

Hence η^* is finite and

$$x_0^*(\alpha, 0) = 0 \text{ if } \alpha < \eta^* \ . \tag{2.11}$$

It then follows from (2.9) and (2.11) that the optimal exercise policy is given by

$$x_0^*(\alpha, W) = x_0^*(\alpha, 0) = (\alpha - \eta^*)^+ \ ,$$

which is of the desired form. □

Corollary 2.1. *For CARA utilities, the average cost of an ESO grant is decreasing in the size of the grant.*

Proof. We prove this by induction on the number of time periods to expiry. We can actually prove a stronger result - the marginal cost of issuing an ESO in the case of CARA utilities is decreasing, i.e., the cost function is concave in grant size. The result holds trivially for $T = 0$, i.e., when the option is expiring immediately. Suppose it holds for $T = m$. Let $C(\alpha, W, m)$ denote the cost of α ESO grants with m periods to go, when the employee has non-option wealth W. (We fix the current price of the underlying to S_0 and the accumulated discount factor is β_0.) We have two possibilities:

- $x_0^* = 0$*:* Then $C(\alpha, W, m+1) = \mathbb{E}^{\mathbb{Q}}[D_1 C(\alpha, W, m)|\mathcal{F}_1]$, and the concavity of cost property follows from induction hypothesis.

- $x_0^* > 0$*:* Then, from Lemma 2.2 it follows that the marginal cost of the ESO grant is $(S_0 - K)^+$ for all $\alpha > \eta^*$, with η^* as defined in Lemma 2.9. Since the stock does not pay dividends, by Jensen's inequality the discounted European call option payoff, i.e., $D_t(S_t - K)^+$ is a sub-martingale under the risk neutral measure $\mathbb{Q}$. Then the marginal cost $(S_0 - K)^+$ for grant size $\alpha > \eta^*$ is less than that for any size $\alpha' < \eta^*$, which is at least $\mathbb{E}^{\mathbb{Q}}[D_1(S_1 - K)^+|\mathcal{F}_1]$.

□

Remark 2.3. *For the exponential utility model, the asymptotic marginal cost of granting an in-the-money ESO, as the size of the total grant* $\alpha \to \infty$*, turns out to be equal to the options intrinsic value, i.e.,* $(S_0 - K)$*. This is the same as the cost at which companies used to expense ESOs prior to the FASB stipulations in 2005. However, the intrinsic value for this model comes out as a marginal cost and not as the average cost, as was used for cost accounting.*

Power or CRRA Utilities

For CRRA or power utilities[7] $U(\cdot)$ in (2.8) takes the form $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$, for some $\gamma > 1$. In Section 2.4, for CARA utilities, the optimal exercise policy was shown to be independent of non-option wealth W and to assume a very simple form. For CRRA utilities, the non-option wealth W will impact exercise policy. Nevertheless, as we show in the Lemma 2.3, the dependence of optimal exercise policy in grant size and non-option wealth level is easily characterized in this case as well.

Lemma 2.3. *For CRRA utilities, the optimal exercise policy has the form*

$$x_0^*(\alpha, W) = \left(\theta\alpha - (1-\theta)\frac{W}{\beta_0(S_0 - K)^+}\right)^+,$$

where θ *is independent of* α*,* W *and* $0 \leq \theta < 1$*.* S_0 *denotes the current stock price and* β_0 *the accumulated discount factor.*

Proof. For this proof, we again fix β, the subjective discount factor for current payoffs and the current stock price S.

[7]This class also includes log utilities, i.e., $U(x) = \ln(x)$, which can also be considered as a power utility in the limiting case $\gamma \to 1$.

We first note that the value function is homogeneous for CRRA utilities. If $V_d(x, \alpha, W)$ denotes the utility derived by following an exercise policy $\mathbf{x}$ given an initial grant α and non-option wealth level W, then we must have

$$V_d(M\mathbf{x}, M\alpha, MW) \;\; = \;\; M^{1-\gamma} \cdot V_d(\mathbf{x}, \alpha, W_0)$$

Also note that, if $\mathbf{x}x$ is a feasible exercise policy to follow for grant α, then $M\mathbf{x}$ must be a feasible policy for the grant $M\alpha$. This means that if $\mathbf{x}^*$ is an optimal exercise policy for a grant of size α and when non-option wealth is $\frac{W}{M}$, then $M \cdot \mathbf{x}^*$, is the optimal exercise policy for a grant of size α and non-option wealth W. Hence

$$x_0^*(M\alpha, W_0) \;\; = \;\; M \cdot x_0^* \left(\alpha, \frac{W}{M} \right) \; . \tag{2.12}$$

Suppose the options have strike K and expiry T. The optimal exercise quantity $x_0^*(\alpha, W)$ must satisfy the following condition:

$$x_0^*(\alpha, W) \;\; = \;\; \inf\{x \mid 0 \leq x \leq \alpha \text{ and } x_0^*(\alpha - x, W + x \cdot \beta_0(S_0 - K)^+) = 0\} \; .$$

Then, using (2.12) we get

$$x_0^*(\alpha, W) \;\; = \;\; \inf \left\{ x \;\mid\; 0 \leq x \leq \alpha \text{ and } x_0^* \left(\frac{\alpha - x}{W + x \cdot \beta_0(S_0 - K)^+}, 1 \right) = 0 \right\} \tag{2.13}$$

Suppose $x_0^* = 0$ for all combinations of α and W. Then the form specified in the lemma trivially applies with $\theta = 0$. If not, then we must have $S > K$. Consider then a certain combination (α, W), say (α_A, W_A) such that

$$x_0^*(\alpha_A, W_A) \;\; \triangleq \;\; x_A^* > 0$$

Using (2.13), then

$$x_0^* \left(\frac{\alpha_A - x_A}{W_A + x_A \beta_0(S_0 - K)}, 1 \right) \;\; = \;\; 0 \tag{2.14}$$

Now, we define the set A_0 and quantity κ as follows:

$$\begin{aligned} A_0 \;\; &\triangleq \;\; \{\alpha \geq 0 \mid x_0^*(\alpha, 1) = 0\} \\ \kappa \;\; &\triangleq \;\; \sup A_0 = \sup\{\alpha > 0 \mid x_0^*(\alpha, 1) = 0\} \; . \end{aligned} \tag{2.15}$$

(2.14) shows that the set A_0 is non-empty. We now show that

- κ must be finite, i.e., A_0 is bounded above and

- If $\alpha < \kappa$, then $x_0^*(\alpha, 1) = 0$.

If either of these claims is not true then as the optimal exercise policy must be continuous in grant size, there would exist $\underline{\alpha}$, $\bar{\alpha}$ such that $\underline{\alpha} < \bar{\alpha}$ and

$$\begin{aligned} x_0^*(\underline{\alpha}, 1) &= 0 \text{ and} \\ x_0^*(\bar{\alpha}, 1) &= 0 \\ \text{but } x_0^*(\alpha, 1) &> 0 \text{ if } \underline{\alpha} < \alpha < \bar{\alpha} \ . \end{aligned}$$

Let $x_0^*(\bar{\alpha} - \delta, 1) = y$ for some $\delta : \bar{\alpha} - \underline{\alpha} > \delta > 0$. Then, using (2.13)

$$\begin{aligned} \frac{\bar{\alpha} - \delta - y}{1 + y\beta_0(S_0 - K)} &= \underline{\alpha} \ , \\ \text{i.e., } y &= \frac{\bar{\alpha} - \delta - \underline{\alpha}}{1 + \underline{\alpha}\beta_0(S_0 - K)} \ . \end{aligned}$$

But, this would mean that

$$\lim_{\alpha \to \bar{\alpha}^+} x_0^*(\alpha, 1) = \frac{\bar{\alpha} - \underline{\alpha}}{1 + \underline{\alpha}\beta_0(S_0 - K)} > 0 \ ,$$

while $x_0^*(\bar{\alpha}, 1) = 0$; i.e., the optimal exercise policy is discontinuous in grant size, which cannot be the case as the value function is continuously differentiable. From (2.13) and (2.15), it follows that if $\frac{\alpha}{W} \geq \kappa$, then

$$\kappa = \frac{\alpha - x_0^*(\alpha, W)}{W_A + x_0^*(\alpha, W) \cdot \beta_0(S_0 - K)} \ , \tag{2.16}$$

$$\text{i.e., } x_0^*(\alpha, W) = \frac{1}{\kappa\beta_0(S_0 - K) + 1}\alpha - \frac{\kappa\beta_0(S_0 - K)}{\kappa\beta_0(S_0 - K) + 1}\frac{W_A}{\beta_0(S_0 - K)}. \tag{2.17}$$

Since $x_0^*(\alpha, 1) = 0$ if $\alpha < \kappa$, using (2.12), we conclude

$$x_0^*(\alpha, W) = 0 \text{ if } \frac{\alpha}{W} < \kappa \ . \tag{2.18}$$

Combining (2.17) and (2.18), we get

$$x_0^*(\alpha, W) \;=\; \left(\frac{\alpha - \kappa \cdot W}{1 + \kappa \cdot \beta_0 (S_0 - K)^+}\right)^+ .$$

which is of the form stated in the lemma, and the parameter $\theta = \frac{1}{1+\kappa\beta_0(S_0-K)}$ is independent of α and W. Note that the exercise quantity given by (2.19) always satisfies the constraint $x_0^* < \alpha$. □

Lemma 2.3 can be used to show that average costs of an ESO portfolio with a single type of options is decreasing when the employee has CRRA utility. We first prove the following useful result, which formalizes the notion that early exercises tend to reduce ESO costs. For this, the following definition will be helpful:

Definition 2.1. *Consider two strategies $\mathbf{x^A}$ and $\mathbf{x^B}$ for exercising an option grant of size α. Strategy $\mathbf{x^A}$ is said to dominate strategy $\mathbf{x^B}$ if it always leaves a greater number of options unexercised i.e.,*

$$\alpha - \sum_{t=0}^{t} x_t^A \;\geq\; \alpha - \sum_{t=0}^{t} x_t^B \;, 0 \leq t \leq T$$

Next, we show that the cost of a grant associated with a dominant strategy is always higher.

Lemma 2.4. *If strategy $\mathbf{x^A}$ to exercise an option grant of size α dominates another strategy $\mathbf{x^B}$, then the option cost C^A associated with strategy $\mathbf{x^A}$ is higher than the cost C^B associated with $\mathbf{x^B}$.*

Proof. For $0 \leq t \leq T-1$, recall

$$\alpha_t^A \;\triangleq\; \alpha - \sum_{s=0}^{t} x_s^A \;;$$

$$\alpha_t^B \;\triangleq\; \alpha - \sum_{s=0}^{t} x_s^B \;.$$

α_t^A and α_t^B are $\mathcal{F}_t$ measureable. As $\mathbf{x_A}$ dominates $\mathbf{x_B}$, we must have $\alpha_t^A \geq \alpha_t^B$. Also, let $P_t \triangleq D_t(S_t - K)^+$. Then P_t is a sub-martingale as the stock does not pay dividends. Now,

$$
\begin{aligned}
C^A - C^B &= \mathbb{E}^{\mathbb{Q}}\left[\sum_{t=0}^{T-1}(x_t^A - x_t^B)P_t + (\alpha_{T-1}^A - \alpha_{T-1}^B)P_T\right] \\
&= \mathbb{E}^{\mathbb{Q}}\left[\sum_{t=0}^{T-1}(x_t^A - x_t^B)P_t + (\alpha_{T-1}^A - \alpha_{T-1}^B)\mathbb{E}^{\mathbb{Q}}[P_T|\mathcal{F}_{T-1}]\right] \\
&\geq \mathbb{E}^{\mathbb{Q}}\left[\sum_{t=0}^{T-1}(x_t^A - x_t^B)P_t + (\alpha_{T-1}^A - \alpha_{T-1}^B)P_{T-1}\right] \qquad (2.19) \\
&= \mathbb{E}^{\mathbb{Q}}\left[\sum_{t=0}^{T-2}(x_t^A - x_t^B)P_t + (\alpha_{T-2}^A - \alpha_{T-2}^B)P_{T-1}\right] \qquad (2.20) \\
&\quad \ldots \\
&= 0\ .
\end{aligned}
$$

In (2.19), we used the fact that $\alpha_{T-1}^A \geq \alpha_{T-1}^B$ and that P_t is a sub-martingale. In (2.20), we substituted $\alpha_{T-1}^A = \alpha_{T-2}^A - x_{T-1}^A$ and $\alpha_{T-1}^B = \alpha_{T-2}^B - x_{T-1}^B$. Note that the inequality is strict if the discounted option payoff process $D_t(S_t - K)^+$ is a strict sub-martingale and there is a non-zero probability of a positive difference in unexercised option positions. □

Remark 2.4. *Option payoff process is guaranteed to be a $\mathbb{Q}$ sub-martingale if the stock does not pay dividends. Lemma 2.4 also holds under a weaker condition - the employee must exercise all her options, whenever the stock reaches a level at which the payoff process no longer remains a sub-martingale under the $\mathbb{Q}$ measure.*

We now use Lemmas 2.3 and 2.4 to show that ESO cost is sub-linear in grant size for CRRA utilities.

Corollary 2.2. *If the employee has CRRA utility, then the average cost of a batch of ESOs, all with the same terms, is decreasing in the size of the grant.*

Proof. Fixing, S_0, T and β_0, from the homogeneity of CRRA utilities, (2.12), it follows that

$$
C(\alpha, W) = \alpha C\left(1, \frac{W}{\alpha}\right)\ .
$$

Thus, to prove that average cost is decreasing in α, it suffices to show that cost of the ESO grant is increasing in initial wealth i.e., W. For this, we appeal to Lemma 2.4 and demonstrate that the number of unexercised options associated with a higher non-option wealth level dominates the number of unexercised options with a lower non-option wealth level on a path-by-path basis using finite induction.

Consider, for an option with strike K and expiry T, two different combinations of grant sizes and non-option wealth (α_0, W_0) and $(\bar{\alpha_0}, \bar{W_0})$ at time 0, such that $\alpha_0 \geq \bar{\alpha_0}$ and $W_0 \geq \bar{W_0}$. We will show that these inequalities are preserved throughout the option's life-time.

Suppose $S_0 \leq K$, then there are no exercises at $t = 0$, and the first combination will continue to dominate the second at $t = 1$. If $S_0 > K$, then using Lemma 2.3, the difference in unexercised options after exercises at time 0 are accounted for will be

$$\begin{aligned}
\alpha_1 - \bar{\alpha}_1 &= \alpha_0 - \left(\theta\alpha_0 - (1-\theta)\frac{W_0}{\beta_0(S_0-K)}\right)^+ - \left(\bar{\alpha_0} - \left(\theta\bar{\alpha_0} - (1-\theta)\frac{\bar{W_0}}{\beta_0(S_0-K)}\right)^+\right) \\
&\geq \alpha_0 - \bar{\alpha_0} - \left(\theta(\alpha_0 - \bar{\alpha_0}) - (1-\theta)\frac{W_0 - \bar{W_0}}{\beta_0(S_0-K)}\right)^+ \\
&\geq (1-\theta)(\alpha_0 - \bar{\alpha_0}) \\
&\geq 0\,.
\end{aligned}$$

Also, the difference in subjectively discounted non-option wealth at $t = 1$ will be

$$\begin{aligned}
W_1 - \bar{W}_1 &= W_0 + \beta_0\left(\theta\alpha_0 - (1-\theta)\frac{W_0}{\beta_0(S_0-K)}\right)^+ (S_0 - K) \\
&\quad - \left(W_0 + \beta_0\left(\theta\alpha_0 - (1-\theta)\frac{\bar{W_0}}{\beta_0(S_0-K)}\right)^+ (S_0-K)\right) \\
&\geq W_0 - \bar{W_0} - \left(-\theta(\alpha_0 - \bar{\alpha}_1) + (1-\theta)\frac{W_0 - \bar{W_0}}{\beta_0(S_0-K)}\right)^+ \beta_0(S_0-K) \\
&\geq W_0 - \bar{W_0} - \left((1-\theta)\frac{W_0 - \bar{W_0}}{\beta_0(S_0-K)}\right)^+ \beta_0(S_0-K) \\
&\geq \theta(W_0 - \bar{W_0}) \\
&\geq 0\,.
\end{aligned}$$

Thus the first combination will always dominate the second at the beginning of the period $t = 1$. By repeating this argument, we see that number of unexercised option associated with the first grant will always dominate the ones associated with the second. Moreover, the difference will become strict whenever there is an exercise. Using Lemma 2.4 then we conclude that when discounted option payoff process is a $\mathbb{Q}$ sub-martingale, the cost function is increasing in the non-option wealth position. This, in turn implies that the average cost is decreasing in the size of the grant. $\square$

We demonstrated in this section, that irrespective of stock price dynamics, when the employee has CARA or CRRA type preferences, her exercise policies would make average ESO costs decreasing in the grant size. We now consider the case when the employee's portfolio consists of multiple types of options, i.e., $N > 1$. The analysis for this case is considerably involved. We begin by showing that there exists a partial and intuitive order of exercises between different option types in the next section.

2.5 Option Exercises with Multiple Option Types

Generally speaking, an employee is likely to have unexercised options of multiple types in her portfolio. However, because of the explosion of number of state variables (options of each type have to be tracked), the problem becomes considerably difficult to analyze. Before, moving on to the question of cost function structure in this case, we consider some properties of relative exercises between options of different types that can help to make the analysis easier. We begin by showing that the model based on expected utility of a measure of terminal wealth, as described in (2.1), guarantees a certain rational order of exercise between competing options. This is summarized in the following lemma:

Lemma 2.5. *The exercise policy of an employee seeking to optimize her expected utility as in (2.1) will satisfy the following properties:*

1. *Any unexercised and expiring in the money options will be exercised.*
2. *For two options with same strike K but different expiries, the ones with the earlier expiry will be exercised in entirety before any of the options with the later expiry are exercised.*
3. *For two options with different strikes but same expiries, the ones with the lower strike will be exercised in entirety before any of the options with the higher strike are exercised.*

Proof. The first property is trivial to show. For Property 2, simply note that if there exists an optimal policy which exercises an option with a later expiry before an option with an earlier expiry, with both having the same strikes, then switching the unexercised earlier expiry options for an equal number of options with the later expiry will result in a dominant strategy. Property 3 also follows from a similar argument, but requires a little more effort. Let i, j be such that $T_i = T_j$ but $K_i < K_j$. Assume that property 3 does not hold. Then for some values of W_t, S_t, α_i, α_j we must have $x^*_{i,t} < \alpha_i$ but $x^*_{j,t} > 0$. Now, consider an alternate strategy $\tilde{x}$ - which is identical to x^* except for following differences:

- $\tilde{x}_{i,t} = x^*_{i,t} + \epsilon$ and $\tilde{x}_{j,t} = x^*_{j,t} - \epsilon$, where $\epsilon = \min\{\alpha_i - x^*_{i,t}, x^*_{j,t}\} > 0$. The employee's wealth at time t due to this strategy change increases by $\Delta = \epsilon \cdot \beta_t(K_j - K_i)$

- If along any sample path, we subsequently have $\tilde{\alpha}_i(s,\omega) < x^*_i(t,\omega)$ for $s > t$, then set $\tilde{x}_{i,t} = \tilde{\alpha}_i(s,\omega)$ and $\tilde{x}_j(t,\omega) = x^*_j(t,\omega) + x^*_i(t,\omega) - \tilde{\alpha}_i(s,\omega)$. Due to this, the employee's payoff in this state reduces by the amount $\Delta(s,\omega)$ such that

$$\begin{aligned}\Delta(s,\omega) &= (x^*_i(t,\omega) - \tilde{\alpha}_i(s,\omega)) \cdot \beta_s((S_s - K_i)^+ - (S_s - K_j)^+ \\ &\leq (x^*_i(t,\omega) - \tilde{\alpha}_i(s,\omega))\beta_{t+1}(K_j - K_i) \ .\end{aligned}$$

Thus, the net change in terminal wealth by employing strategy $\tilde{x}$ over x^* for any sample path ω is given by

$$\begin{aligned}\Delta W &= \Delta - \sum_{s=t+1}^{T_i} \Delta(s,\omega) \\ &\geq \epsilon\beta_t(K_j - K_i) - \sum_{s=t+1}^{T_i} (x^*_i(t,\omega) - \tilde{\alpha}_i(s,\omega))\beta_{t+1}(K_j - K_i) \\ &\geq \epsilon(K_j - K_i)(\beta_t - \beta_{t+1}) > 0 \ ;\end{aligned}$$

since $\sum_{s=t+1}^{T_i}(x^*_i(t,\omega) - \tilde{\alpha}_i(s,\omega)) = \epsilon$. Thus $\tilde{x}$ dominates terminal payoffs over x^* on a path-by-path basis and hence x^* cannot be an optimal exercise policy for (2.1). $\square$

Although, exercise properties mentioned in Lemma 2.5 appear intuitive and somewhat obvious, these properties - especially property 3, are not satisfied by all proposed models of exercise behavior. For example, consider the model proposed in Jain and Subramanian [75], in which the employee exercises to optimize utility of intertemporal consumption. In their model, employee's optimization problem takes the form

$$\max \mathbb{E}\left[\sum_{t=1}^{T} U(t, P_t)\right] \ .$$

where P_t denotes the payoff received at time t from option exercises. We illustrate a simple example in Appendix A, that shows that under this model Property 3 of

Lemma 2.5 can be violated[8].

Lemma 2.4 imposes a partial order of relative exercise between different vested options. However, it cannot predict the relative order of exercise between two options with strikes K_1 and K_2 and expiries T_1 and T_2 when $K_1 < K_2$, but $T_1 > T_2$. The exact order of exercise will depend on the believed stock dynamics and one can even have an interleaving of the exercises of two options. This makes the general case of multi-period multi-type ESO costing quite challenging.

In Section 2.6, we now turn again to the key question of interest - the nature of the cost associated with a portfolio of ESOs when the employee has multiple types of options. We consider the case of CARA utilities as they offer a considerable ease in understanding this cost structure.

2.6 ESO costs for portfolios with multiple ESO types

In Section 2.4, we observed that for most common utility functions the average cost of an ESO grant is decreasing in grant size, irrespective of the stock-price dynamics. This suggests that the cost of an ESO option portfolio is sub-linear i.e., it is less than the sum of its parts. In this section, we seek to examine if the analogous property of sub-additivity holds for portfolios with multiple ESO types. Unfortunately, it turns out that in this case, the option cost can be super-linear for certain stock price processes. However, under an "acceptable" set of stock price dynamics, as we will clarify shortly, the sub-linearity property can be shown to hold for CARA utilities. We also demonstrate that when there are multiple types of options, new grants can have a retrospective effect on already issued ESOs (if the grant is unanticipated), by making them cost more or less.

We consider only the CARA class of utilities in this section. CARA utilities offer considerably simplify the analysis because future exercises become independent from past exercises for these models.

We start with proving a useful property of exercise policies associated with CARA utilities, that will help us to better understand the costs associated with different parts of an ESO portfolio.

[8] Note that one would expect an employee to exercise an option with lower strikes first even when faced with liquidity related issues.

Lemma 2.6. *Suppose the employee has a CARA type utility function and α_1 vested ESOs with expiry T_1 and α_2 vested ESOs with expiry T_2 where $T_1 < T_2$. Both sets of options have the same strike K. Then, her exercise policy and associated cost for α_2 options with expiry T_2 is the same as in the case when she is made only this grant and not the α_1 options with expiry T_1.*

Proof. To simplify notation, we consider two employees, A and B who have the same CARA type preferences and identical beliefs about the stock price process S_t and also use the same subjective discounting β_t. Employee A has α_1 options with expiry T_1 as well as α_2 options with a later expiry T_2. Employee B on the other hand has only α_2 options with expiry T_2. All options have the same strike K. Then the lemma states that A and B will have identical policies for exercising the options with expiry T_2.

From Lemma 2.4 it follows that A will have exercised all α_1 options with expiry T_1 before any option with expiryy T_2 is exercised. Then, at the beginning of any period, $A's$ portfolio will contain

1. some $y > 0$ options with expiry T_1 and α_2 options with expiry T_2 OR

2. only unexercised options with expiry T_2.

In the latter is true then as exercise policy of ESOs in case of CARA utilities are independent of the non-option wealth level, further exercises of options with expiry T_2 by A, occur independently and uninfluenced by prior exercises of options with expiry T_1. Hence whenever A exercises options with expiry T_2, B must also exercise an identical number of her options with expiry T_2, provided they had the same number of options with expiry T_2 at the beginning of the period. It also follows that they will have identical exercises from that point onwards.

Thus to show the result of the lemma, all we need to show is that whenever A leaves some options with expiry T_1 unexercised, B does not exercise any of her options.

To see why this is true, suppose at some point, A finds it optimal to leave $\alpha' > 0$ options with expiry T_1 and all α_2 options with expiry T_2 unexercised. Now consider another hypothetical employee C, identical to A and B in preferences, beliefs and investment opportunities but having $\alpha' + \alpha_2$ options with expiry T_2 at this juncture. We argue that C will leave at least α_2 of her options unexercised. Indeed, if this is not the case then we have the following contradiction: as C has a strictly dominating ESO portfolio over A, she must realize at least as much utility as A expects at that point. If C exercises more than α' of her options, it is clear that her strategy can be copied by A and hence their expected utilities at that point must be equal. Then $A's$ exercise policy which can always be implemented by C, is also optimal for C and should be the one chosen by C as it requires fewer exercises. Hence it must be optimal for C to leave more than α_2 of her options unexercised. From Lemma 2.2, it follows that then B should not exercise any of her options, when A does not exercise options that have expiry T_2.

This also shows that the cost associated with the grant with expiry T_2 is not affected by any grant of options with the same (or even lower) strike and an earlier expiry for CARA utilities. This establishes the result we sought to prove. □

Lemma 2.7. *Suppose the employee has CARA type utility function and α_1 vested ESOs with strike K_1 and α_2 vested ESOs with maturity K_2, where $K_1 < K_2$. Both sets of options have the same expiry T. Then, her exercise policy and associated cost for α_2 options with strike K_2 is the same as in the case when she is made only this grant and not the α_1 options with strike K_1.*

Proof. The proof is broadly similar to that of Lemma 2.7 and omitted. □

Lemma 2.6 shows that for CARA utilities, when an employee has a portfolio of two different types of ESOs differing only in their expiries, the cost of the ESO with a later expiry is essentially the same as it would have been when the employee was made a grant of only that type of options. It is easy to see by a similar reasoning that even when the employee has CRRA utilities, the exercise for ESOs with later expiries will get "delayed" by presence of any options with earlier expiries. This usually means that such a grant will cost more to the company then it would have on a stand-alone basis. If the property of sub-additivity were to hold, then it must be true that the exercise of options with the earlier expiry gets pushed forward and the consequent fall in their cost makes up for any rise in the cost of options with the later expiry. Rather surprisingly, this does not hold in general and thus the sub-linearity property does not extend directly to ESO portfolios with multiple types of grants. The following example illustrates the issue.

Example 2.2. *We consider an employee with the CARA utility function $U(W) = -\exp(-cW)$, with $c = 0.00025$. This employee has $\alpha_1 = 100$ options expiring at time $T_1 = 1$ and $\alpha_2 = 100$ options expiring at $T_2 = 2$. All options are vested and have strike $K = 90$. Further, the dynamics of stock price currently at 100 are as illustrated in Figure 2.2.*

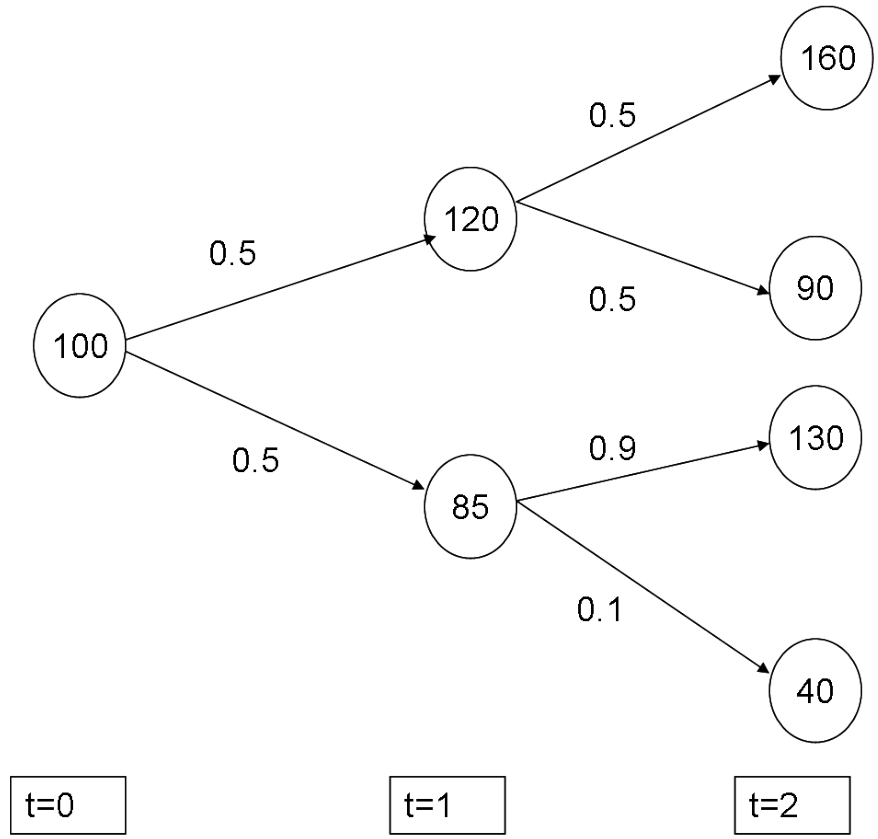

Figure 2-1: Employee's belief about stock price dynamics. Numbers on arrows indicate a probability of transition while the circles enclose realized stock price.

We assume both the risk-free as well as subjective discounting factors to be constant at 1 *for simplicity. It turns out that that for this particular problem, the optimal exercise strategy for the employee, with respect to the options expiring at* $T = 1$ *is to not exercise any of them at* $t = 0$*. The corresponding fair value cost to the employer for this part of the grant as a result is* $100 \cdot 12.88 = 1288$*. If however, the employer was made a grant of only* 100 *options expiring at* $T_1 = 1$ *and none expiring at* T_2*, then the employee would have been compelled to exercise* 7.58 *options immediately and the rest at* $t = 1$*. The resulting cost would have been* $7.58 \cdot 10 + 92.42 \cdot 12.88 = 1263$*. As we know from Lemma 2.7 the cost associated with the option grant corresponding to expiry* $T = 2$ *remains the same as it would have been had the grant been made in isolation.*

Thus, in this particular case, the cost of the total option grant turns out to be more than the sum of its parts! In this example, the presence of longer duration options helps the employee to in essence "diversify" the risk of payoffs with options expiring at $T = 1$*. In the event of a price decline at time* $t = 1$*, the presence of ESOs maturing* *at* 2 *will still allow her to receive a high payoff with a high likelihood. This enables her to hold on to some of the options expiring at* $T = 1$*, which she would have otherwise exercised.*

In this same example, if we consider a more risk-averse employee with value of $c = 0.001$*, then this behavior flips. This employee, for the same stock price process, will exercise all* 100 *options expiring at* $T = 1$ *if she also had the options expiring at* $T = 2$*. On the other hand, if she did not have the options expiring latter, she would have exercised only* 23 *options. It is easy to see that, the option cost of the portfolio will be sub-linear or less than the sum of costs associated with the grants made in isolation.*

Example 2.2 shows that cost structure of portfolio with multiple types of ESO grants can be difficult to generalize.

It turns out that we can trace the possibility of supper-additivity of option portfolios to that of the use of different types of options for "diversification". For most asset dynamic models this should not be the case. We now formalize this notion. We first fix an option type i to simplify notation. We define for CARA utilities, the Certainty Equivalent and Incremental Certainty Equivalent of an option grant as follows:

Definition 2.2. *Let the optimal utility of an employee with CARA utility and risk aversion parameter* c *for an endowment of* α *options of the type* i *be* $U^*(\alpha)$*. Then, the 'Certainty Equivalent' of this grant,* $h(\alpha)$ *is defined as*

$$h(\alpha) \;=\; -\frac{\ln(-U^*(\alpha))}{\beta c} \;,$$

where as before β *is the period's subjective discount factor.*

Certainty equivalent can thus be viewed as a measurement of utility in "cash" units, i.e, the amount of money needed to realize the same utility.

Definition 2.3. *Suppose the optimal utility for an employee with CARA utility and risk aversion parameter* c *holding a portfolio* P *of options is* $U^*(P)$ *and in case he is made an additional grant of* α *options of type* i*, the same changes to* $U^*(\alpha, P)$*. Let* β_t *be the accumulated discounting factor. If the current discounting factor is* β*, then the "Incremental Certainty Equivalent" of this grant,* $h_t(\alpha|P)$ *is defined as*

$$h_t(\alpha|P) \;=\; -\frac{\ln(-U^*(\alpha, P))}{\beta_t c} + \frac{\ln(-U^*(P)))}{\beta_t c} \;.$$

Similarly, we define the incremental certainty equivalent of the portfolio P as

$$h_t(P|\alpha) \;=\; -\frac{\ln(-U^*(\alpha,P))}{\beta_t c} + \frac{\ln(-U^*(\alpha))}{\beta_t c} \;.$$

We then define the "option non-diversifiability" condition as follows:

Definition 2.4. *The stock price process considered by the employee is " option non-diversifiable" if the incremental certainty equivalents of any two option grants irrespective of their size or type are co-monotone.*

Remark 2.5. *Most commonly encountered stock price processes are independent "returns" processes. For such processes, it is easy to see that Incremental Certainty Equivalents are increasing functions of the current stock price and thus should be co-monotone.*

We now show that when the option non-diversifiability condition is satisfied, the cost associated with an ESO portfolio is in fact sub-additive, i.e., less than the sum of its parts. We first show that presence of additional grants speeds up exercises of ESOs in general.

Theorem 2.1. *For a CARA employee, if the stock price process considered by the employee is option non-diversifiable then for any endowment P and option i and grant size α:*

- $$h_0'(\alpha) \;\geq\; h_0'(\alpha|P) \;,$$
- *The optimal exercise strategy corresponding to the option type i in presence of additional portfolio P is dominated by the optimal exercise strategy when not endowed with P.*

Proof. We prove this by induction on the time to expiry of option i, T_i. For $T_i = 0$, the proposition trivially holds as both $h_0(\alpha)$ and $h_0(\alpha|P)$ are given by $(S_{T_i} - K_i)^+$, and either all or none of the options will be exercise in either case. Suppose the proposition is true for $T_i = k$. Then, we show the same should hold for $T_i = k+1$. Suppose with $T_i = k+1$ periods to go, in the case when the employee also has an additional endowment P, the optimal exercise policy leaves her with a part P' of the portfolio P and $\bar{\alpha}$ options of type i. Then, as it is sub-optimal for the employee to exercise any more options of type i, the first order optimality conditions imply

$$\mathbb{E}[-cU(\bar{\alpha}, P') \cdot \beta_0 (S_0 - K_i)^+ + U_\alpha(\bar{\alpha}, P')] \;\; \leq \;\; 0 \; ,$$

Hence,

$$\begin{aligned} h_0'(\bar{\alpha}|P) = h_0'(\bar{\alpha}|P') \;\; &= \;\; \frac{\mathbb{E}[\beta_1 \cdot \exp(-\beta_1 c \cdot (h_1(\bar{\alpha}) + h_1(\bar{\alpha}|P'))) \cdot h_1'(\bar{\alpha}|P')]}{\mathbb{E}[\exp(-\beta_1 c \cdot (h_1(P') + h_1(\bar{\alpha}|P')))]} \\ &\geq \;\; \beta_0 (S_0 - K)^+ \; . \end{aligned} \tag{2.21}$$

Now, using the induction hypothesis,

$$\begin{aligned} & \frac{\mathbb{E}[\beta_1 \cdot \exp(-\beta_1 c \cdot (h_1(\bar{\alpha}) + h_1(\bar{\alpha}|P'))) \cdot h_1'(\bar{\alpha}|P')]}{\mathbb{E}[\exp(-\beta_1 c \cdot (h_1(P') + h_1(\bar{\alpha}|P')))]} \\ \leq \;\; & \frac{\mathbb{E}[\beta_1 \cdot \exp(-\beta_1 c \cdot (h_1(\bar{\alpha}) + h_1(\bar{\alpha}|P'))) \cdot h_1'(\bar{\alpha})]}{\mathbb{E}[\exp(-\beta_1 c \cdot (h_1(P') + h_1(\bar{\alpha}|P')))]} \; . \end{aligned} \tag{2.22}$$

We now define a new probability measure $\mathbb{R}$, such that its Radon-Nikodym derivative w.r.t. the original measure $\mathbb{P}$ is given by

$$\frac{d\mathbb{R}}{d\mathbb{P}} \;\; = \;\; \frac{\exp(-\beta_1 c \cdot h_1(\bar{\alpha}))}{\mathbb{E}[\exp(-\beta_1 c \cdot h_1(\bar{\alpha}))]} \; .$$

Then, from (2.21) and (2.22), we have

$$\begin{aligned} h_0'(\bar{\alpha}|P) \;\; &\leq \;\; \frac{\mathbb{E}[\exp(-\beta_1 c \cdot (h_1(\bar{\alpha}) + h_1(P'|\bar{\alpha}))) \cdot \beta_1 \cdot h_1'(\bar{\alpha})]}{\mathbb{E}[\exp(-\beta_1 c \cdot (h_1(\bar{\alpha}) + h_1(P'|\bar{\alpha})))]} \\ &= \;\; \frac{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha})) \cdot \beta_1 \cdot h_1'(\bar{\alpha})]}{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha}))]} \\ &= \;\; \lim_{\Delta \to 0} \frac{1}{\Delta} \frac{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha})) \cdot \beta_1 \cdot h_1(\bar{\alpha} + \Delta|\bar{\alpha})]}{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha}))]} \\ &\leq \;\; \lim_{\Delta \to 0} \frac{1}{\Delta} \frac{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha}))] \cdot \mathbb{E}^{\mathbb{R}}[\beta_1 \cdot h_1(\bar{\alpha} + \Delta|\bar{\alpha})]}{\mathbb{E}^{\mathbb{R}}[\exp(-\beta_1 c \cdot h_1(P'|\bar{\alpha}))]} \qquad (2.23) \\ &= \;\; \frac{\mathbb{E}[\exp(-\beta_1 c \cdot h_1(\bar{\alpha})) \cdot \beta_1 \cdot h_1'(\bar{\alpha})]}{\mathbb{E}[\exp(-\beta_1 c \cdot h_1(\bar{\alpha}))]} \; . \qquad (2.24) \end{aligned}$$

In (2.23), we have made use of the fact that as $\beta_1 h_1(\bar{\alpha} + \Delta|\bar{\alpha})$ and $\beta_1 h_1(P|\bar{\alpha})$ are co-monotone by the assumption of option non-diversifiability, $\exp(-c\beta_1 h_1(\bar{\alpha} + \Delta|\bar{\alpha}))$ and $\beta_1 h_1(P|\bar{\alpha})$ must be negatively correlated.

From, (2.21) and (2.24) it follows that in the absence of the additional endowment

P, the employee must leave $\tilde{\alpha} \geq \bar{\alpha}$ options unexercised. This proves the second part of the hypothesis.

Now to show the first part, we consider two cases:

1. The employee does not exercise any type i options when endowed with additional portfolio P. In this case, the option holder must also not exercise any options, when not endowed with P. Thus $\bar{\alpha} = \tilde{\alpha} = \alpha$ and the first part of the induction hypothesis follows from (2.24).

2. The employee makes an exercise in the presence of P. In this case, $h(\alpha|P) = (S - K_i)^+$, which is a lower bound on the marginal certainty equivalent and hence the first part of the induction hypothesis trivially holds.

□

Corollary 2.3. *For CARA employees, the marginal cost to the employer of any option grant is less when the employee has a portfolio of some unexercised options compared to the case when she has none. Thus, the cost of the option portfolio is sub-additive in its components.*

Proof. This follows directly from Theorem 2.1 and Lemma 2.4. □

Remark 2.6. *It should, in fact, be easy to show a relatively stronger condition. If a portfolio P can be parceled in two sub-portfolios P_1 and P_2, then*

$$C(P) \;\; \leq \;\; C(P_1) + C(P_2) \; .$$

2.7 Summary

In this chapter, we demonstrated that one should not consider the employee decision associated with exercising an ESO in isolation, but rather in the context of the entire non-tradeable ESO portfolio that she carries. We sought to qualitatively analyze how adopting such a point of view would impact the exercise behavior of the employee and consequently the implied cost to the issuing company using a simple utility based model as described in (2.1). We observed that in general, the cost function can be sub-linear or super-linear in the size of the portfolio. However, for utility functions commonly considered in literature, such as CARA and CRRA utilities, when the ESO portfolio has only one type of option, the average cost of the grant can be shown to be decreasing in grant size, irrespective of stock price process dynamics.

The problem corresponding to ESO portfolios with options of multiple types, especially multiple maturities, is more complex. We first showed that the proposed model ensures a "rational" order of exercise between different options. Then, we showed how this property results in certain simplifications for CARA type utility functions. However, in general, nothing can be said about the cost of an ESO portfolio vis-a-vis the cost of its individual components for aribtrary stock price beliefs. Under additional reasonable restrictions on the stock price process, we showed that for CARA utilities the cost due to ESO allocations will be sub-additive for an arbitrary portfolio of options. The restriction that is required on stock price processes for this to hold is in fact quite intuitive - it effectively says that ESOs should not have diversification benefits with respect to one another. Note that no such assumption was necessary while dealing with the case of options of the same type.

The model in in (2.1) is however not very suitable for practical use. This is because in the most general case, it will entail solving a high dimensional dynamic programming problem just to obtain the employee's exercise policy. While the CARA utility model has some nice properties and limits the state space to just unexercised options of various kinds, it does not scale well as the number of different types of options in the employee's portfolio increases. Also restricting usability are the several model parameters such as perceived or real world stock price dynamics and utility function specifications that will be needed to use this model.

Our task going forward is then to develop a practically useable model to incorporate the portfolio effect on exercise policies. The ultimate goal of this model should be to estimate the cost of the ESO portfolio for the employer. It will be also desirable to have an attribution of the portfolio cost to its different components. The latter is important because as we demonstrated, a new grant can affect the exercise and the cost of already issued options, if it is not anticipated by the employee. Further, if the employer wishes to transfer some of its ESO related liabilities to an external agency, such a cost split-up estimate will be crucial for the external agency.

This motivates us to look for an alternate modeling approach, which is the topic for Chapter 3.

Chapter 3

Tractable Models for Pricing Employee Stock Options

3.1 Introduction

In Chapter 2, we observed that ESOs should ideally be priced at a portfolio level, as an employee's decision whether to exercise an ESO is likely to be influenced by the presence of other ESOs in her portfolio. As a result, the approach of costing ESOs on a stand-alone basis is probably flawed and an ESO portfolio's cost may be very different from what one would obtain when each component is priced individually. Also, we observed that under reasonable assumptions, the portfolio cost function is sub-additive and hence piecemeal costing of ESOs will in general overestimate their cost.

The model presented in Chapter 2, though simple and appealing, is computationally demanding and involves several subjective parameters that would be difficult to obtain or even estimate using limited stock option exercise data that a company would typically have. These shortcomings motivate us to look in a relatively new direction to tackle the problem of modeling ESO exercises.

In our proposed model, the employee uses a risk-management based framework to make decisions about exercising ESOs. More specifically, an employee treats her ESO portfolio as an investment portfolio and rebalances it periodically to manage the associated risk that cannot be hedged. This rebalancing must occur through exercising of options in lieu of trading them, as she is unable to do so because of contractual reasons. We measure unhedgeable risk in terms of the short-term variance it introduces to the employee's portfolio. This is similar in spirit to the Markowitz mean-variance portfolio optimization framework, the oldest formal risk-management toolkit known to finance for this purpose. Such an approach has not been used in the literature before to the best of our knowledge and provides for a common framework to analyze employee behavior for all stock-related compensation. It gives a joint exercise model for all the ESOs in an employee's portfolio and one which can be solved very efficiently.

We find that the exercise behavior implied by this risk-management conisderation can be computed using a simple optimization routine. We show that the employee exercises an option, possibly partially, when a certain barrier function associated with the option, which we call the "delta-barrier function" falls below a stochastically varying threshold. Exercises based on this criteria are also shown to be consistent with the partial relative order or exercises between different options as shown in Lemma 2.5 in Chapter 2, where this order was established from very different considerations.

Our exercise model takes into account only the impact of unhedgeable risk on employee's exercise decisions. This, we believe is the most important factor that leads to early exercises. Some factors such as termination or quitting and vesting would also affect pricing. The model we propose is still flexible enough to incorporate these effects and with a little loss of computational tractability. Whenever applicable we will remark how this may be done in practice, but for the most part, we retain focus on demonstrating how coupled exercise behavior can be modeled in a tractable way.

We also propose another related model for ESO exercises, where the employee measures risk in terms of unhedgeable volatility of her portfolio. We show that this simplifying modeling assumption leads to an all-or-none exercise by the employees naturally, an unstated assumption in bulk of the ESO costing literature to date. This also implies a linear pricing rule for ESO portfolios. This model provides an attractive alternative for companies to estimate cost, when data is limited or for accounting purposes. Another benefit is that under this model, the payoff that employee gets upon exercising can be closely approximated by a martingale process. This allows us to get narrowly spaced analytical bounds on the resulting cost of the option to the employee.

A key advantage of these models is that it uses very few subjective parameters that can make cost estimation and calibration cumbersome. Apart from a few parameters, which are representative of employee's risk-averseness and can be calibrated from exercise data, the model refers to quantities that should be available from market data. This has an added benefit in terms of hedging the ESO cost risks, for parties who intend to do so.

The key contributions in this chapter are:

(a) We provide a concrete and tractable framework to price ESOs as a portfolio based on a risk-management approach.

(b) The exercise strategy that we derive is barrier based where the employee exercises the option when an explicit function of the stock and option parameters hits a threshold level. While, barrier based ESO exercise models, have been examined and analyzed in the literature before, most explicit barrier functions considered hitherto in the literature are ad-hoc and/or motivated primarily by the ease of computation/analysis. We provide an investment decision rationale for the proposed barrier function.

(c) Our exercise model, being a joint exercise model for a portfolio of options, predicts how different options will be exercised in relation to each other. Our model is parsimonious in terms of number of parameters making it easier to calibrate and verify for use in practice. The model also guarantees a pecking order of exercise between options.

(d) We also provide a simplified model, again inspired by a risk-management based framework, that naturally leads to an all-or-none exercise by the employee and a linear pricing rule for ESO portfolios. Further, for this model, we are able to derive analytical bounds that bracket the cost of the option within a fairly tight interval.

Chapter Layout

In Section 3.2, we describe the model used in this paper. In Section 3.3, the Risk Management framework used to model exercise behavior for ESOs is motivated. We the describe the myopic mean-variance optimization based decision model and derive the employee's exercise behavior for the same. We define the delta-barrier exercise function which plays an important role in the exercise consideration in Section 3.4 and also list its important properties. In Section 3.5, we describe the myopic mean-volatility decision framework. A special case of this framework makes the exercise decisions with respect to different ESOs in the portfolio independent of each other and the portfolio cost additive. Based on this, we derive analytical bounds on the cost of an ESO in Section 3.6. Finally, in Section 3.7 we summarize the results and the models presented to price ESOs.

3.2 Model

As remarked in Chapter 2, we believe ESOs are modeled more appropriately as investment instruments rather than consumption ones. We therefore take an investment-decision perspective for exercise models.

We model the problem of option exercise by employee as a portfolio optimization problem. The employee then naturally considers the decision to exercise an option or not in the larger context of her overall portfolio.

Employee's Portfolio

Similar to the model in Section 2.2 in Chapter 2, we assume that the employee holds a portfolio of N types of call options. We allow the portfolio to also include options that are expected to be issued in future. Let α_i denote the number of options of type i held by the employee. Without loss of generality, we assume that the 1^{st} N_1 of the N options have already been issued (but not necessarily vested). The i^{th} issued option where $1 \leq i \leq N_1$, is characterized by the tuple (K_i, V_i, T_i), with K_i denotes the strike, V_i the vesting time, and T_i the expiry time for option i. Note that this means that the employee can exercise an option of type i at any time between V_i and T_i. We can easily include restricted stock grants in this model, as a special type of call option by setting $K_i = 0$, and V_i as the time till which the stock must be held. For restricted stock, T_i would be set to ∞ or the end of the planning horizon for the employee.

The remaining $N_2 = N - N_1$ ESOs are the ones that have not been yet issued but are anticipated by the employee and hence would affect her decision process. These options, expected to be granted in the future, will not have an absolutely set strike. However, the level of moneyness at which these options will be granted is assumed to be certain and known, i.e., the strike is fixed as a multiple of the prevailing stock price on the date of issue[1]. We characterize an unissued option i, where $N_1 + 1 \leq i \leq N$, by the tuple (I_i, m_i, V_i, T_i). Here, I_i denotes the issue date of the i^{th} anticipated option, while m_i is the preset ratio $\frac{S_{I_i}}{K_i}$ and V_i and T_i denote respectively the vesting and expiry times of the option as before. In addition to the ESO portfolio, we also assume that the employee holds some non-option related wealth W_t. This wealth W_t may also have some correlation with the employing company's stock.

[1]For example, the future option could be known to at the money(ATM) at the time of issuance.

Stock Price Dynamics

The stock price S_t is assumed to follow a geometric Brownian Motion as in the Black-Scholes framework with volatility σ. We assume that there are no dividends to keep the analysis simple. The model may be easily extended to incorporate continuous dividends without much difficulty though.

If μ is the expected return on the stock, then

$$dS_t = \mu S_t dt + \sigma S_t dZ_t \ .$$

The risk-free rate r is also assumed to be non-stochastic. In the risk-neutral world,

$$dS_t = r S_t dt + \sigma S_t dZ_t^Q \ .$$

Let $C_{i,t}$ denote the Black-Scholes value of the option i and $E_{i,t} = (S_t - K_i)^+$, the payoff upon exercise at time t. Also let $\delta_{i,t}$ denote the "delta" of this option or the sensitivity of its Black-Scholes option value to S_t. For options that have already been issued, i.e., for $i : 1 \leq i \leq N_1$,

$$\begin{aligned}
E_{i,t} &= (S_t - K_i)^+ \ , \\
C_{i,t} &= S_t N(d_{i,t}) - K e^{-r(T_i - t)} N(d_{i,t} - \sigma\sqrt{T_i - t}) \ , \\
\delta_{i,t} &= N(d_{i,t}) \ , \\
\text{where, } d_{i,t} &= \frac{\ln(\frac{S_t}{K_i}) + r(T_i - t)}{\sigma\sqrt{T_i - t}} + \frac{1}{2}\sigma\sqrt{T_i - t} \ .
\end{aligned}$$

For those options for which the strike is yet to be set, i.e., for $j : N_1 + 1 \leq j \leq N$,

$$\begin{aligned}
C_{j,t} &= S_t \left(m_j N(\bar{d}_{j,t}) - e^{-r(T_j - I_j)} N(\bar{d}_{j,t} - \sigma\sqrt{T_j - I_j}) \right) \ , \qquad (3.1) \\
\delta_{j,t} &= m_j N(\bar{d}_{j,t}) - e^{-r(T_j - I_j)} N(\bar{d}_{j,t} - \sigma\sqrt{T_j - I_j}) \ , \\
\text{where, } \bar{d}_{j,t} &= \frac{\ln(m_j) + r(T_j - I_j)}{\sigma\sqrt{T_j - I_j}} + \frac{1}{2}\sigma\sqrt{T_j - I_j} \ .
\end{aligned}$$

From (3.1), we see that unissued options, in terms of their Black-Scholes value behave very much like restricted stock until the issue date (or the time until which the strike is fixed). We will appeal to this fact later in our analysis.

Exercise Behavior

We characterize an employee in terms of her risk-averseness and exposure to the company's stock[2].

In our setting, the employee can neither sell the option nor hedge it by selling short the underlying stock. These are obvious restrictions that apply to ESO holders. The only way the ESO holder can reduce the risk associated with the option payoff is by exercising it. We do allow for limited proxy hedging in this model. We will assume that the employee has no inside-information about stock prices that she can exploit for personal gains.

Also, in our model the employee does not explicitly take into account the possibility of being terminated or quitting voluntarily departure or being forced to exercise options to meet liquidity needs in future while making exercise related decisions. These effects can be incorporated by adding them as exogenous shocks to the model.

For our basic model, we ignore the "reset" and "reload" features that may be present in ESO terms. As we shall see later, our model can be easily augmented to account for these.

We also assume that the employee is taxed immediately on the proceeds received upon any exercise[3] and the applicable marginal tax rate remains constant throughout the planning horizon. Under this assumption, taxes do not make a structural difference to our problem and hence we assume without loss of generality that there are no taxes (we can effectively replace all grants α_i by $\alpha_i(1-f)$, f being the marginal tax rate, to account for the taxes and reduce the general problem to the case when there are no taxes, see also Aboody [2].).

The key determinant of the employee's decision process concerning ESO exercises

---

[2]Employees have a disproportionate exposure to the company's stock as a bulk of their current and future wealth is related to the stock performance of the company. Hence the entire risk associated with company's stock value (and not just the systemic risk, i.e., the market correlated risk) would matter to the employee.

[3]There are more than one types of ESOs and they can defer in their tax and accounting treatments. The most commonly issued type of ESOs are of the type called 'Non-qualified Stock Options' and are subject to the kind of tax treatment assumed here. Specifically, these options are treated in the same way as compensation but are taxable at exercise. This means that for these options, the value $(S_t - K)$ realized upon exercise is treated as income irrespective of whether the stocks received as a result of exercise are sold or not. Moreover, this value is assessed at the applicable income tax rate for the employee. Certain other options called the Incentive Stock Options differ in their tax treatment. For them, the tax is assessed only when the stocks received as a result of the exercise are actually sold. If the stock is held for a sufficiently long time (typically a year), upon selling, the gains realized (over the option strike price), if any, are assessed at the applicable capital gains tax rate. Since the latter is typically much less than marginal income tax rates, an employee holding an incentive stock option often has an incentive to convert the option to a stock and hold on to the stock purely due to tax reasons. In this paper, we do not treat the case of Incentive Stock Options.

is then her need to manage the risk associated with her ESO portfolio. We propose two decision making frameworks based on risk-management or portfolio optimization approach.

(a) Mean-Variance Optimization - Th mean variance optimization methodology proposed by Markowitz [90] is probably the most commonly approach used to model a risk - reward tradeoff. Here, we assume that the employee uses mean-variance optimization to adjust her ESO portfolio on a day-to-day basis.

(b) Mean-Volatility Optimization - This framework is also motivated by a risk management perspective. It has an advantage over the mean-variance optimization for pricing ESOs because under certain additional assumptions, it leads to an all-or-none exercise strategy for options, and makes the portfolio pricing problem separable. This allows us to estimate the cost of a portfolio of options as a sum of the costs of its components.

Unlike the utility based models in Chapter 2, the exercise models that will be presented in this chapter are not suitable to get an idea of the subjective value of ESO grants to the employee. The models that we propose only give us a handle on the employee's exercise behavior when granted a portfolio of options. Using this we can estimate the cost of this grant to the company issuing the options.

Upon exercise of ESOs, as the granting company creates fresh new stocks, there is also a dilution effect, see Black and Scholes [18]. We also neglect the effects of such "dilution" since, in practice, the number of ESOs exercised will be very small compared to the number of shares outstanding for a typical company, making this effect relatively insignificant. We compute the cost associated with an ESO as simply the expected value of payoffs received from the exercise under the risk neutral measure $\mathbb{Q}$.

We now describe our risk-management based exercise models in greater detail.

3.3 Risk Management Models

In this section, we propose a method to model employee exercise behavior, that is both computationally as well as conceptually simple. To motivate the model, we start with a discrete time setting, where the employee divides her time horizon into evenly spaced periods of length Δ for making exercise related decisions. As described below, a simplifying assumption that we make is that the employee manages her portfolio risk myopically.

- Instead of assuming trading restrictions on the ESOs for their entire duration, for deciding whether to exercise an option in the current period or not, we assume the employee weighs in the effect of trading restrictions on the options for only that period. This is the myopic enforcement of the trading constraints. Under this view, the employee can hedge the option position after time $t + \Delta$, and thus lock in its fair market value at the end of the period, which for type i option is $C_{i,t+\Delta}$. Suppose, at time t, the employee exercises $x_{i,t}$ of the $\alpha_{i,t}$ of the type i options she had at the beginning of the period. As a result, she will be left with $\alpha_{i,t+\Delta} = \alpha_{i,t} - x_{i,t}$ type i unexercised options at the beginning of the next period $t + \Delta$. Then the (time-discounted) value realized from the ESO portfolio is

$$V = \sum_{i=1}^{N} x_{i,t} E_{i,t} + \exp(-r\Delta) \sum_{i=1}^{N} \alpha_{i,t+\Delta} C_{i,t+\Delta} \ .$$

- The reason an employee may exercise her options earlier than she would have in absence of trading restrictions, is because of the risk related to the option position during the current period that she could not hedge. In the absence of this risk, the time t fair value of the total discounted payoff received via her strategy of exercising $x_{i,t}$ type i options at t is given by

$$\mathbb{E}^{\mathbb{Q}}[V] = \sum_{i=1}^{N} x_{i,t} E_{i,t} + \mathbb{E}^{\mathbb{Q}} \left[\exp(-r\Delta) \sum_{i=1}^{N} \alpha_{i,t+\Delta} C_{i,t+\Delta} \right] \tag{3.2}$$

$$= \sum_{i=1}^{N} x_{i,t} E_{i,t} + \sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t}) C_{i,t} \ . \tag{3.3}$$

However, because the employee cannot hedge the option position by short selling the stock, some of this value is at risk and the employee accounts for it by penalizing the fair value obtained in (3.3) by a term that is proportional to the variance in the realized value under her exercise strategy for the period.

$$\mathrm{var}^{Q}(V) = \mathrm{var}^{Q} \left(\exp(-r\Delta) \sum_{i=1}^{N} \alpha_{i,t+\Delta} C_{i,t+\Delta} \right) \ .$$

- This effectively means that at each period the employee solves a one-step mean-variance portfolio optimization, but using the risk-neutral distribution instead of subjective or actual probabilities[4]. We characterize the employee's risk aversion to unhedgeable risk by a parameter λ, which can be interpreted as the penalty imposed on unhedgeable risk per period. The employee's objective function is then to

$$\max \mathbb{E}^{\mathbb{Q}}[V] - \frac{\lambda}{2}\mathrm{var}^{\mathbb{Q}}(V) \ . \tag{3.4}$$

The employee's problem is thus a standard one stage portfolio optimization problem in a risk-neutral setting where the employee maximizes the unhedgeable variance adjusted value of her ESO portfolio at the end of the period. The trade-off that the employee faces is to exchange some ESOs for their intrinsic value to cut the portfolio variance at the expense of some of its "fair" value. The variance metric makes the option exercise problem a true portfolio problem as different entities of the portfolio are correlated.

- The different ESOs in the employee's portfolio are all correlated because their value depends on the employer's stock price. The myopic setting enables us to get a handle on quantifying this dependence as it can be now modeled through the option's "delta" which is its local sensitivity to the stock price S_t.

For small Δ, using the first order Taylor expansion, the discounted value of the call option i at $t + \Delta$ can be approximated as

$$e^{-r\Delta}C_{i,t+\Delta} \approx C_{i,t} + \delta_{i,t} \cdot (e^{-r\Delta}S_{t+\Delta} - S_t) \ .$$

Thus, the variance of V is given by

$$\mathrm{var}^{\mathbb{Q}}(V) = \left(\sum_{i=1}^{N}(\alpha_{i,t} - x_{i,t})\delta_{i,t}\right)^2 \sigma^2 S_t^2 \Delta \ .$$

The portfolio optimization problem that the employee wishes to solves is to maximize the short-term variance penalized value of her portfolio and can be stated as:

[4]Subjective or actual forward distributions pose a difficult estimation problem as well, making their use undesirable in a practical setting. Also under the assumption that the risk due to S_t is completely unhedgeable and hence uncorrelated with the market, the average rate of return μ on S_t should be equal to the risk-free rate r under the CAPM or APT theory. This will lead to an identical problem for the employee.

$$
\begin{aligned}
\max \quad & \sum_{i=1}^{N} x_{i,t}E_{i,t} + \sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})C_{i,t} - \\
& \frac{\lambda\Delta}{2}\left(\sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})\delta_{i,t}\right)^2 \sigma^2 S_t^2; \\
\text{s.t.} \quad & 0 \le x_{i,t} \le \alpha_{i,t} \ , \\
& x_{i,t} = 0 \text{ if } t < V_i \ . \qquad (3.5)
\end{aligned}
$$

One drawback of the above formulation is that it may recommend exercises for options that are not in the money. We fix this by adding a constraint that only in the money options can be considered for exercise. With this additional constraint, problem (3.5) is equivalent to the following quadratic optimization problem.

$$
\begin{aligned}
\max \quad & -\sum_{i=1}^{N} x_{i,t}(C_{i,t} - E_{i,t}) - \frac{\lambda\Delta}{2}\sigma^2 S_t^2 \left(\sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})\delta_{i,t}\right)^2 ; \qquad (3.6) \\
\text{s.t.} \quad & 0 \le x_{i,t} \le \alpha_{i,t} \ , \\
& x_{i,t} = 0 \text{ if } S_t \le K_i \text{ or } t < V_i \ . \qquad (3.7)
\end{aligned}
$$

The problem in (3.7), being a quadratic programming problem with linear constraints (LCQP), can be readily solved using standard convex optimization techniques. As we shall see in Section 3.4, it also has an intuitive solution in terms of how options are chosen for exercises. We first see how the model can be augmented to allow partial or proxy hedging of the employer stock related risk as well as incorporate other unhedgeable risk in the employee's portfolio without making any structural change in the problem (3.7) that the employee needs to solve.

Partial hedging with the myopic Mean-variance Optimization Model

Until now, we assumed the option risk is completely unhedgeable and the employee does not possess any illiquid non-option wealth. The myopic mean-variance maximization model can be naturally extended to include possible partial hedging by the employee as well as other assets in the employee's portfolio that cannot be hedged or traded. We now assume that there is a market tradeable factor Z_t^M that has a correlation ρ with Z_t^M. Thus, the variation of S_t can then be written as

$$\frac{dS_t}{S_t} = r \cdot dt + \sigma\rho dZ_t^M + \sigma\sqrt{1-\rho^2}\bar{Z^Q}_t$$

The factor Z_t^M denotes the tradeable market risk while $\bar{Z^Q}_t$ denotes the idiosyncratic risk associated with S_t and is uncorrelated with Z_t^M. Let Y_t denote an appropriately sized zero net cost portfolio, with a volatility σS_t in value and exposure to Z_t^M. Then,

$$dY_t = \sigma S_t dZ_t^M$$

in the risk neutral setting.

The employee may also have other unhedgeable and illiquid assets in her investment portfolio. These may also add to her exposure to the employer's stock price. An example would be pension fund contributions. We can incorporate these our model as well. Let W_t be the time t fair value of the other non-tradeable assets that the employee owns. Let β_0 denote the exposure of return on these assets to return on the stock S_t and σ_0 denote the uncorrelated risk in the portfolio. We again assume a factor representation for W_t,

$$\frac{dW_t}{W_t} = rdt + \beta_0\sigma dZ_t^Q + \sigma_0 dZ_t^Y \ .$$

To keep notation simple, we will assume that the factor corresponding to Z_t^Y is idiosyncratic and cannot be hedged against, but the general case can be dealt with in a similar fashion.

The problem (3.7) then becomes

$$\begin{aligned}
\max_{\{x_{i,t}\},y_t} \quad & W_t + \sum_{i=1}^{N}\alpha_{i,t}C_{i,t} - \sum_{i=1}^{N}x_{i,t}(C_{i,t}-E_{i,t}) - \frac{\lambda\Delta}{2}\Sigma^2 \ ; \\
\text{s.t.} \quad & \Sigma^2 = \sigma^2 S_t^2\left(\rho\left(\beta_0\frac{W_t}{S_t} + \sum_{i=1}^{N}(\alpha_{i,t}-x_{i,t})\delta_{i,t}\right) + y_t\right)^2 \\
& \quad + (1-\rho^2)\left(\beta_0\frac{W_t}{S_t} + \sum_{i=1}^{N}(\alpha_{i,t}-x_{i,t})\delta_{i,t}\right)^2 + \sigma_0^2 W_t^2 \ , \\
& 0 \le x_{i,t} \le \alpha_{i,t} \ , \\
& x_{i,t} = 0 \text{ if } S_t \le K_i \text{ or } t < V_i \ .
\end{aligned} \tag{3.8}$$

This is equivalent to the problem

$$
\begin{array}{ll}
\max\limits_{\{x_{i,t}\}} & -\sum\limits_{i=1}^{N} x_{i,t}(C_{i,t}-E_{i,t}) - \frac{\lambda\Delta}{2}(1-\rho^2)\sigma^2 S_t^2 \left(\beta_0 \frac{W_t}{S_t} + \sum\limits_{i=1}^{N}(\alpha_{i,t}-x_{i,t})\delta_{i,t}\right)^2 \;; \\
\text{s.t.} & 0 \le x_{i,t} \le \alpha_{i,t} \;, \\
 & x_{i,t} = 0 \text{ if } S_t \le K_{i,t} \text{ or } t < V_{i,t} \;.
\end{array} \tag{3.9}
$$

Thus in effect, we add $\beta_0 \frac{W_t}{S_t}$ of restricted stock to the employee's portfolio, and the problem (3.9) is structurally similar to (3.7). Note also that ungranted but anticipated options also behave like restricted stock with the exposure as given by (3.1) and these may also be incorporated in the model likewise.

The term λ that corresponds to the penalty applied to the one period variance in value that cannot be hedged against and the length of the period Δ for which the short-selling restriction is imposed appear in conjunction. In the limiting case one can make the short selling constraint less severe by reducing the time step Δ, while simultaneously increasing the penalty term for the unhedgeable risk resulting from this constraint, in such a way that the product $\lambda\Delta$ is held constant and thus come up with a continuous version of the myopic mean-variance optimization problem.

Another advantage of the model from a practical point of view is that the individual's risk averseness and ability to hedge as well as the discretization time-step are all captured by a single hyperparameter

$$\upsilon \triangleq \lambda\Delta(1-\rho^2).$$

In practice, υ should be obtained empirically from employee stock option exercise data.

We can thus consider the problem as stated in (3.10) below as the prototype of the employee's exercise problem.

$$
\begin{array}{ll}
\max & -\sum\limits_{i=1}^{N} x_{i,t}(C_{i,t}-E_{i,t}) - \frac{\upsilon}{2}\sigma^2 S_t^2 \left(\sum\limits_{i=1}^{N}(\alpha_{i,t}-x_{i,t})\delta_{i,t}\right)^2 \;; \\
\text{s.t.} & 0 \le x_{i,t} \le \alpha_{i,t} \;, \\
 & x_{i,t} = 0 \text{ if } S_t \le K_i \text{ or } t < V_i \;.
\end{array} \tag{3.10}
$$

In the next section, we examine some of the interesting properties of exercise policies that are obtained as solutions of (3.10).

3.4 Exercise Behavior under Myopic Mean Variance Optimizing Policy

We now consider the solution to the problem in (3.10).

For notational convenience, henceforth, we will drop the suffix t for each quantity in this section, with the understanding that the optimal policy x_i^* and the quantities discussed in this section are computed for each time t.

We now characterize the nature of the exercise policy that the employee will follow to maximize the myopic mean variance criterion. Let G denote the set of vested options that are in the money. G represents the set of options that are eligible for immediate exercise.

Lemma 3.1. *Let G denote the set of options that have vested and are in the money. If $i, j \in G$ and $\frac{C_i - E_i}{\sigma S \sigma_i} < \frac{C_j - E_j}{\sigma S \delta_j}$, then options of type i must be all exercised before any of type j is exercised.*

Proof. Let x_i^* denotes the optimal number of type i options to be exercised.

The "delta" of the portfolio at optimality can then be defined as

$$\delta^* \triangleq \sum_{i=1}^{N} (\alpha_i - x_i^*)\delta_i \ .$$

The first order optimality conditions for Problem (3.10) then require that:

$$\begin{aligned} C_i - E_i &\geq (\lambda \Delta \delta^* \sigma^2 S^2)\delta_i \text{ if } i \in G \text{ and } x_i^* = 0 \ , \\ C_i - E_i &= (\lambda \Delta \delta^* \sigma^2 S^2)\delta_i \text{ if } i \in G \text{ and } 0 < x_i^* < \alpha_i \ , \\ C_i - E_i &\leq (\lambda \Delta \delta^* \sigma^2 S^2)\delta_i \text{ if } i \in G \text{ and } x_i^* = \alpha_i \ . \end{aligned}$$

Hence, if $\frac{C_i - E_i}{\sigma S \delta_i} < \frac{C_j - E_j}{\sigma S \delta_j}$, then any option of type i cannot be exercised unless all options of type j have been exercised. □

Lemma 3.1 immediately motivates the following simple algorithm to solve the optimization problem (3.10).

Algorithm 3.1. Algorithm to find Optimal Exercise Policy

1. Initialize $x_i^* := 0$ for all i, $\delta^\dagger := \sum_{i=1}^N \alpha_i \delta_i$

2. Find G, the set of indices of exercisable options

3. Sort the options in G in increasing order of

$$\zeta_i \triangleq \frac{C_i - E_i}{\delta_i} \tag{3.11}$$

4. Loop through options in G in increasing order of ζ_js. For each option j

If $\zeta_j \geq \upsilon\sigma S\delta^\dagger$, then
Stop. The current values of all x_i^* are optimal and $\delta^\dagger = \delta^*$

else if $\zeta_j < \upsilon S\delta^\dagger$ but $\zeta_j \geq \upsilon\sigma S(\delta^\dagger - \alpha_j\delta_j)$, then
Set $x_j^* := \frac{\delta^\dagger - \frac{\zeta_j}{\upsilon\sigma S}}{\delta_j}$ and then $\delta^\dagger := \frac{\zeta_j}{\upsilon\sigma S}$
Stop. The current values of all x_i^* and δ^* are optimal and $\delta^\dagger = \delta^*$

else
Set $x_j^* := \alpha_j$ and $\delta^\dagger := \delta^\dagger - \alpha_j\delta_j$

Note, that the portfolio's delta $\delta^\dagger$ (and hence its instantaneous variance) as well as its value monotonically decrease throughout the algorithm. Algorithm 3.1 works by decreasing the delta of the portfolio until the myopic mean-variance objective reaches its optimal value.

In general, the myopic mean-variance maximization model would recommend both partial as well as complete exercises. As can be seen from Algorithm 3.1, at any time, the model recommends an all or none exercise policy for all but one option. However, in practice as the stock price varies continuously, at any given time the method would almost always recommend no exercise or a partial exercise of one particular option.

In the next section we see that the myopic mean-variance based exercise strategy guarantees the same logical order of exercises between options as the one we established in Chapter 2 from a utility maximization consideration.

The delta-barrier function and its properties

The quantity ζ_j as defined in (3.11) plays an important role in determining and sequencing options for exercise. ζ_j is a function of the option's strike and time to expiry. It also depends on the current stock price. ζ_j can thus be interpreted as an exercise barrier function associated with each option.

Definition 3.1. *The "delta-barrier" function of an option with strike K and time to expiry T is defined as*

$$B(S,K,T) \;=\; \frac{C(S,K,T)-(S-K)^+}{\sigma S\delta(S,K,T)} \; . \tag{3.12}$$

Here $\delta(S,K,T)$ is the delta of a call option with expiry T and strike K.

The delta-barrier function admits an intuitive interpretation as a pseudo-Sharpe ratio for the option. (Note that the delta-barrier is a ratio of the option premium (C-E) over the short-term volatility of this premium.) Algorithm 3.1 shows that the employee always exercises options in an increasing order of their pseudo-sharpe ratios.

Definition 3.2. *An exercise strategy is a "delta-barrier based exercise strategy" if it satisfies the following property:*

There exists a global, possibly stochastic, threshold $\nu \geq 0$ such that the strategy recommends to (partially or completely) exercise an option with strike K and time to expiry T if and only if

(a) the option is in the money and

(b) $B(S,K,T) \leq \nu$.

The myopic mean-variance optimizing exercise strategy is thus a "delta-barrier" based exercise strategy. The delta-barrier function has certain monotonicity properties with respect to its arguments. As a consequence, a delta-barrier based exercise strategy guarantees the same pecking order in which options will be exercised as we obtained from a different argument in Lemma 2.5 in Chapter 2. We demonstrate this through a series of results on the properties of the delta-barrier function.

Lemma 3.2. *The function $\tilde{B}(S,K,T) \triangleq \frac{C(S,K,T)-S+K}{\sigma S\delta(S,K,T)}$ is monotone in S and decreases with the same.*

Proof. We have,

$$\begin{aligned} C &= SN(d) - Ke^{-r\tau}N(d-\sigma\sqrt{T}) \text{ , where ,} \\ d &= \frac{\ln(\frac{S}{K})}{\sigma\sqrt{T}} + \left(\frac{r}{\sigma}+\frac{\sigma}{2}\right)\sqrt{T}. \end{aligned}$$

Note that both the numerator and the denominator of $\tilde{B}(S,K,\tau)$ are always non-negative as $C \geq (S-K)^+ \geq (S-K)$, C being the Black-Scholes value of the option.

Further the denominator $\sigma SN(d)$ is an increasing function of S as d is increasing in S and so is $N(\cdot)$. The numerator, on the other hand, is a decreasing function of S as

$$\frac{\partial}{\partial S}(C - S + K) \quad = \quad N(d) - 1 < 0 \ .$$

It then follows that $\tilde{B}(S, K, t)$ is decreasing in S . □

Corollary 3.1. *Between two ESOs having the same expiry, a delta-barrier based exercise strategy first exercises the option with a higher value of the ratio* $\frac{S}{K}$*; i.e., it exercises first the option that is deeper in the money.*

Proof. The proof follows directly from Lemma 3.2. Observe that for in-the-money options $\tilde{B}(S, K, T) = B(S, K, T)$. Moreover, $B(\cdot)$ (and $\tilde{B}(\cdot)$) are homogeneous in the sense that $B(S, K, T) = B(\frac{S}{K}, 1, T)$. It then follows from the definition of a delta-based exercise strategy that between two options having the same expiry, such a strategy would first exercise an option with the lower strike. □

Lemma 3.3. *When the option is in the money, i.e.,* $S \geq K$*; the delta-barrier function* $B(S, K, T)$ *is increasing in* T*.*

Proof. This property requires a bit more effort to prove. We define the function

$$b(m, y) \quad = \quad B(mK, K, y^2) = B(m, 1, y^2) \ .$$

It follows that the delta-barrier function is increasing in T if $b(m, y)$ is increasing in y. Note that we operate under the case $m > 1$.

Now,

$$\begin{aligned} b(m, y) &= \frac{1}{m\sigma}\frac{C - E}{N(d)} \ . \\ \Rightarrow \frac{\partial b}{\partial y} &= \frac{1}{m\sigma}\left(\frac{\Theta}{N(d)} + \frac{(C - E)\Phi(d)}{(N(d))^2}\left(\frac{\ln m}{\sigma y^2} - \left(\frac{r}{\sigma} + \frac{\sigma}{2}\right)\right)\right) \ , \\ \text{where, } \Theta &= m\Phi(d)\sigma + 2ryN(d - \sigma y)e^{-ry^2} . \end{aligned}$$

Note that $\Theta \geq 0$ and $(C - E) \geq 0$. Thus, it follows that if

$$\frac{\ln m}{\sigma y^2} - \left(\frac{r}{\sigma} + \frac{\sigma}{2}\right) \geq 0 \ ,$$

then the delta-barrier function is increasing in τ. Hence, we need only consider the case when

$$\frac{\ln m}{\sigma y^2} - \left(\frac{r}{\sigma} + \frac{\sigma}{2}\right) < 0 \; . \tag{3.13}$$

In this case,

$$\frac{\partial b}{\partial y} = \frac{\Phi(d)}{m\sigma(N(d))^2}\left(mN(d)\sigma + 2rye^{-ry^2}\frac{N(d)N(d-\sigma y)}{\Phi(d)} + (C-E)\left(\frac{\ln m}{\sigma y^2} - \frac{r}{\sigma} - \frac{\sigma}{2}\right)\right) \; .$$

The quantity in the brackets,

$$Q = mN(d)\sigma + 2rye^{-ry^2}\frac{N(d)N(d-\sigma y)}{\Phi(d)} + (C-E)\left(\frac{\ln m}{\sigma y^2} - \frac{r}{\sigma} - \frac{\sigma}{2}\right) \; ,$$

is increasing in m when $m \geq 1$ and the condition (3.13) is satisfied. This can be verified easily by taking 1^{st} order derivatives and noting that $N(d), N(d-\sigma y)$ should increase with m, while the quantity $C-E$ decreases with m and so does $\Phi(d)$ when $m > 1$. It then follows that if we can show that $Q \geq 0$ for $m = 1$, then $\frac{\partial b}{\partial y}$ should be positive for all $m \geq 1$. At $m = 1$, we have $Q = Q_0$ with

$$\begin{aligned} Q_0 &= N(d)\sigma + 2rye^{-ry^2}\frac{N(d)N(d-\sigma y)}{\Phi(d)} - \left(N(d) - e^{-ry^2}N(d-\sigma y)\right)\left(\frac{r}{\sigma} + \frac{\sigma}{2}\right) \\ \Rightarrow yQ_0 &= -N(d)\left(\frac{r}{\sigma} - \frac{\sigma}{2}\right)y + e^{-ry^2}N(d-\sigma y)\left(\left(\frac{r}{\sigma} + \frac{\sigma}{2}\right)y + 2ry^2\frac{N(d)}{\Phi(d)}\right) \; . \end{aligned}$$

Again, if $f \triangleq (\frac{r}{\sigma} - \frac{\sigma}{2})y \leq 0$, it will trivially follow that $Q_0 > 0$, hence we consider the case when $f > 0$. Note that $2ry^2 = d^2 - f^2$ when $m = 1$. Thus, we have

$$\begin{aligned} yQ_0 &= -N(d)f + e^{-\frac{d^2-f^2}{2}}N(f)\left(d + \frac{N(d)}{\Phi(d)}(d^2 - f^2)\right) \\ &= -N(d)f + \frac{\Phi(d)N(f)}{\Phi(f)}d + \frac{N(d)N(f)}{\Phi(f)}(d^2 - f^2) \\ &= \frac{N(d)N(f)}{\Phi(f)}\left(\left(\frac{\Phi(d)}{N(d)}d + d^2\right) - \left(\frac{\Phi(f)}{N(f)}f + f^2\right)\right) \; . \end{aligned}$$

Since $d \geq f$, it will follow that $yQ_0 \geq 0$ if we can show that the function

$$f(x) \triangleq \frac{\Phi(x)}{N(x)}x + x^2$$

is increasing in x from $x \geq 0$. This in turn follows from the fact that the function $g(x) \triangleq \frac{\Phi(x)}{N(x)} + x \geq 0$ for $x \geq 0$ and is also increasing in x, a fact proved in Lemma

B.1 in Appendix B. □

Corollary 3.2. *Between options having the same strike, any delta-barrier based exercise strategy completely exercises the one with the shorter before exercising an option with a longer expiry.*

Proof. This is a direct consequence of the Lemma 3.3. □

To summarize, a delta-barrier based exercise strategy naturally satisfies the following desirable properties that one would expect of a rational exercise behavior as obtained in Lemma 2.5 in Chapter 2.

1. An option which is at the money at expiry will be exercised. This is because the barrier function B is zero for such options and by definition, the threshold $\upsilon > 0$.

2. Between two options that are in the money and have the same expiry, the strategy always exercises the option that has the lower strike first.

3. Between two options that are in the money and have the same strike, the strategy always exercises the option that has the shorter expiry first.

Figure 3-1 shows how the delta-barrier function varies with moneyness i.e., $\frac{S}{K}$ and time to maturity(for $r = 0.05$ and $\sigma = 0.2$ (annualized)).

Under the myopic mean-variance optimization framework, the option exercise problem for the employee is a simple quadratic problem. This enables us to obtain the cost of an entire ESO portfolio by simply simulating a price process path for the stock. Also, the same simulations can be used to allocate the cost of the portfolio amongst its different components by simply keeping tracks of the exercise along each sample path.

The exercise threshold ν for the delta-barrier based exercise strategy that we derived from the mean-variance optimization framework is in general stochastic and depends on the overall portfolio structure. This makes the option pricing problem non-linear. While this, as we argued in Chapter 2, represents a truer picture of ESO portfolio cost, a linear model is simple and more interpretable to report. In the next section, we see how we can derive a linear model for pricing ESOs by a suitable modification of the risk management framework presented in this section.

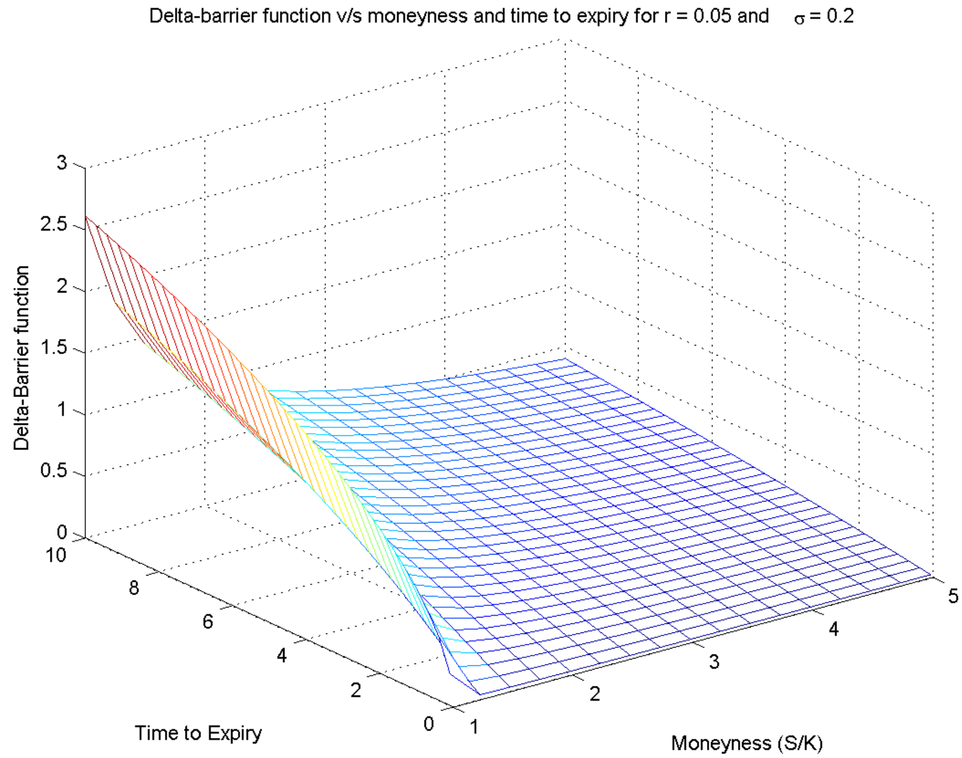

Figure 3-1: Variation of Delta-Barrier function w.r.t. moneyness and time to maturity.

3.5 Myopic Mean-Volatility Based Exercise Model

The myopic mean-variance based exercise model allows one to model parsimoniously the effects of partial exercises and coupling of option exercises. Under this model the ESO portfolio cost will be non-linear. Based on our analysis in Chapter 2, such non-linear models are likely to present a truer picture of the ESO costs. However they might be unsuitable for accounting and reporting purposes. Moreover costs still have to be recovered via simulation of the stock price, which does not provide a great insight to the price process. These properties result from using a variance based optimization scheme.

We now propose a variant of the myopic mean-variance based optimization model, the "myopic mean-volatility optimization" model, where the employee penalizes the instantaneous volatility in his/her portfolio instead of the instantaneous variance. As a side-note, for normal distributions, the problem of penalizing the volatility instead of variance in the utility function is akin to Value at Risk (VaR) or Conditional Value at Risk (CVaR) constrained optimizations, that organizations often use to manage their risk.

We will see that with certain additional assumptions, the problem of finding the optimal exericse policy for ESOs becomes separable for this model. An immediate related consequence is options are exercised as all or none and we get a linear pricing rule for options, where the cost of an option portfolio can be obtained as the sum of its parts. Another surprising and extremely useful benefit of the model is that we get simple analytical formulae for bounds on the cost of an ESO.

The only difference that we make in our model is that the employee now penalizes the fair value of her portfolio by its short-term unhedgeable volatility rather than variance. Thus her objective as described in (3.4) now changes to

$$\max \mathbb{E}^{\mathbb{Q}}[V] - \chi\sqrt{\mathrm{var}^{\mathbb{Q}}(V)} \,. \tag{3.14}$$

Here χ is a risk-averseness parameter, and Δ as before is the length of a time-period.

The problem (3.9) then becomes:

$$\begin{aligned}
\max_{\{x_{i,t}\},y_t} \quad & W_t + \sum_{i=1}^{N} \alpha_{i,t} C_{i,t} - \sum_{i=1}^{N} x_{i,t}(C_{i,t} - E_{i,t}) - \chi\sqrt{\Delta}\Sigma \ ; \\
\text{s.t.} \quad & \Sigma^2 \ = \ \sigma^2 S_t^2 \left(\rho \left(\beta_0 \frac{W_t}{S_t} + \sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})\delta_{i,t} \right) + y_t \right)^2 \\
& \qquad + (1-\rho^2) \left(\beta_0 \frac{W_t}{S_t} + \sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})\delta_{i,t} \right)^2 + \sigma_0^2 W_t^2 \ , \\
& 0 \le x_{i,t} \le \alpha_{i,t} \ , \\
& x_{i,t} = 0 \text{ if } S_t \le K_i \text{ or } t < V_i \ .
\end{aligned} \tag{3.15}$$

For the optimal choice of y_t, the problem in (3.15) reduces to

$$\begin{aligned}
\max_{\{x_{i,t}\}} \quad & -\sum_{i=1}^{N} x_{i,t}(C_{i,t} - E_{i,t}) - \chi\sqrt{\Delta}\Sigma \ ; \\
\text{s.t.} \quad & \Sigma^2 \ = \ \sigma^2 S_t^2 (1-\rho^2) \left(\beta_0 \frac{W_t}{S_t} + \sum_{i=1}^{N} (\alpha_{i,t} - x_{i,t})\delta_{i,t} \right)^2 + \sigma_0^2 W_t^2 \ , \\
& 0 \le x_{i,t} \le \alpha_{i,t} \ , \\
& x_{i,t} = 0 \text{ if } S_t \le K_i \text{ or } t < V_i \ .
\end{aligned} \tag{3.16}$$

Just as in the model presented in Section 3.3, the term χ that corresponds to the penalty applied to the one period volatility in value due to hedging restrictions and the length of the period Δ for which the short-selling restriction is imposed appear in conjunction. Again, in the limiting case one can make the short selling constraint less severe by reducing the time step Δ, while simultaneously increasing the penalty term for the unhedgeable risk resulting from this constraint in such a way that the product $\chi\sqrt{\Delta}$ is held constant and thus effectively use the myopic mean-volatility optimization problem in (3.16) in a continuous time setting.

We now show that the optimal exercise policy for (3.16) is a delta-barrier based exercise policy. Moreover, if the only source of unhedgeable risk in the employee's portfolio comes through the exposure to the employer's stock, then the threshold in the delta-barrier based exercise policy becomes a constant and the problem becomes separable.

Proposition 3.1. *The optimal exercise strategy corresponding to the problem in (3.16) is a "delta-barrier" based strategy. Also if $\sigma_0 = 0$, then the exercise threshold becomes fixed and the optimal exercise strategies for different option types become decoupled.*

Proof. We will again drop the suffix t, to ease notation. Let x_i^* denote the optimal number of type i options to be exercised. We define the quantity δ^* and the variance at optimality as follows

$$\begin{aligned}
\delta^* &= \beta_0 \frac{W}{S} + \sum_{i=1}^{N} (\alpha_i - x_i^*)\delta_i \ ; \\
(\Sigma^*)^2 &= (\delta^*)^2 (1-\rho^2)\sigma^2 S^2 + \sigma_0^2 W^2 \ .
\end{aligned}$$

Let G denote the set of options that are vested and are in the money, i.e., can be considered for exercise. Then the first order optimality conditions imply, at optimality:

$$\begin{aligned}
C_i - E_i &\leq \chi\sqrt{\Delta}\frac{\beta^*(1-\rho^2)\sigma^2 S^2}{\Sigma^*}\delta_i \text{ if } i \in G \text{ and } x_i^* = 0 \ , \\
C_i - E_i &= \chi\sqrt{\Delta}\frac{\beta^*(1-\rho^2)\sigma^2 S^2}{\Sigma^*}\delta_i \text{ if } i \in G \text{ and } 0 < x_i^* < \alpha_i \ , \\
C_i - E_i &\geq \chi\sqrt{\Delta}\frac{\beta^*(1-\rho^2)\sigma^2 S^2}{\Sigma^*}\delta_i \text{ if } i \in G \text{ and } 0 < x_i^* < \alpha_i \ .
\end{aligned}$$

Hence there exists a threshold ν such that an option i is exercised only if it is in the money and its delta barrier function $\xi_i = \frac{C_i - E_i}{\sigma S \delta_i}$ satisfies $\xi_i < \nu$.

Now consider the case, when $\sigma_0 = 0$. Problem (3.16) (with the time t suffixes dropped) in this case becomes

$$\begin{aligned}
\max_{\{x_i\}} \quad & -\sum_{i=1}^{N} x_i(C_i - E_i) - \chi\sqrt{\Delta}\sqrt{1-\rho^2}\left(\beta_0 \frac{W}{S} + \sum_{i=1}^{N}(\alpha_i - x_i)\delta_i\right)\sigma S \ ; \\
\text{s.t.} \quad & 0 \le x_i \le \alpha_i \ , \\
& x_i = 0 \text{ if } i \notin G. \qquad (3.17)
\end{aligned}$$

This is a Linear Optimization problem. Let

$$\nu \ = \ \chi\sqrt{\Delta}\sqrt{1-\rho^2} \ .$$

Then a solution to (3.17) is given by

$$\begin{aligned}
x_i^* &= \alpha_i \text{ if } \xi_i = \frac{C_i - E_i}{\sigma S \delta_i} < \nu \text{ and } i \in G \ , \\
x_i^* &= 0 \text{ if } \xi_i = \frac{C_i - E_i}{\sigma S \delta_i} \ge \nu \text{ or } i \notin G \ . \qquad (3.18)
\end{aligned}$$

Thus, the exercise policy in this case is an all or none policy. Moreover, the threshold $\nu = \chi\sqrt{\Delta}\sqrt{1-\rho^2}$ is a constant and independent of the portfolio structure. This shows that the optimal exercise policy for each option can be determined independently. □

As the exercise policy for (3.16) is delta-barrier based, the rational exercise order properties stated in Lemma 2.5 in Chapter 2 will hold for this version of the employee's problem as well. With the exercise policy obtained through solving (3.18), the ESO costing problem is not only decoupled, i.e., the problem can be solved for each option separately but also becomes linear i.e., the cost of the portfolio of ESOs now will simply be a sum of the cost of its components.

We term the exercise strategy obtained in the special case when $\sigma_0 = 0$, as described in (3.18) as a "fixed threshold delta-barrier based exercise strategy". Again we observe that the model parameters, χ, Δ and ρ can be combined in a single hyperparameter ν given by

$$\nu \ = \ \chi\sqrt{\Delta}\sqrt{1-\rho^2} \ . \qquad (3.19)$$

ν can be interpreted as the target pseudo Sharpe-ratio, desired by the employee of an option to keep it in her portfolio. Intuitively, this strategy recommends exercise to the employee when the instantaneous volatility in the option value becomes significantly high compared to the incremental value of holding the option vis-a-vis exercising it, i.e., its pseudo Sharpe-ratio drops below a fixed threshold. Only ν matters for pricing options and this is the parameter that should be empirically calibrated for implementation.

Lemma 3.2 showed that the delta-barrier function is decreasing in S (for $S > K$). The barrier function is also homogeneous in S and K. Thus, given an exercise threshold ν, there exists a critical multiple M_t such that, if the ratio $\frac{S_t}{K}$ of the stock price to the option's strike exceeds this multiple then the option will be exercised. In our model, M depends only on the option's time to expiry.

Figure 3.5 shows the critical exercise price by strike multiple as a function of time to expiry for various values of ν under the fixed-threshold delta-barrier based exercise strategy (at 20% annualized volatility in stock returns σ and risk free-rate $r = 5\%$ annualized). The relationship is almost linear, especially when option has some time to expiry.

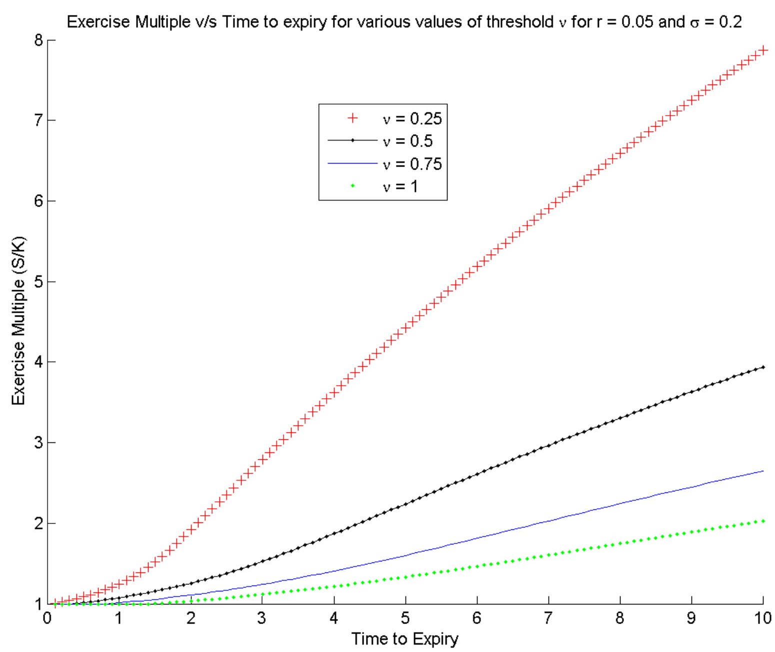

Figure 3-2: Exercise multiple v/s and time to expiry for different values of ν.

Barrier policies where the employee exercises her ESOs when the stock price to strike multiple exceeds a barrier function have been suggested in the literature before. For example, in the Hull and White [72] model, the employee exercises the stock price hits a certain multiple. Note that this model does not consider the impact of time to expiry on the option exercises. Another such method is proposed in Cvitanic, Weiner and Zapatero [47], in which the employee exercises an ESO whenever the ratio of the stock price to the option's strike exceeds a target multiple that declines exponentially with the residual life of the option. This target multiple also has a discontinuity at the option expiry time. These barrier based exercise strategies are ad-hoc and motivated primarily from analytical considerations. The fixed threshold delta-barrier based exercise strategy proposed here, in (3.18), was derived from a simplified but explicit model of employee behavior.

In addition, as we shall see in the next section, this model of exercise behavior also offers computational advantages. In particular, we can obtain tight analytical bounds on the implied cost of an ESO for the employer.

3.6 Pricing ESOs under the fixed threshold delta-barrier exercise strategy

In the previous section, we saw that under the myopic mean-volatility optimization framework, when the only risk that the employee is unable to hedge is the one corresponding to her employer's stock, the portfolio cost of ESOs becomes linear and can be obtained by adding the cost of all constituent options. Moreover each ESO in the portfolio is exercised independently of the other and according to the fixed threshold delta-barrier based exercise strategy as described by (3.18). In this section, we attempt to price analytically the cost C^{ν} of the ESO for this exercise strategy, given an exercise threshold ν. Suppose at time $t = 0$, the employee has an option with strike K and maturity T. Let the prevailing stock price be S_0. We first under the case that the *option has already vested*, i.e., the employee can exercise the option right away if she desires to do so. To recap, under the fixed threshold delta-barrier based exercise policy, the employee exercises the option when the following conditions are satisfied

- The option is in the money, i.e., $S_t > K$,
- The delta-barrier function $B_t = \frac{C_t - E_t}{\sigma \delta_t S_t}$ is less than or equal to the exercise threshold ν.

Because, the stock price process is continuous, the barrier function B_t, which is a continuous function of the stock price S_t will also follow a continuous path. Hence if $B_0 > \nu$, and if the option is exercised at a subsequent time τ; we can expect the following condition to be satisfied at the time of exercise τ.

$$\begin{aligned} B_\tau &= \nu \\ \Rightarrow C_\tau - E_\tau &= \nu\sigma S_\tau \delta_\tau \ . \end{aligned} \tag{3.20}$$

Unfortunately, condition (3.20) does not hold strictly and there is a probability that $B_t < \nu$ at exercise. To see how this can happen, consider a sample path where the option remains out of money up to time τ. The delta-barrier function however could still decrease and it might be the case that $B_\tau < \nu$. If now the option moves into the money, the delta-barrier function will decrease even further and under the fixed threshold exercise policy, the option must be exercised at this point. The condition (3.20) will be violated in this case. Fortunately, such an event can occur only with a small probability and close to option expiry. We can in fact show that (3.20) is violated if and only if the option is exercised in the window $(T - T_N, T)$, where T_N is a constant that depends on the risk free-rate r, volatility σ and the exercise threshold ν. Also when the violation does occur it is small in magnitude. As a result, we can derive upper and lower bounds on the cost of the option to the issuer which will be fairly close to each other and thus give a fair indication of the actual cost to the employee.

Lemma 3.4. *Given an exercise threshold ν, risk-free rate r and volatility level σ, there exists a unique critical expiry T_N such that*

- *$B(S,t) \geq \nu$ for all out of the money options if $t < T - T_N$ and*
- *If $t > T - T_N$ and the option has not yet been exercised, then the option will be exercised as soon as it gets in the money.*

Proof. Let us define T_N as that expiry for which $B(K,K,T-T_N) = \nu$. Note that

$$B(K,K,T-T_N) = 1 - e^{-rT_N} \frac{N\left(\left(\frac{r}{\sigma} - \frac{\sigma}{2}\right)\sqrt{T_N}\right)}{\sigma N\left(\left(\frac{r}{\sigma} + \frac{\sigma}{2}\right)\sqrt{T_N}\right)} \ .$$

and thus depends neither on K nor on T but only on r, σ and ν. From Lemma 3.3, it follows then $B(K,K,t) > \nu$ for $t < T_N$. Also for $t = T - T_N$, the barrier function for ATM stock price is exactly ν. Hence, for $t > T - T_N$ from Lemma 3.3 and Lemma 3.2, $B(S,t) \leq \nu$ for all values of $S \geq K$. This completes the proof. □

Lower bound on the cost C^ν of a vested ESO

Let τ be the time the barrier $B(t)$ first hits the threshold ν before time T and that the option is in the money, i.e., the exercise time of the option. If the option expires unexercised, then we set $\tau = \infty$. Let $\tau^* = \min\{\tau, T\}$. Then τ^* is a stopping time.

Lemma 3.5.

$$\mathbb{E}^Q[e^{-r\tau}(C_\tau - E_\tau)] \quad = \quad C_0 - C_0^\nu \ , \tag{3.21}$$

where C_t^ν denotes the cost function of issuing the ESO (as incurred by the employer), and C_t denotes the Black-Scholes price of the Call option that has the same terms as the ESO.

Proof. Let us define A as the event that the employee exercises her option during its lifetime, i.e., $\tau = \tau^*$.

$$\begin{aligned}
\mathbb{E}^Q[e^{-r\tau^*}C_{\tau^*}] &= P(\tau^* = \tau)\mathbb{E}^Q[e^{-r\tau^*}C_{\tau^*}|\tau = \tau^*] + P(\tau^* \neq \tau)\mathbb{E}^Q[e^{-r\tau^*}C_{\tau^*}|\tau \neq \tau^*] \\
&= P(\tau < \infty)\mathbb{E}^Q[e^{-r\tau}C_\tau|\tau < \infty] + P(A^c)\mathbb{E}^Q[e^{-rT}C_T|A^c] \ .
\end{aligned}$$

Note A^c is the event $\tau \neq \tau^*$. Now, as an expiring in the money option is always exercised under a delta-barrier based exercise policy. $A^c \Rightarrow S_T < K$ and hence in this case, the ESO was never exercised i.e., $\tau^* = T$. Also, since $S_T < K$, $C_T = 0$. It then follows that

$$\mathbb{E}^Q[e^{-r\tau^*}C_{\tau^*}] \quad = \quad \mathbb{E}^Q[e^{-r\tau}C_\tau] \ .$$

As C_t is the price of a tradeable asset, $e^{-rt}C_t$ is a martingale. Since τ^* is a stopping time, by the optional stopping theorem, $\mathbb{E}^Q[e^{-r\tau^*}C_{\tau^*}] = C_0$.

Finally, note that $\mathbb{E}^Q[e^{-r\tau}E_\tau] = C_0^\nu$ by definition. Hence, we conclude that

$$\mathbb{E}^Q[e^{-r\tau}(C(\tau) - E(\tau))] \quad = C_0 - C_0^\nu \ .$$

□

Lemma 3.6.

$$\mathbb{E}^Q[e^{-r\tau}\delta(\tau)S_\tau] \quad = \quad S_0\delta_0 \ . \tag{3.22}$$

Proof. Again, let A be the event that the employee exercises her option during its lifetime, i.e., $\tau = \tau^*$. We have

$$\begin{aligned}
\mathbb{E}^Q[e^{-r\tau^*} S_{\tau^*}\delta_{\tau^*}] &= P(\tau^* = \tau)\mathbb{E}^Q[e^{-r\tau^*} S_{\tau^*}\delta_{\tau^*}|\tau = \tau^*] + P(\tau^* \neq \tau)\mathbb{E}^Q[e^{-r\tau^*} S_{\tau^*}\delta_{\tau^*}|\tau \neq \tau^*] \\
&= P(\tau < \infty)\mathbb{E}^Q[e^{-r\tau} S(\tau)\delta(\tau)|\tau < \infty] + P(A^c)\mathbb{E}^Q[e^{-rT} S_T\delta_T|A^c] \ .
\end{aligned}$$

As an expiring in the money option will be exercised, $A^c \Rightarrow S_T < K$. Hence conditioned on A^c, $\delta_T = 0$. This is because $\delta(t) = N\left(\frac{\ln(\frac{S_t}{K})+r(T-t)}{\sigma\sqrt{T-t}} + \frac{1}{2}\sigma\sqrt{T-t}\right)$. Hence if $S_T < K$ $\delta(T) = 0$ (in limit). It then follows,

$$\mathbb{E}^Q[e^{-r\tau} S_\tau\delta_\tau] = \mathbb{E}^Q[e^{-r\tau^*} S_{\tau^*}\delta_{\tau^*}] \ .$$

Now,

$$\begin{aligned}
e^{-rt} S_t\delta_t &= e^{-rt}(C_t + Ke^{-r(T-t)}N(d_t + \sigma\sqrt{T-t})) \\
&= e^{-rt}\left(\mathbb{E}_t^Q\left[e^{-r(T-t)}(S_T - K)^+\right] + \mathbb{E}_t^Q\left[e^{-r(T-t)}K \cdot 1_{\{S_T > K\}}\right]\right) \\
&= \mathbb{E}_t^Q[e^{-rT} S_T \cdot 1_{\{S_T > K\}}] \ .
\end{aligned}$$

Thus, the function $e^{-rt}S_t\delta_t$ is in fact a martingale under the risk neutral measure $\mathbf{Q}$. Since, τ^* is a stopping time it then follows

$$\mathbb{E}^Q[e^{-r\tau^*} S_{\tau^*}\delta_{\tau^*}] = S_0\delta_0 \ .$$

□

Lemma 3.7. *Under the fixed threshold delta-barrier exercise policy, the cost of the ESO to the employer is lower bounded by*

$$\begin{aligned}
C_0^\nu &\geq C_0^{\nu-} \triangleq \max(0, C_0 - \nu\sigma S_0\delta_0) \\
&= S_0 N(d_0)(1 - \nu\sigma) - Ke^{-rT}N(d_0 - \sigma\sqrt{T}) \ . \qquad (3.23)
\end{aligned}$$

Proof. Under the fixed threshold delta-barrier exercise policy, the following inequality always holds.

$$\begin{aligned}
B_\tau &\leq \nu \\
\Rightarrow E_\tau &\geq C_\tau - \nu\sigma S_\tau\delta_\tau \ . \\
\text{Hence, } \mathbb{E}^Q[e^{-r\tau}E_\tau] &\geq \mathbb{E}[e^{-r\tau}C_\tau] - \nu\sigma\mathbb{E}[e^{-r\tau}S_\tau\delta_\tau] \\
\Rightarrow C_0^\nu &\geq C_0 - \nu\sigma S_0\delta_0 \ .
\end{aligned}$$

□

Upper bound on the cost C_0^ν of a vested ESO

Lemma 3.8. *If $T > T_N$ then,*

$$C_0^\nu \le C(S, K, T - T_N) \ ,$$

where $C(S, K, T - T_N)$ denotes the price of the call option with same strike K but maturity $T - T_N$.

Proof. From Lemma 3.4, it follows that

- If the option is in the money at time $T - T_N$, then it will be exercised at that time if it already has not been,
- If the option is exercised after time $T - T_N$, then it was out of money at time $T - T_N$. It will then be exercised as soon as the option gets in money since the barrier function B_t will always be less than ν for an ATM option for $t > T - T_N$. It then follows that the payoff for exercise in $[T - T_N, T]$ will be less than ϵ for any $\epsilon > 0$ and the contribution to the option costs for an exercise in this interval can be neglected.

Then, we can consider the option exercise strategy as a (possibly sub-optimal) way of exercising an American Call option with strike K and expiry $T - T_N$. This gives us the desired bound. □

The bound given by Lemma 3.4, although simple turns out to be rather weak in practice. The following result allows us to compute a much stronger bound.

Refined Upper Bound on the cost C_0^ν of a vested ESO

Lemma 3.9. *Let $T' = T - T_N$. Then*

$$C_0^\nu \le C_0^{\nu+} \triangleq \max\left(S_0 - K, \mathbb{E}^Q[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+]\right) \ . \tag{3.24}$$

Proof. We know on any sample path where the option was exercised for a non-zero payoff i.e., before time $T' = T - T_N$,

$$E_\tau = C_\tau - \nu\sigma S_\tau \delta_\tau \ .$$

From the martingale property,

$$
\begin{aligned}
e^{-r\tau} E_\tau &= \mathbb{E}^Q_\tau[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})] \\
&\leq \mathbb{E}^Q_\tau[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \ .
\end{aligned}
$$

As before, let A denote the event that the option was exercised sometime during its life. Then,

$$
\begin{aligned}
C_0^\nu = \mathbb{E}^Q[e^{-r\tau} E_\tau; A] &\leq \mathbb{E}^Q[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+; A] \\
&\leq \mathbb{E}^Q[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \ .
\end{aligned}
$$

In Appendix B, the following expression for $\mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+]$ is derived in terms of bivariate normal distributions:

$$
\begin{aligned}
C_0^\nu \leq\ & C_0 - \sigma\nu S_0\delta_0 \\
& + K e^{-rT} N_2\left(d(S_0, K, T) - \sigma\sqrt{T}, -d(S_0, K, T') + \sigma\sqrt{T'}, -\sqrt{\frac{T'}{T}}\right) \\
& -(1 - \nu\sigma) S_0 N_2\left(d(S_0, K, T), -d(S_0, K, T'), -\sqrt{\frac{T'}{T}}\right) \ ,
\end{aligned} \tag{3.25}
$$

where

$$
d(S, K, T) = \frac{ln(\frac{S}{K})}{\sigma\sqrt{T}} + \left(\frac{r}{\sigma} + \frac{\sigma}{2}\right)\sqrt{T} \ .
$$

$N_2(a, b, \rho)$ denotes the probability $\Pr(X \leq a, Y \leq b)$ for two jointly normal random variables X and Y, each having zero mean and unit variance, and correlation ρ. □

Exercise Threshold ν	***Maturity Shortening*** T_N	***Lower Bound*** $C_0^{\nu-}$	***Upper Bound***	***Refined Upper Bound*** $C_0^{\nu+}$
0.00	0.000	25.21	25.21	25.21
0.25	0.096	21.42	24.82	21.44
0.50	0.378	17.63	23.64	17.72
0.75	0.839	13.84	21.65	14.05
1.00	1.473	10.05	18.73	10.43
1.25	2.280	6.26	14.65	6.85

Table 3.1: Cost Estimates (Upper bounds and Lower bounds) for an ATM ESO for different values of the exercise threshold.

Table 3.1 shows the bounds obtained by this method for a range of exercise thresholds on at-the-money vested ESO with strike 100 (for $r = 0.05$, $\sigma = 0.2$ (both annualized) and $T = 4$ years). In general, the lower bound obtained is a better estimate of the true cost, as it differs from the actual cost only when there has been an exercise after time $T - T_N$. In this case, the lower bound contribution naively adds a negative payoff to the cost instead of a zero payoff. However, both the magnitude of the negative payoff as well as the probability of exercise in the interval $(T - T_N, T)$ are small. As we can see the lower and refined upper bounds are quite close to each other. Both bounds get tighter as ν decreases and converge to the Black-Scholes value for $\nu = 0$. Also, the bounds also converge together to 0 and intrinsic value respectively for the option for that value of ν for which $T_N = T$.

Pricing unvested ESOs

In this section, we consider the case when the options have not yet vested, but will vest at a future time $T_0 > 0$. In presence of the constraints imposed due to the vesting feature, the fixed threshold delta-barrier based exercise strategy would exercise options as follows:

(a) If at the vesting time T_0, the barrier-delta function is less than the exercise threshold i.e., $B_{T_0} \leq \nu$, and the option is in the money then it is exercised immediately at T_0.

(b) If at time T_0, the barrier-delta function is above the exercise threshold ν, then the option is exercised whenever it is in money and $B_{T_0} \leq \nu$.

Vesting provision in fact increases the cost of the option to the employer. This is because the vesting constraint can only cause the employee to delay her exercise of the ESO. This effect of vesting on ESO cost has been observed before, see Ingersoll [73]. At time T_0, when the ESO vests, if the stock price then exceeds the threshold L where L satisfies $B(L, K, T - T_0) = \nu$ then the option is exercised immediately. To obtain the cost of the option with vesting, we must then replace the option cost for the sample paths when $S_{T_0} > L$ to the realized cost which is $e^{-rT_0}(S_{T_0} - K)^+$.

Lemma 3.10. *Under the delta-barrier exercise strategy with a fixed threshold ν, the cost W_0^ν of the ESOs with maturity T, vesting lag T_0 and strike K can be bounded as follows: Let $T' = T - T_N$ where T_N is as defined in Lemma 3.4. Suppose $T' \geq T_0$ and*

L be such that

$$B(L, K, T - T_0) \;=\; \nu. \qquad (3.26)$$

Then,

$$\begin{aligned} Y_0^\nu &\geq Y_0^{\nu-} = \Delta V \; + C_0 - \nu\sigma S_0 \delta_0; \\ Y_0^\nu &\leq Y_0^{\nu+} = \Delta V + \mathbb{E}^Q[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \text{ , } \textit{where,} \\ \Delta V &= \mathbb{E}^Q[e^{-rT_0}(S_{T_0} - K - C_{T_0} + \nu\sigma S_{T_0}\delta_{T_0})^+] \; . \end{aligned} \qquad (3.27)$$

Proof. It is clear that at T_0, if $S_{T_0} \geq L$ the option will be exercised, else it will behave like a regular ESO, with no vesting from that point. Since $T - T_0 \geq T_N$, using Lemma 3.4, we get $L \geq K$. Further, from Lemma 3.2,

$$S_{T_0} - K \;>\; C_{T_0} - \nu\sigma S_{T_0}\delta_{T_0} \text{ for } K > L \; .$$

Then,

$$\begin{aligned} Y_0^\nu &= \mathbb{E}^Q[e^{-rT_0}(S_{T_0} - K); S_{T_0} \geq L] + \mathbb{E}_{T_0}^Q[e^{-rT_0}C_{T_0}^\nu; S_{T_0} < L] \\ &= \mathbb{E}^Q[e^{-rT_0}(S_{T_0} - K - C_{T_0} + \nu\sigma S_{T_0}\delta_{T_0}); S_{T_0} \geq L] + \mathbb{E}^Q[e^{-rT_0}(C_{T_0} - \nu\sigma S_{T_0}\delta_{T_0}); S_{T_0} \geq L] \\ &\quad + \mathbb{E}_{T_0}^Q[e^{-rT_0}C_{T_0}^\nu; S_{T_0} < L] \\ &= \mathbb{E}_{T_0}^Q[e^{-rT_0}(S_{T_0} - K - C_{T_0} + \nu\sigma S_{T_0}\delta_{T_0})^+] + \mathbb{E}^Q[e^{-rT_0}(C_{T_0} - \nu\sigma S_{T_0}\delta_{T_0}); S_{T_0} \geq L] \\ &\quad + \mathbb{E}_{T_0}^Q[e^{-rT_0}C_{T_0}^\nu; S_{T_0} < L] \\ &= \Delta V + \mathbb{E}^Q[e^{-rT_0}(C_{T_0} - \nu\sigma S_{T_0}\delta_{T_0}); S_{T_0} \geq L] + \mathbb{E}_{T_0}^Q[e^{-rT_0}C_{T_0}^\nu; S_{T_0} < L] \; . \end{aligned}$$

Then, using Lemma 3.7, we have

$$\begin{aligned} Y_0^\nu &\geq \Delta V + \mathbb{E}^Q[e^{-rT_0}(C_{T_0} - \nu S_{T_0}\delta_{T_0}); S_{T_0} \geq L] + \mathbb{E}^Q\left[\mathbb{E}_{T_0}^Q[e^{-rT_0}(C_{T_0} - \nu\sigma S_{T_0}\delta_{T_0})]; S_{T_0} < L\right] \\ &= \Delta V + C_0 - \nu\sigma S_0\delta_0 \; . \end{aligned}$$

And using Lemma 3.9 we get,

$$\begin{aligned} Y_0^\nu &\leq \Delta V + \mathbb{E}^Q[e^{-rT_0}(C_{T_0} - \nu S_{T_0}\delta_{T_0}); S_{T_0} \geq L] \\ &\quad + \mathbb{E}^Q\left[e^{-rT_0}\mathbb{E}_{T_0}^Q[e^{-r(T'-T_0)}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+]; S_{T_0} < L\right] \\ &\leq \Delta V + \mathbb{E}^Q[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \; . \end{aligned}$$

This completes the proof. □

It may be noted that ΔV can be expressed in terms of bivariate normal distributions, once we have the value of threshold L. L however must be computed numerically by inverting the delta-barrier function, which should be relatively straight forward as this function is monotonic. Also, when $T' < T_0$, then the value of the vested option simply becomes equal to that of a call option with expiry T_0 under the delta-barrier exercise strategy.

Incorporating Termination and Reload Effects

We can extend the delta-barrier method to account for termination effects at the expense of some analytical tractability. With a reload feature, the employee upon exercising her ESOs gets new ESOs aside from the exercise proceeds. To incorporate reload options in our models, we can simply add the additional value of reload options to the exercise payoff in our myopic setting. This feature can thus be accounted with little additional computational complexity.

Similarly termination effects can be incorporated in the model exogenously. A common approach used is to model termination as an arrival in a Poisson process (with possibly time inhomogeneous arrival rates). Since the exercise strategy used by the employee is easily solvable, we can price the ESO in presence of the possibility of early termination by using direct Monte Carlo simulations.

3.7 Conclusions and Future Directions

In this chapter, we proposed a myopic risk management based framework to model exercise behavior for Employee Stock Options (ESOs). This framework leads to a tractable method to compute the cost of an ESO portfolio. This was directly motivated by the Markowitz portfolio optimization problem. The resulting exercise behavior is governed by the level of a barrier function which depends on the time value of the option, i.e., the option's premium over its intrinsic value and the instantaneous volatility in the option price, both evaluated at their Black-Scholes' values. Using this exercise behavior, we can price the cost of ESOs to the issuing company. We also showed that in general, neither would a risk-averse employee exercise all her options at the same time nor would the cost of a portfolio of options be equal to the sum of its parts. However, this becomes the case under certain model assumptions which lead to an exercise behavior where the employee exercises her option when it is in money and the barrier function is below a threshold which depends on the employee's risk-averseness. These assumptions also allow us to derive tightly spaced analytical bounds on the cost of the option to the issuers. We also indicated how these results can be extended to account for vesting and exogenous termination.

In terms of future work, it will be interesting to see if one can build a dynamic but tractable model that can take into account portfolio effects for expensing options. Also the delta-barrier function implies an artificial shortening of the option lifetime which is an undesirable modeling side-effect. A model which better accounts for the remaining value of the option's time value will be useful and so will be tighter bounds on the option's cost. Finally, an empirical analysis to see how well the predicted exercise behavior matches with observed behavior will be of great interest.

Chapter 4

Variable Annuities with Guaranteed Lifetime Payouts

4.1 Introduction

Advances in medicine and better mechanisms to contain and deal with man-made and natural calamities have led to a remarkable increase in life expectancies over the last few decades. People are living longer and spending a significant part of their lives in retirement, especially in the developed countries. At the same time, the social security systems have come under a cloud. Particularly in the US, many influential voices in both political and academic circles have expressed concerns about the sustainability of the government sponsored social security system. For example, see the opinion expressed by John Snow [111]. Moreover, life-style and health-care related costs have continuously spiralled upwards, especially for the elderly - who are a high risk group and have made life in retirement not only longer but also more expensive. As a result of all this, the idea that individuals and households need to have a systematic retirement plan for and by themselves, is now gaining more and more traction.

The private sector industry has come up with several innovations and offerings in the retirement solutions space. Insurance companies have been offering equity linked pension schemes with portfolio insurance for quite some time. Variable Annuity (VA) products are widely popular in the US across all demographics as investment and tax-planning instruments. VA Sales in the were expected to cross $180 billion for the year 2007 according to research by Milliman. Similar products are also very popular in other developed markets such as the UK and Japan and are gaining grounds in the emerging markets as well. Companies offering VA products are now embedding them with options that will allow investors to also use them as a steady and assured source of income during retirement. More than 95% VA products now offer some sort of financial guarantees with them. These guarantees generally offer the investor a form of downside protection against market risk and thus helps her secure her retirement nest-egg.

Over time, the options provided with VAs have become more innovative and exotic. One of the earliest embedded options made available with VAs was the "Guaranteed Minimum Death Benefit" scheme, see Milevsky and Posner, [93] for details. GMDB entitles the investor's beneficiary to collect a minimum benefit (usually the initial amount invested in the VA fund with a small appreciation rate) upon the investor's death. Thus the GMDB is like a stochastic maturity put option. Insurance companies subsequently introduced schemes that would enable the investor to collect the scheme's benefits during her own lifetime, unlike the GMDB. Examples of these are the Guaranteed Minimum Accumulation Benefits (GMAB) family and Guaranteed Minimum Income Benefits (GMIB) family (also known as Guaranteed Annuity Options(GAO)). GMAB guarantees to the investor a certain minimum level of capital appreciation over pre-specified horizons and is again a complex variant on the basic put option. Under the GMIB/GAO scheme, which became especially popular in the UK, investor's money is invested in a VA fund for a fixed duration and upon maturity, which typically occurs at retirement age, the underlying account value can be either withdrawn or annuitized at a guaranteed payout rate. GMIB/GAO can be considered as an equity-quantoed (equity denominated) interest rate option.

More recently, companies have been offering the "Guaranteed Minimum Withdrawal Benefit" (GMWB) schemes. Under this scheme, the investor's capital is invested through a VA fund in an asset-mix of her choice. The investor is guaranteed that upon maturity, again linked usually to the retirement age, she will be able to take at least an x % of her initial investment every year for N years, no matter how the underlying investment performs[1]. According to Milliman's third annual Guaranteed Living Benefits (GLB) survey of leading U.S. VA carriers, election rates of the GMWB scheme or its variants have increased steadily from from 24% in 2004 to 29% in 2005, to 40% in 2006, and then to 43% during the first six months of 2007 and the GMWB family has now become the most popular of all VA products.

Insurance companies have now also started offering a lifetime benefit feature with GMWB and other VA linked options, enabling the investor to simultaneously manage both *financial* as well as *longevity* related risks. While a lifetime benefit feature can be

[1] Typically $x \cdot N = 100$ so that the total guarantee is equal to the original investment.

offered with any of the GMAB, GMIB or GMWB schemes, according to Milliman's third annual GLB survey, that offered with the GMWB is emerging as the most popular of these and continues to wrest market share from the other two. GMWB with lifetime withdrawals is commonly known as "Guaranteed Lifetime Withdrawal Benefits" (GLWB) or "Guaranteed Withdrawal Benefits" (GWB) for life. GWB for life typically guarantees to the investor withdrawal amounts that are indexed to her VA account's high water-mark value (called the benefit base) for her entire life. A key advantage of the GMWB/GWB family of schemes over other VA related options available in the market is that in these schemes, the underlying investment can continue to have market exposure even when the withdrawals start and thus has a better growth opportunity. In contrast, if exercised the GMIB/GAO annuitizes the investor's account value, effectively converting it into fixed income instruments upon maturity. The GWB for life usually also has a ratchet like feature commonly known as step-ups, where the benefit base used to calculate the guaranteed withdrawal amounts is periodically and automatically raised to the VA account value, if the latter exceeds the same.

Unlike exchange traded or over the counter options, the investor pays for the options elected with VAs in a piece-meal way over several years, typically as a fixed fraction of the underlying VA account value or the defined "protection" level[2]. Currently, insurance companies are offering the GWB for life at annual fees in the range of roughly 50 to 90 basis points. The fee is typically indexed to the benefit base that is used for calculating the guaranteed withdrawal amounts[3].

If fairly priced, the GWB for life option is an attractive retirement solution for investors as it allows them to manage the risks related to their own longevities, which cannot be mitigated at an individual level. Further, the ability to stay invested in the market while in retirement would allow investors to better cope with the inflation related risk, which becomes significant as the retirement lifespans get longer[4].

For the long term sustainability of the GWB, it is also important that the companies offering it remain profitable and viable. For example, GAO schemes were

[2]This type of arrangement is in part due to the insurance industry conventions and in part to ease the burden of large payments from clients. The fee structure impacts both pricing and risk management of these products.

[3]Schemes differ in terms of promised withdrawal rates as well as features such as the frequency of "step-ups" in the benefit base and the indexing of the fees. The website http://www.annuityfyi.com provides a long but non-comprehensive list of leading insurance companies offering GWB for life like products.

[4]Equity markets are known to be better hedges against inflation as compared to fixed income instruments like annuities over a long run, see, for example, Bodie [20].

launched in the UK in a high interest rate environment, but as they drew closer to maturity, interest rates plummeted while the stock markets soared, forcing at least one company to close its product to new buyers (Chu and Kwok, [39]). Also, if the GWB is significantly underpriced or raises the possibility of a debilitating loss for the underwriting company, then the related credit-worthiness issues should make potential clients skeptical of the GWB. This is because payouts from the GWB for life option, if they happen at all, will occur only 20 to 30 years after the client has enrolled into the option. Hence the continued claims paying ability of the company underwriting the GWB guarantee is considerably important. Because VA based solutions are so widely used and can have a huge impact on market (over 55% of the estimated trillion dollar plus worth of VA assets are beleived to be in equity markets according to VA Data Research Services (VARDS)), this is an issue that should concern regulators as well.

Our goal for this chapter as well as Chapter 5 is to analyze the cost and risk of underwriting the GWB for life guarantee. Our analysis suggests that concerns about GWB will not be entirely misplaced because the product entails considerable risk from a large number of factors that should pose serious challenges in its valuation and risk management.

4.1.1 Related Work

Brennan and Schwartz [24] and Boyle and Schwartz [23] were one of the first to extend the option pricing methodology to insurance contracts where the time or frequency of payouts are linked to investor's death.

The GMDB option, which is essentially a stochastic maturity put option, has been extensively analyzed in the literature; for example by Milevsky and Polsner [93], Mudavanhu and Zhuo [97]. Chu and Kwok [38] analyze the GMAB type principle protection scheme with reset features. Moreover, Milevsky and Salisbury [95], Dai, Kwok and Zong [48], Siu [110] also consider how policy-holders can strategically exercise options embedded in a VA and their impact on the implied cost. Biffis [17], Chu and Kwok [39], Boyle and Hardy [22], Ballotta and Haberman [10], Pelsser [100] have analysed the GMIB/GAO type option feature in VAs.

"Equity Indexed Annuities" (EIAs) constitute another family of products similar to VAs with embedded options and allow investors to claim a limited or ("capped") upside in equity markets, with a floor protection. Practically, EIAs are different from VAs because unlike the VA based options which are derived from privately managed funds, EIAs are contractually linked to the published market indices. However, both families of products are analyzed using similar methods and models. EIAs have also been studied extensively: see for example, Gaillardetz and Lin [59], Buetow [26], Kijima and Wong [80], Siu [110], Cheung and Yang [36].

The GWB for life is a relatively recent addition to the market space, and is yet to be analyzed in detail in the literature. A product similar to the GWB for life - the GMWB which guarantees withdrawals over a fixed number of years, has been analyzed in Milevsky and Salisbury [95], Dai, Kwok and Zong [48], Chen, Vetzal and Forsyth [34]. Milevsky and Salisbury [95] point out that GWMB is like a Quanto Asian Put Option in a Black-Scholes economy and find that it is possibly severely underpriced. Dai, Kwok and Zong [48] analyze the same product from an investor's perspective and focus on deriving the optimal withdrawal policies, again in a Black-Scholes framework. Recently Chen, Vetzal and Forsyth [34] also consider the impact of optimal "withdrawal" strategies as well as jump risks in the context of GMWB. These papers do not consider the ratchet or step-up feature that is common in GWB for life products. Another recent work by Hoz, Kling and Rub, [69], presents a numerical analysis of the GWB for life with ratchet like features using the general contingent claim analysis framework for VAs as outlined by Bauer, Kling and Russ [13]. Again the valuation is in a Black-Scholes world. Thus, valuations of the GWB/GMWB family have so far been based primarily on the basic Black-Scholes framework with constant interest rates.

Coleman et. al [40] investigate effects of jumps and stochastic volatilities in hedging the GMDB with a ratchet or step-up like feature. Coleman, Li and Patron [41] consider hedging against both interest rate and equity related risks in the context of long-duration VA liabilities.

Insurance products that offer financial guarantees over an investor's lifetime also face risks related to population mortality (or longevity) apart from the risk due to the market factors. Lately, researchers have begun to question the commonly made assumption that population mortality risk can be considered to be statistically diversifiable. Cairns, Blake and Dowd [28], Biffis [16] and Milvesky, Promislow and Young [94] have argued that population mortality related risk may consist of a systemic component to it[5] and hence may not be entirely diversifiable. Biffis and Millossovich [17] suggest interesting ways to jointly model mortality risks as well as various financial risks in the context of GAOs.

[5]For example, medical breakthroughs or natural or man-made calamities can systematically impact a population's longevity.

4.1.2 Goals, findings and Contributions

In this chapter, we seek to develop a basic understanding of the GWB for life feature and investigate the associated risk-factors and their severity.

- We formally introduce the GWB for life feature and examine how such a product would have performed, if it were offered in the past.

- We first analyze a continuous time version of the GWB for life assuming retirement lifespans are exponentially distributed and Black-Scholes asset price dynamics. This model allows us to obtain an analytical expression for GWB value. While this analysis cannot be used as an absolute valuation tool for the GWB for life, it provides us valuable insights into the possible sources of risk. We find that:

  - The GWB guarantees become more expensive for the company as the volatility of the underlying VA fund increases. Also, the product value varies significantly depending on the investor age at inception thus creating an adverse selection bias risk.

  - The GWB value has a convex relationship with interest rates making it susceptible to volatilities in interest rates as well.

  - GWB for life also has a sizeable risk related to investor pool longevities.

4.1.3 Chapter Layout

In Section 4.2, we formally describe the product specifics and other stylized features that we consider in this chapter as well as the next. Section 4.3 then provides an idea of how such a product might have fared (from the point of view of the company offering the product) in a historical context. Next, in Section 4.4, we derive an analytical expression for the fair value of the GWB for life guarantee under the simplifying assumptions mentioned earlier. We then use this analysis with some numerical examples to get a sense of the relative price of GWB for life and the different sources of risk to value in Section 4.5. We conclude with a summary of the findings in Section 4.6.

4.2 Product Description

In this section, we formally describe the Guaranteed Withdrawal Benefit (GWB) for life feature. Typically, this option is available as an add-on or a rider feature with a VA account at a fee premium.

- All guaranteed payments and fees are defined in terms of two state variables - one is simply the underlying VA account value and the other is referred to as the "benefit base". The benefit base is used to determine the guaranteed withdrawal amount for a year. We use B_n and C_n to respectively denote the benefit base and the account value net of withdrawals, if any, at the n^{th} anniversary.

- Suppose at time 0, an investor aged A opens an account with the GWB for life feature with an initial investment C_0. This capital is invested in an asset mix of investor's choice through a VA fund. All dividends and distributions are assumed to be reinvested. For simplicity, we construct a reference index S_t for the VA fund to tracks its total returns. The initial value of the benefit base is set as $B_0 = C_0$.

- There is a minimum waiting period W and a retirement age A_R defined in the contract. The investor can start taking withdrawals from her account from the $(T+1)^{st}$ anniversary, where $T = \max(A_R - A, W)$. We assume A, A_R, W to be all integers. For $n > T$, the investor is guaranteed to be able to take a withdrawal of $q \cdot B_{n-1}$ at the n^{th} anniversary. The insurance company is responsible for covering any shortfall in case the account value falls below the guaranteed withdrawal level, i.e., $q \cdot B_{n-1}$. This is the GWB guarantee and q is termed as the guaranteed withdrawal rate.

- If the withdrawal taken at the n^{th} anniversary does not exceed the contractual guarantee (i.e., 0 for the first T anniversaries and $q \cdot B_{n-1}$ thereafter), the benefit base B_n at the n^{th} anniversary is set to the higher of B_{n-1} and the contract value after withdrawals (if any), i.e., C_n. This is the step-up (also sometimes known as ratchet) feature. If the withdrawals exceed the contractual guarantee, then B_n is set to the lower of B_{n-1} and C_n.

- The investor is allowed to withdraw during any year the greater of the gains in her account value over the previous year or a certain fraction b of her account value without any penalties. Withdrawals in excess of these and the guaranteed level, however, result in the imposition of a surrender charge. In addition, large withdrawals may also result in tax penalties.

- Every year, the investor is charged a fraction h of the benefit base for the year, i.e., $h \cdot B_{n-1}$ at the n^{th} anniversary, as fees for the GWB guarantee. We assume that the fees are charged separately to the investor rather than being deducted from the account[6].

- Upon investor' death, the residual account value is returned to a beneficiary.

While the above description does not fit any one specific product in the market, it captures the key features of this class of products. Annual fees for GWB for life option are usually the same for all participating investors, irrespective of their age or chosen asset mix. Most insurance companies offering GWB also leave room to unilaterally increase GWB fees after sales. We do not analyze the value of this option to the GWB provider. Any tendency to increase the premium in contracts already made poses a reputation risk besides going against the very spirit of insurance and will make potential clients wary.

Appendix C provides a numerical illustration of the evolution of GWB over an investor's lifetime for a hypothetical sequence of the reference VA fund returns and withdrawals by the investor.

GWB valuation is clearly dependent on how the investor chooses to withdraw. In principle, the investor can strategize her withdrawals. Milevsky and Salisbury [95], Dai, Kwok and Zong [48] and Chen, Vetzal and Forsyth [34] investigate the implications of "optimal" dynamic withdrawals by investors in the context of the GMWB and find them to increase the GMWB costs substantially. However, we believe that in practice, especially for the GWB for life type product, investors are unlikely to follow such optimal dynamic policies. This is because:

- "Optimal" withdrawal policies typically recommend the investor to withdraw her investment out completely, when the protection guarantee is out-of-the money. The papers cited above consider only the surrender charges that the insurance company levies on large withdrawals to evaluate the cost of large

[6]In practice, fees related to VA products are typically deducted from the investor's account every year and thus affect the account value. The assumption about fees being charged separately allows us to isolate the cash inflows and outflows associated with the product. This makes the break-even fee computation, which otherwise would need solving a fixed point problem, much easier. As the fee involved is small, typically few tens of basis points a year, we do not expect this approximation to alter our results in a significant way.

withdrawals for the investors. In practice, investors will also incur high indirect costs in terms of taxes on the excess distributions and this is likely to make taking large strategic withdrawals unattractive for investors.

- Investors are also more risk-averse, unable to hedge risks due to their own longevities and less equipped than institutions like insurance companies to hedge financial risk.
- Withdrawing less than the allowed amount is also likely to be sub-optimal for the investor as the guarantees are valid for her lifetime.

Also, n practice, most investors follow simple thumb rules rather than complex dynamically optimal strategies to manage their investments. In our analysis, we will primarily focus on the case where the investor withdraws the contractually guaranteed amount at each anniversary[7]. This is the maximum withdrawal the investor can take without causing the step-ups to reverse. Let R^s_{n+1} denotes the return on the underlying VA fund for the period $(n, n+1]$. The dynamics of the account value C_n and the benefit base B_n for a steady contract specified rate of withdrawal are given by:

$$\begin{aligned} C_{n+1} &= (C_n \cdot R^s_{n+1} - q_{n+1} B_n)^+ \ ; \\ B_{n+1} &= \max(B_n, C_{n+1}) \ , \end{aligned} \tag{4.1}$$

where

$$q_n = \begin{cases} 0 \ , & \text{if } n \leq T \ , \\ q \ , & \text{if } n > T \ . \end{cases} \tag{4.2}$$

So far, we have not incorporated the mortality related randomness in our model. We assume that the mortality process is independent of the market dynamics, and the insurance company is risk-neutral with respect to it. We denote by $\mathbb{Q}$, the measure on the expanded sample space containing both the securities market and the investors' mortalities and which is obtained by combining the risk neutral pricing measure and the mortality laws. While we do not price the mortality risk into the GWB, we investigate the magnitude of the implied risk. This can be used to compute a premium or risk capital. See Milevsky, Promislow and Young, [94] for suggestions about pricing mortality risk.

In the subsequent parts of this chapter as well as Chapter 5, we provide increasingly refined analyses of the GWB product. Before proceeding to a formal analysis

[7]Later, in Section 5.3 in Chapter 5, we consider an alternate dynamic withdrawal strategy.

using the risk-neutral pricing machinery, we first perform a small but interesting back-testing experiment and examine how liabilities and revenues for a GWB guarantee underwriter would have looked like, had it been offered in the past.

4.3 GWB - A Hypothetical Historical Analysis

We examine in this section how GWB would have fared had it been offered at various times in this and the last century. We set the GWB product parameters as follows- guaranteed withdrawal rate $q = 6\%$, minimum waiting period $W = 3$ years, retirement age $A_R = 65$ years, and fee rate $h = 0.65\%$.

We then calculate the total values (discounted back to the time of account opening) of the liabilities and fees for the insurance company arising from the GWB feature if an investor aged 60 opened an account with an initial investment of $100,000$ at the beginning of each of the 1124 months from January 1871 to December 1972. We consider three possible asset allocation mixes, 20% equities, 60% equities and 100% equities by value. The balance is assumed to be invested in bonds, whose month-on-month returns, we assume are the same as the prevailing long term interest rates. We also assume that the VA portfolio is re-balanced monthly to get the desired asset mix composition and dividends from equities are re-invested, i.e., there are no distributions. We consider three values for realized investor longevity - 85 years, 90 years and 95 years.

To compute the cash-flows involved, we use the monthly data for S&P composite levels, dividends and long-term interest rates, as extrapolated by Shiller and available online from the website [1] for the period - January 1871 to December 2007. Since the corresponding benchmark short-term interest rates were not available for the entire period, we use a flat annual rate of 3.5% for discounting all cash-flows.

We first consider the performance of the GWB, if the month when the VA account was opened was one of the 277 months from December 1949 to December 1972. This ensures that we consider the markets for the relatively stable post second world war interval only.

Rather surprisingly, for all asset-mix choices, in not one of these 277 scenarios, the insurance company underwriting the GWB would have had to finance any part of the guaranteed withdrawals, even if the investor went on to live for 95 years. Thus, in all cases, the minimum withdrawal guarantee was superfluous and never drawn upon! Table 4.1 summarizes the distribution of the net-value earned to the company from the product, i.e., the difference between the total discounted value of the fees collected and the payouts made over the product's lifespan. This suggests a healthy 20% to 30% margin on sales on average. Also, it suggests that the longer the investor lives and the more aggressive her VA fund choice, the higher are the revenues for the company on average[8].

Equity Expo-sure	***Investor Age at Death***	***Average***	***Std.***	***minimum***	***5%le value***	***Median***	***95%le value***	***maximum***
20%	85	17,145	2,541	14,259	14,471	15,806	22,467	22,677
	90	20,554	3,524	16,568	16,787	19,317	27,552	27,751
	95	23,970	4,308	18,708	19,006	22,741	32,004	32,286
60%	85	19,978	2,699	16,401	16,747	19,110	25,791	26,669
	90	24,510	3,638	19,480	20,210	23,804	33,308	34,361
	95	29,709	4,753	23,118	24,095	28,881	42,141	43,572
100%	85	23,509	6,889	16,042	16,655	21,214	43,566	45,767
	90	28,965	7,986	19,930	20,776	25,080	51,064	54,000
	95	35,565	9,118	23,719	25,922	32,119	57,311	60,933

Table 4.1: Distribution statistics of GWB for life net value in different scenarios for an account started between Dec. 1949 and Dec. 1972 (for initial investment of 100,000).

If however, we consider the entire range of the data available[9], i.e., a total of 1224 vintages with the month of account opening ranging from January 1871 to December 1972, the results look very different. Table 4.2 shows key statistics for the total discounted value of the payouts that the company would have had to finance while Table 4.3 shows the same for the total discounted net value, i.e., fees less payouts, that the company booked.

Equity Expo-sure	***Investor Age at Death***	***Average***	***Std.***	***minimum***	***5%le value***	***Median***	***95%le value***	***maximum***
20%	85	0	0	0	0	0	0	0
	90	124	939	0	0	0	5,045	13,876
	95	1,630	4,025	0	0	0	19,549	29,182
60%	85	579	3,613	0	0	0	20,340	44,813
	90	2,669	9,642	0	0	0	51,255	75,390
	95	5,898	16,436	0	0	0	79,621	101,371
100%	85	9,734	27,822	0	0	0	130,526	160,970
	90	17,268	42,238	0	0	0	182,010	231,579
	95	25,672	55,316	0	0	0	225,359	291,843

Table 4.2: Distribution statistics of GWB for life payouts in different scenarios for an account started between Jan. 1871 and Dec. 1972 (for initial investment of 100,000).

[8]Note that these are average revenues and do not represent fair values.

[9]The corresponding time-frame would include two world wars and the Great Depression.

Equity Exposure	*Investor Age at Death*	*Average*	*Std.*	*minimum*	*5%le value*	*Median*	*95%le value*	*maximum*
20%	85	15,004	1,880	12,070	13,173	14,569	22,199	22,677
	90	17,011	2,781	5,269	14,645	16,305	27,232	27,751
	95	17,424	5,319	(8,252)	8,653	17,320	31,636	32,286
60%	85	18,378	4,512	(17,936)	13,392	17,855	27,698	29,486
	90	19,511	9,209	(44,945)	1,293	20,367	33,046	35,122
	95	19,350	15,793	(67,686)	(19,154)	22,156	39,785	43,572
100%	85	14,989	25,965	(115,372)	(50,905)	19,273	51,043	58,948
	90	12,475	40,230	(170,711)	(89,250)	22,063	64,825	74,888
	95	8,929	53,827	(223,944)	(120,486)	25,229	76,652	88,656

Table 4.3: Distribution statistics of GWB for life net value in different scenarios for an account started between Jan. 1871 and Dec. 1972 (for initial investment of 100,000).

These tables show that while still profitable on average, the GWB for life is definitely not a low-risk money spinner. The net-income values are considerably smaller and the insurance company would have had to burden guaranteed withdrawals in quite a few scenarios resulting in substantial losses (assuming no hedging). Also, average net value now decreases with the aggressiveness of the chosen asset mix and higher investor longevities no longer appear beneficial to the company.

The statistics for shortfall and net-value given in Table 4.2 imply that they have a skewed distribution in which large losses occur with small probabilities. This suggests that model estimation risk is important for GWB as it is the extreme events that drive most of the losses and the pricing model used must be able to accurately capture them.

We end this section with a caveat - in general, any historical analysis for GWB can be regarded as best as only instructive, because the GWB not only has a skewed distribution for realized value but also has a long duration and financial data is typically not stationary over such durations.

We now turn to valuing GWB using risk neutral pricing methods.

4.4 Black Scholes Model with Continuous Step-ups and Exponential Mortality - CBSME Model

As (4.1) indicates, GWB has discrete cash-flows and step-ups. Considerable analytical simplicity however is offered by considering a continuous time version of the GWB. In this section, we consider this case and derive an analytical expression for the GWB value. This is useful in generating insights in to the risk factors associated with the

GWB.

- Step-ups as well as all cash-flows, i.e., withdrawals and fee payments are made continuously.

- Investor's lifespan in retirement is exponentially distributed with mean $\frac{1}{\lambda}$.

We assume that S_t, the reference index for the chosen VA fund, follows a geometric Brownian motion with constant volatility σ and the risk free rate r rate is also constant. This is basically the Black-Scholes model [18].

The investor starts taking withdrawals from time $t > T \triangleq \max(W, A_R - A)$ at the rate $q_t B_t$ where,

$$q_t = \begin{cases} 0 \,, & \text{if} \quad t \leq T \,, \\ q \,, & \text{if} \quad t > T \,. \end{cases}$$

The investor also (separately) pays fees to the company at rate h indexed to B_t. For ease of reference, we shall henceforth refer to this model, which treats GWB as a continuous time instrument and prices it under the Black Scholes framework assuming exponential retirement lifespans, as the CBSME model.

In this model, the only state variables that affect the GWB value are B_t, C_t and the investor age or equivalently time t. The dynamics for B_t and C_t are given as:

$$dC_t = \begin{cases} rC_t dt - q_t B_t dt + C_t \sigma dZ_t^Q & \text{if} \quad C_t \geq 0 \,, \\ 0 & \text{if} \quad C_t = 0 \,. \end{cases} \tag{4.3}$$

$$dB_t = \mathbf{1}_{\{C_t = B_t\}} (dC_t)^+ \,. \tag{4.4}$$

Here Z_t^Q is a $\mathbb{Q}$ Brownian motion. Now, let $L(C, B, t)$ denote the fair value of the protection offered to the investor through the GWB scheme at t as a function of the state variables. This is simply the part of the withdrawals that are borne by the insurance company. Note that if $H(C, B, t)$ denotes the fair value of all withdrawals made by the investor or her beneficiary, then the following equality holds

$$L(C, B, t) = H(C, B, t) - C \,. \tag{4.5}$$

Let $G_h(C, B, t)$ denote the value of the revenue stream from the investor at time t as a function of the state variables. We also define a normalized revenue stream as $G(C, B, t) \triangleq \frac{G_h(C,B,t)}{h}$.

Our goal is to find the functions $L(C, B, t)$ and $G(C, B, t)$.

We first prove the following property that reduces the dimensionality of the problem.

Proposition 4.1. *The functions $L(C, B, t)$ and $G(C, B, t)$ are homogeneous B, C and satisfy the following relations:*

$$\begin{aligned} L(C,B,t) &= B \cdot L\left(\frac{C}{B}, 1, t\right) \ ; \\ G(C,B,t) &= B \cdot G\left(\frac{C}{B}, 1, t\right) \ . \end{aligned} \tag{4.6}$$

Proof. Let the tuple (B_t, C_t) denote an investor's benefit base and account value at time t. For a given sample path ω of the evolution of the market and mortality factors, let $(B_s(\omega), C_s(\omega))$ be the value of this tuple at some time $s \geq t$ along ω. From the homogeneity of the system dynamics as given by (4.4), it follows that had the investor started with a benefit base and account value combination $(y \cdot B_t, y \cdot C_t)$ for some $y > 0$, then the sample path ω values of these quantities at some time $s \geq t$ would have been $(y \cdot B_s(\omega), y \cdot C_s(\omega))$. Moreover, the cashflows that the company incurs, i.e., the shortfalls and revenues at any time t are also homogeneous functions of the benefit base B_t and account value C_t. From these, it follows that for any $y < 0$,

$$\begin{aligned} L(y \cdot C, y \cdot B, t) &= y \cdot L(C, B, t) \ ; \\ G(y \cdot C, y \cdot B, t) &= y \cdot G(C, B, t) \ . \end{aligned}$$

Hence for $B > 0$, we must have

$$\begin{aligned} L(C,B,t) &= B \cdot L\left(\frac{C}{B}, 1, t\right) \ ; \\ G(C,B,t) &= B \cdot G\left(\frac{C}{B}, 1, t\right) \ . \end{aligned}$$

□

In light of Proposition 4.1, we define the functions $l(x, t)$ and $g(x, t)$ as follows

$$\begin{aligned} l(x,t) &= L(x,1,t) \ \ldots \ , 0 \leq x \leq 1 \ ; \\ g(x,t) &= G(x,1,t) \ \ldots \ , 0 \leq x \leq 1 \ . \end{aligned} \tag{4.7}$$

The quantity x can be interpreted as the benefit base capitalization ratio $\frac{C}{B}$. The functions $l(x,t)$ and $g(x,t)$ are to be interpreted as the values of respectively the liabilities and revenue streams, normalized by the benefit base, before a step-up operation is performed.

It turns out that for $t \geq T$, it is easier to find $L(C,B,t)$ and $G(C,B,t)$ (or equivalently $l(x,t)$ and $g(x,t)$) by solving a differential equation. Once this is done, $L(C,B,t)$ and $G(C,B,t)$ can be determined for $t < T$ by taking a risk neutral expectation of their respective values at time T. For this we will need the joint distribution for (B_T, C_T) conditional on (B_t, C_t) for $t \leq T$ under the risk neutral measure $\mathbb{Q}$. We therefore break the pricing of GWB into two different phases:

1. **Phase 1**: When $t < T$. There are no withdrawals in this phase.

2. **Phase 2**: When $t \geq T$ and the investor has started taking withdrawals.

We now describe each of these in detail beginning with Phase 2.

Value of Cashflows in Withdrawal Phase (Phase 2)

Here, because of the assumption that the investor has an exponentially distributed residual life in retirement, $l(x,t)$ and $g(x,t)$ as defined in (4.7) become independent of time t for $t \geq T$. We summarize and prove this in the following proposition.

Proposition 4.2. *For $t \geq T$,*

$$\begin{aligned} l(x,t) &= l(x,T) \; ; \\ g(x,t) &= g(x,T) \; . \end{aligned} \tag{4.8}$$

Proof. Note that both the market returns process and the mortality process are memoryless for $t \geq T$. Moreover the dynamics of the state variables as given (4.4) also do not depend on t for $t \geq T$ as $q_t = q$ is a constant for this period. It then follows that for $t \geq T$, the system has no memory and hence

$$\begin{aligned} L(C,B,t) &= L(C,B,T) \; ; \\ G(C,B,t) &= G(C,B,T) \; . \end{aligned}$$

Using (4.7), the result in (4.8) then follows immediately. □

For $t \geq T$, then $l(x,t)$ and $g(x,t)$ can be regarded as functions of just one variable. To ease notation, for this particular subsection, where we are working under the case $t \geq T$, we shall denote them as simply $l(x)$ and $g(x)$ respectively.

Lemma 4.1. *1. The function $l(\cdot)$ satisfies the following 2^{nd} order differential equation*

$$\frac{1}{2}\sigma^2 x^2 l'' + (rx-q)l' - (r+\lambda)l \;=\; 0 \; ; \tag{4.9}$$

and the boundary conditions

$$l(0) \;=\; \frac{q}{r+\lambda} \; , \tag{4.10}$$
$$l'(1) \;=\; l(1) \; . \tag{4.11}$$

2. The function $g(\cdot)$ satisfies the following 2^{nd} order differential equation

$$\frac{1}{2}\sigma^2 x^2 g'' + (rx-q)g' - (r+\lambda)g + 1 \;=\; 0 \; ; \tag{4.12}$$

and the boundary conditions

$$g(0) \;=\; \frac{1}{r+\lambda} \; , \tag{4.13}$$
$$g'(1) \;=\; g(1) \; . \tag{4.14}$$

Proof. We will prove the result only for $l(x)$, as the proof for $g(x)$ follows along identical lines. We know that in phase 2, $L(C,B,t) = L(C,B,T) = B \cdot l(\frac{C}{B})$. Let A be the event that the investor passes away in the interval $(t, t+dt]$. Thus $\mathbb{P}(A) = \lambda dt + o(dt^2)$. When $0 < C < B$, $dB = 0$. Then by Ito's Lemma,

$$\begin{aligned} dL &= \mathbb{P}(A)\cdot(0-L) + (1-\mathbb{P}(A))\cdot\left(\frac{\partial L}{\partial C}dC + \frac{1}{2}\frac{\partial^2 L}{\partial C^2} < dC \cdot dC >\right) \\ &= -\lambda L dt + \frac{\partial L}{\partial C}\left((rC - qB)dt + \sigma C dZ_t^Q\right) + \frac{\partial^2 L}{\partial C^2}\sigma^2 C^2 dt \; . \end{aligned}$$

Here, $< dC \cdot dC >$ is short-hand for $\mathbb{E}^{\mathbb{Q}}[dC^2]$. Under the risk-neutral measure, we must have

$$\mathbb{E}^{\mathbb{Q}}[dL] \;=\; rLdt - (qBdt - L)^+ \; ,$$
$$\text{i.e., } (rC - qB)\frac{\partial L}{\partial C} + \frac{1}{2}\sigma^2 C^2 \frac{\partial^2 L}{\partial C^2} - \lambda L \;=\; rL \; . \tag{4.15}$$

Using the property that $L(C,B,t) = B \cdot l(x)$, with $x \triangleq \frac{C}{B}$, we get

$$\begin{aligned}(rx-q)l' + \frac{1}{2}\sigma^2 x^2 l'' - \lambda l &= rl \;, \\ \text{i.e., } \frac{1}{2}\sigma^2 x^2 l'' + (rx-q)l' - (r+\lambda)l &= 0 \;.\end{aligned}$$

For the first boundary condition, we simply note that

$$\begin{aligned}l(0) &= L(0,1,T) \\ &= \mathbb{E}^{\mathbb{Q}}\left[\int_0^\infty exp(-(r+\lambda)t)\cdot q dt\right] \\ &= \frac{q}{r+\lambda} \;.\end{aligned}$$

The second boundary condition is more technical and represents a "smooth pasting condition". We give a rough proof:

Suppose step-ups can occur only after a delay Δ. Let $B_t = B$, $C_{t+\Delta} = C$ and $L(C_{t+\Delta}, B_{t+\Delta}) = L(C, \max(C,B)) \triangleq \tilde{L}_B(C)$. The function $\tilde{L}_B(\cdot)$ is assumed to be continuously differentiable. Then, it follows that

$$\text{if } C > B,\; \frac{d\tilde{L}_B(C)}{dC} = \frac{dL(C,C)}{dC} = l(1) \;.$$

By continuity,

$$\frac{d\tilde{L}_B(C)}{dC}\Big|_{C=B} = l(1) \;. \tag{4.16}$$

Now, since B_t is continuous,

$$\begin{aligned}\lim_{\Delta\to 0} \tilde{L}_B(C) &= Bl\left(\frac{C}{B}\right) \;, \\ \text{hence, } \lim_{\Delta\to 0} \frac{d\tilde{L}_B(C)}{dC} &= l'\left(\frac{C}{B}\right) \;. \end{aligned} \tag{4.17}$$

Putting (4.16) and (4.17) together, we get the boundary condition (4.11). □

The differential equations in Lemma 4.9 have the following general solutions

$$l(x) = C_1 x^{-k} e^{-\frac{q}{\sigma^2 x}} \mathrm{WhM}\left(k, m, \frac{2q}{\sigma^2 x}\right) + C_2 x^{-k} e^{-\frac{q}{\sigma^2 x}} \mathrm{WhW}\left(k, m, \frac{2q}{\sigma^2 x}\right) , \tag{4.18}$$

$$g(x) = D_1 x^{-k} e^{-\frac{q}{\sigma^2 x}} \mathrm{WhM}\left(k, m, \frac{2q}{\sigma^2 x}\right) + D_2 x^{-k} e^{-\frac{q}{\sigma^2 x}} \mathrm{WhW}\left(k, m, \frac{2q}{\sigma^2 x}\right) + \frac{1}{r+\lambda} ; \tag{4.19}$$

where

$$k = \frac{r}{\sigma^2} - 1 ,$$

$$m = \sqrt{\left(\frac{1}{2} + \frac{r}{\sigma^2}\right)^2 + 2\frac{\lambda}{\sigma^2}} .$$

The functions $\mathrm{WhW}(k, m, z)$ and $\mathrm{WhM}(k, m, z)$ are hypergeometric functions that solve the Whittaker Differential Equation (see Mathworld, [91]):

$$\frac{d^2u}{dz^2} + \frac{du}{dz} + \left(\frac{k}{z} + \frac{\frac{1}{4} - m^2}{z^2}\right) u = 0 .$$

Some basic properties of these functions, as well as the limits $\lim_{z\to\infty} \mathrm{WhM}(k, m, z)$ and $\lim_{z\to\infty} \mathrm{WhW}(k, m, z)$ are provided in Appendix D. Using these properties, we can find the constants C_1, C_2, D_1, D_2 satisfying the boundary conditions given in Lemma 4.1 as

$$C_1 = \frac{q}{r+\lambda} \cdot \left(\frac{2q}{\sigma^2}\right)^k \frac{\Gamma(\frac{1}{2} - k + m)}{\Gamma(1+2m)} ; \tag{4.20}$$

$$C_2 = C_1 ; \frac{(\frac{1}{2} + k + m)\mathrm{WhM}(k+1, m, \frac{2q}{\sigma^2}) + \mathrm{WhM}(k, m, \frac{2q}{\sigma^2})}{\mathrm{WhW}(k+1, m, \frac{2q}{\sigma^2}) - \mathrm{WhW}(k, m, \frac{2q}{\sigma^2})} ; \tag{4.21}$$

$$D_1 = 0 ; \tag{4.22}$$

$$D_2 = \frac{\frac{1}{r+\lambda}}{\mathrm{WhW}(k+1, m, \frac{2q}{\sigma^2}) - \mathrm{WhW}(k, m, \frac{2q}{\sigma^2})} . \tag{4.23}$$

Value of Cashflows at Inception and in Phase 1

The values of liabilities and the normalized revenue streams for the GWB product at time T, when the withdrawals start will be given by

$$L(C_T, B_T, T) = B_T \cdot l\left(\frac{C_T}{B_T}\right) ;$$

$$G(C_T, B_T, T) = B_T \cdot g\left(\frac{C_T}{B_T}\right) ,$$

where the functions $l(\cdot)$, $g(\cdot)$ and the various parameters in the formulae are given by (4.18), (4.19) and (4.20)-(4.23). Since there are no cash outflows involved in Phase 1 or during the waiting period, when the investor does not withdraw, it follows that, if I_t denotes the indicator variable that the investor is alive at time t, then

$$\begin{aligned} L(C_0, B_0, 0) &= \mathbb{E}^{\mathbb{Q}}[I_T e^{-rT} L(C_T, B_T, T)] \\ &= e^{-rT}\mathbb{E}[I_T]\mathbb{E}^{\mathbb{Q}}\left[B_T \cdot l\left(\frac{C_T}{B_T}\right)\right] . \end{aligned} \tag{4.24}$$

For the revenue stream, we have

$$\begin{aligned} G(C_0, B_0, 0) &= \mathbb{E}^{\mathbb{Q}}\left[\int_0^T e^{-rt} I_t B_t dt + I_T e^{-rT} G(C_T, B_T, T)\right] \\ &= \int_0^T \left(e^{-rt} \cdot \mathbb{E}[I_t] \cdot \mathbb{E}^{\mathbb{Q}}[B_t]\right) dt + e^{-rT} \cdot \mathbb{E}[I_T] \cdot \mathbb{E}^{\mathbb{Q}}\left[B_T \cdot g\left(\frac{C_T}{B_T}\right)\right] . \end{aligned} \tag{4.25}$$

From (4.24) and (4.25), it follows that, to compute the value of the product at time 0, we need:

- The joint distribution of B_T and C_T under the risk neutral measure.
- The marginal distribution of B_t under the risk neutral measure.

Fortunately, both these quantities are computable, using properties of Brownian Motion. It is easier to work with the transformed processes

$$c_t \triangleq \ln\left(\frac{C_t}{C_0}\right) ; \tag{4.26}$$

$$\begin{aligned} b_t &\triangleq \sup_{u:0\le u\le t} c_u \\ &= \ln\frac{(\sup_{u:0\le u\le t} C_u)}{C_0} = \ln\left(\frac{B_t}{B_0}\right) . \end{aligned} \tag{4.27}$$

The equality in (4.27) follows from the fact that $\ln(\cdot)$ is a monotonous function and that $C_0 = B_0$.

It can be shown that the joint distribution of c_t, b_t under $\mathbb{Q}$ is given by

$$\begin{aligned} f^{Q}_{c_t,b_t}(z,m) &= \begin{cases} \frac{2(2m-z)}{\sigma^2 t} \cdot \frac{1}{\sigma\sqrt{t}}\Phi\left(\frac{2m-z}{\sigma\sqrt{t}}\right) \cdot \exp\left(\frac{\nu}{\sigma^2}z - \frac{\nu^2}{\sigma^2}t\right) & \ldots, \quad m \ge z , \\ 0 & \ldots, \quad m < z ; \end{cases} \\ \text{where, } \nu &= r - \frac{1}{2}\sigma^2 . \end{aligned} \tag{4.28}$$

The function $\Phi(x) \triangleq \frac{1}{\sqrt{2\pi}}\exp(-\frac{x^2}{2})$ is the standard normal density function. The marginal distribution of b_t on the other hand is given by

$$f_{b_t}^Q(m) \quad = \quad \frac{2}{\sigma\sqrt{t}}\Phi\left(\frac{m-\nu t}{\sigma\sqrt{t}}\right) - \frac{2\nu}{\sigma^2}\exp\left(\frac{2\nu m}{\sigma^2}\right)\cdot \mathrm{N}\left(\frac{-m-\nu t}{\sigma\sqrt{t}}\right) \quad \ldots, m \geq 0. \tag{4.29}$$

Here $\mathrm{N}(\cdot)$ denotes the standard normal cumulative density function. Equations (4.28) and (4.29) are derived in Appendix E. These distributions then allow us to compute the expectations in (4.24) and (4.25) and consequently $L(C, B, 0)$ and $G(C, B, 0)$.

Note that (4.28) and (4.29) are respectively the joint distribution of (b_T, c_T) and marginal distribution of b_t conditional on $b_0 = c_0$.

For evaluating the value of the product at an intermediate time s during Phase 1, such that, $0 < s < T$, we would need these distributions conditional on general values of b_s, c_s satisfying $c_s \leq b_s$. This can be in fact obtained readily from (4.28) and (4.29) by noting that

$$\begin{aligned} b_t &= \sup_{u:0\leq u\leq t} c_u \\ &= \max\{b_s, \sup_{u:s\leq u\leq t} c_u\} \ . \end{aligned}$$

Unfortunately, although we know the joint distributions for b_t and c_t (and hence effectively B_t and C_t), through (4.28), the integrals in (4.24) and (4.25) do not have a closed form representation and must be evaluated numerically.

Note that, we have assumed exponential mortality model for the investor only in Phase 2. Investor mortality rates during Phase 1, to be used in (4.24) and (4.25) can be arbitrary.

We now use the expressions derived in this section, to compute the value and sensitivity of the GWB for life product for a typical offering.

4.5 Numerical Results

In this section, we use the results derived in Section 4.4 for some numerical computations. As the CBSME model considered therein is an approximation of the actual GWB product, we focus more on risk and sensitivity analysis rather than absolute valuations. The experimental set-up is as follows:

Product Parameters

- We set the minimum waiting period $W = 3$ years, the retirement age $A_R = 65$ years and the guaranteed withdrawal rate $q = 6\%$. These are indicative of typical offerings.

- For computing the net value of the GWB, we will assume that the fees are charged at the rate $h= 0.65\%$ of the benefit base and take a reference initial investment of $C_0 = 100$. Thus the net value figures that we obtain can be interpreted to have units of percentage of sales of VAs for which the GWB for life was elected.

Investors' Profiles

- We consider investor age at inception, i.e., A to vary in the range 50 to 70 years.

- Investor chooses an asset-mix for investment by selecting a target level of exposure (by value) that her portfolio will have to equities and we denote the same by α. We consider four levels for α - $20\%, 40\%, 60\%$ and 80%. The balance of the portfolio will be invested in relatively less volatile instruments, which we refer to as "bonds"[10]. The VA fund is also assumed to be continuously rebalanced to maintain the taget exposures.

- We also consider the overall value of GWB for the insurance company for sales across different investor cohorts and investment plans. For this, we assume that the distribution of clients' age at inception, A, weighted by their initial investment amount C_0 is uniformly distributed in the range 50 to 70. Further, for each cohort, the fraction of investment in VA funds with values of α as $20\%, 40\%, 60\%$ and 80% is $0.1, 0.4, 0.4$ and 0.1 respectively. We refer to this portfolio of clients with the stated distribution of age at inception and asset-mix selection as the "meta-portfolio" and this is again indicative of a typical VA client pool.

- For estimating mortality risk, we use the year 2008 mortality table[11] published by the Pension Benefit Guaranty Corporation (PGBC) and which is used to value annuities under the Employee Retirement Income Security Act (ERISA)

[10] These need not be interest rate or zero-coupon bonds.

[11] These mortality rates have been obtained as a 50-50 blend of mortality rates for healthy males and females in the US.

Section 4050 available from [101]. We assume that the mortality rates remain stationary, i.e., different cohorts experience the same mortality rate at the same ages. We convert this agewise table into a continuous time mortality or hazard rate function by modeling investor death as the arrival of a time inhomogeneous Poisson process with piecewise constant intensities, each piece being of length one year[12]. For convenience, the relevant mortality rates published in this table and the implied hazard rates and average residual life (for ages 49 and above) are listed in Appendix F.

For Phase 2, i.e., for $t > T$, we then set the intensity to be the inverse of the expected residual life of the investor at T, as implied from the hazard rate function. Thus, in Phase 2, we model death as the arrival of a time homogeneous Poisson process.

For illustration, Table 4.4 gives the maximum number of Phase 1 and average number of Phase 2 years for five different cohort ages - $50, 55, 60, 65$ and 70. For the cohort who starts at ages 50, 55 and 60, withdrawals will start at age 65. Because of the minimum waiting period constraint, withdrawals for the cohorts aged 65 and 70 at inception will start at ages 68 and 73 respectively.

Cohort Age	***Max. Phase 1 years (T)***	***Avg. Withdrawal Years ($\frac{1}{\lambda}$)***
50	15	20.38
55	10	20.38
60	5	20.38
65	3	18.01
70	3	14.27

Table 4.4: Maximum Phase 1 years and Average Phase 2 years for select cohorts.

Asset Dynamics

We assume that both equity markets as well as bonds have log-normal returns and are uncorrelated. The Black Scholes volatility of the equity market returns σ_e, is set to 20% (annualized), while the same for the bonds, σ_b, is set to 2% (annualized). The one year risk-free rate is taken to be 3.5%. Because the VA fund is rebalanced continually, the VA fund index S_t will follow a geometric Brownian motion consistent with the assumptions of Section 4.4. Its volatility σ for an exposure α to equities is

[12] This is commonly known as De-Moivre's approximation.

obtained as

$$\sigma^2 \;=\; \alpha^2 \cdot \sigma_e^2 + (1-\alpha)^2 \cdot \sigma_b^2 \quad . \tag{4.30}$$

The following table gives the effective volatility for the four asset allocation mixes that we consider:

Equity Exposure (α)	***Effective Volatility (σ)***
20%	4.31%
40%	8.09%
60%	12.03%
80%	16.01%

Results

For each cohort or age-group, we compute the value of Liabilities and Revenue base at inception for an account started with an investment of 100, i.e., the values $L(100, 100, 0)$ and $G(100, 100, 0)$. These are reported in Tables 4.5 and 4.5 respectively. These are used to compute the break-even fee h_0, or the fee at which the Net Present Value will become 0 and the NPV, assuming a fee of $h = 0.65\%$. Results for each cohort and asset mix are summarized in Tables 4.5 and 4.5.

	Exposure to equities			
Cohort Age	**20%**	**40%**	**60%**	**80%**
50	10.38	14.97	21.80	30.57
55	10.55	14.98	21.38	29.33
60	10.77	14.89	20.53	27.29
65	8.64	11.96	16.44	21.77
70	5.28	7.66	10.97	14.98

Table 4.5: Liability Values for different cohorts and asset mixes under the CBSME Model.

From these results, it would appear that the typical fees of around 65 bps charged by the company grossly underprices the product. However, we re-emphasize that the CBSME model used here analyzes a continuous time version of the actual product and is not comparable to the same at an absolute level. Besides, the exponential model for retirement lifespans is also rather crude as is evidenced by Figure 4-1 that shows how average residual life varies with age.

	Exposure to equities			
Cohort Age	**20%**	**40%**	**60%**	**80%**
50	2684	2870	3126	3427
55	2192	2345	2545	2776
60	1712	1829	1975	2138
65	1423	1515	1628	1754
70	1247	1325	1420	1526

Table 4.6: Revenue Values/Fees for different cohorts and asset mixes under the CB-SME Model.

	Exposure to equities			
Cohort Age	**20%**	**40%**	**60%**	**80%**
50	0.39%	0.52%	0.70%	0.89%
55	0.48%	0.64%	0.84%	1.06%
60	0.63%	0.81%	1.04%	1.28%
65	0.61%	0.79%	1.01%	1.24%
70	0.42%	0.58%	0.77%	0.98%

Table 4.7: Break-even fees for select cohorts and asset mixes under the CBSME Model.

	Exposure to equities			
Cohort Age	**20%**	**40%**	**60%**	**80%**
50	7.07	3.69	-1.48	-8.30
55	3.70	0.26	-4.83	-11.29
60	0.36	-3.00	-7.69	-13.39
65	0.61	-2.11	-5.85	-10.37
70	2.82	0.95	-1.74	-5.06

Table 4.8: Net value of the GWB product for select cohorts and asset mixes under the CBSME Model.

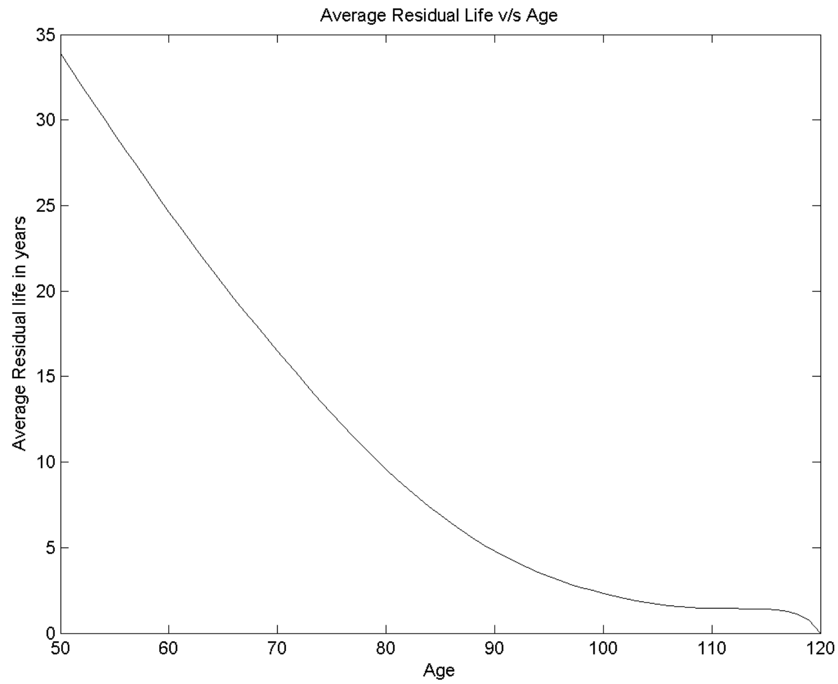

Figure 4-1: Average Residual life as a function of Age

Nevertheless, since the version of GWB considered by the CBSME model is structurally similar to the actual product, we can use the results of Section 4.4 to get an idea of primary risk factors and their magnitudes. The availability of a near closed form solution considerably speeds up the computations needed for this investigation.

Figure 4-2, shows how the break even fee varies with cohort age while Figure 4-3 shows the variation of net-value (at $h = 0.65\%$ fees) with cohort age for different asset mix choices.

We observe that:

- In general, the more aggressive the asset mix, the more expensive would be the GWB for life guarantee. A decomposition of the net-value as provided in Tables 4.5 and 4.5 for selected cohorts reveals that in fact the payout liabilities increase sharply with the VA fund volatility. This would mean that the risk capital requirements will also be higher for GWB guarantee associated with the more volatile VA funds. The corresponding changes in the net value and break-even fees are more subdued because the fee structure which is indexed to the benefit base provides a greater upside in revenues for more aggressive asset mixes and helps to somewhat offset the increases in the liabilities.

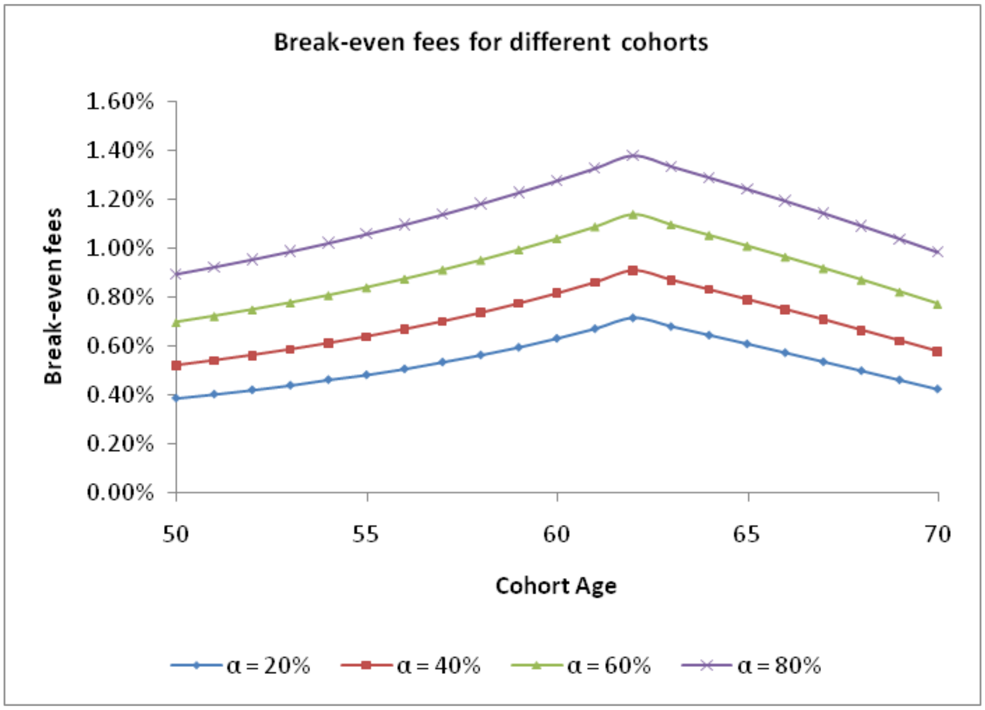

Figure 4-2: Break Even fee as a function of Age A for different asset mix choices under the CBSME Model.

- Cohort age is a significant determinant of the value of the product. The ideal time for an investor to opt for GWB is probably just before retirement so as to minimize the duration of Phase 1 and maximize the same for which the withdrawal guarantee applies, i.e., Phase 2. Because of the waiting period $W = 3$ years, the number of years for which withdrawals are guaranteed starts decreasing with the investor age at inception, A, for $A \geq 62$. However, for an investor with $A < 62$ years, the number of phase 2 years is the same as what it would have been had she started at age 62 years, because she must reach the retirement age $A_R = 65$ years before withdrawals can start. On the other hand, the number of years for which the insurance company can collect fees from the investor is more. This leads to a fall in the break-even fees and an increase in the net value for the company[13]. As a result cohorts aged 62 are the most expensive for the insurance company.

[13] Because of the step-up feature, the costs and revenues associated with the GWB actually have a slightly more complex relationship with cohort age. However, as step-ups tend to increase the value of both the withdrawal guarantees as well as the fee base, the effects somewhat offset each other. This also suggests that the step-up feature is of limited additional value to the investor because the fees are indexed to the benefit base and hence they also get stepped up with the level of the guarantee.

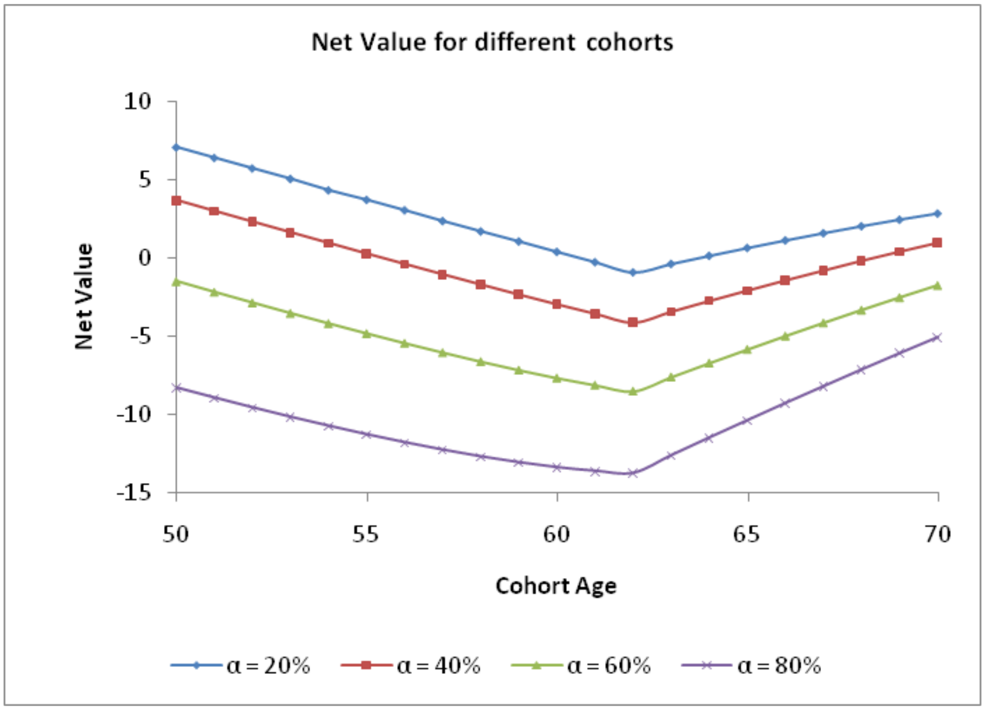

Figure 4-3: Net Value as a function of Age for different asset mix choices under the CBSME Model.

We find that there is a significant risk of adverse selection of clientele in the one price fits all approach that the companies have been using to market and sell the GWB. In the long run, companies must differentiate in fees based on the investor age and their fund selections. A few companies have started charging a premium for GWB guarantees on more aggressive asset mixes. However price differentials across investor ages are also needed.

Next, we consider the sensitivity of the entire client portfolio, i.e., the meta-portfolio described earlier to shifts in interest rates and mortality rates. Figures 4-4 and 4-5 show how respectively the break-even fees and the net-value fluctuate with the interest rate r. As expected, the GWB guarantee becomes more expensive as interest rates decrease. Further, the relationship between net value and interest rates is convex[14]. This indicates that interest rate volatility will work to lower GWB valuations and must be considered in pricing it.

Finally, we compute the net value and break even fees for the meta-portfolio if all average residual lives (during Phase 2) were to change by ± 1. The numbers,

[14]An intuitive reason for this convex nature is that as the risk-free rate decreases, not only does the possibility of a withdrawal shortfall increase but also the discount factor that would apply to the resulting payouts decreases.

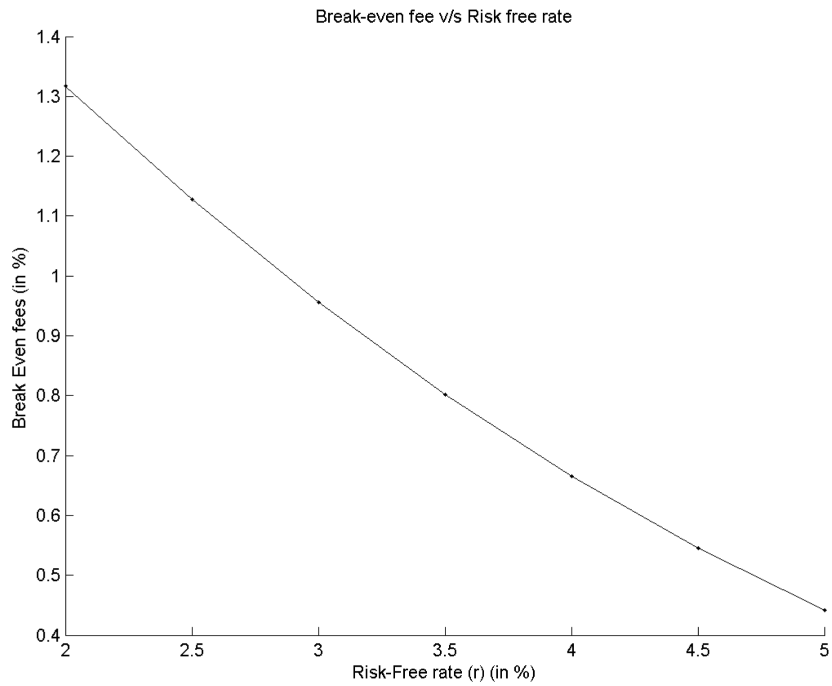

Figure 4-4: Break even fee for the meta-portfolio as a function of risk-free rate under the CBSME Model.

presented in Table 4.9, show that the mortality related risk is substantial. To put things in perspective, the average residual lives at 65 for healthy males and females in the US differ by more than 2 years.

Change in Avg. Residual Life	***Break even fee***	***Net value at h=0.65%***
none	0.80 %	-3.09
+1	0.85 %	-4.08
-1	0.76 %	-2.10

Table 4.9: Impact of changes in average withdrawal years on break-even fee and net value of the meta-portfolio under the CBSME Model.

Thus, the scenarios considered in Table 4.9 are well within the realms of possibility. Also, the effective longevities for the insurance company can be different from the population longevities and difficult to estimate if the amount of capital invested and investor age are dependent. Further, because there are no liquid mortality sensitive instruments, the resultant risk can be difficult to manage.

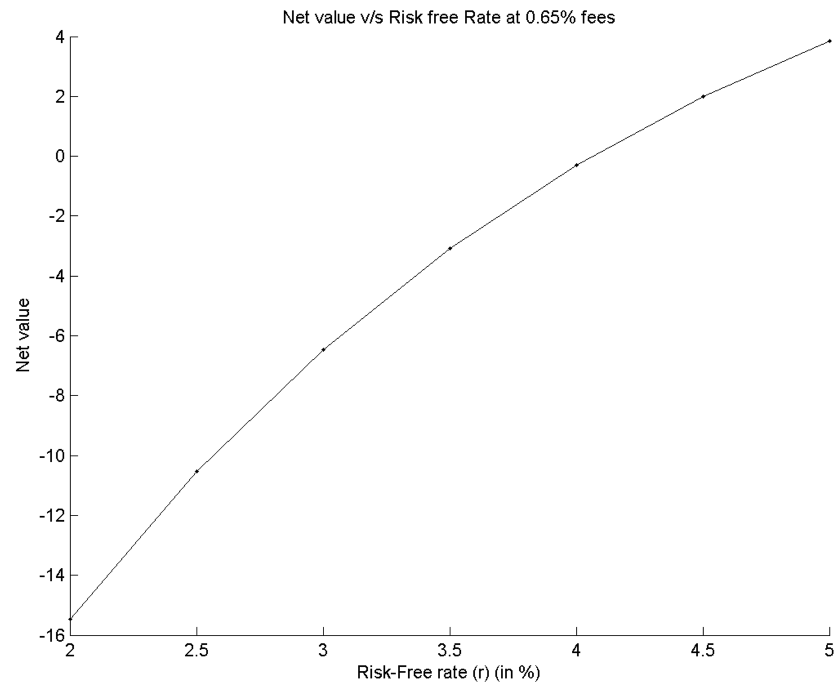

Figure 4-5: Net Value of the meta-portfolio as a function of risk-free rate for the CBSME Model.

4.6 Summary

In this chapter, we introduced the GWB for life feature - an exotic option that insurance companies have been offering to investors to plan for their retirement. The GWB feature allows an individual to stay invested in the market and capture the upside while being assured of a steady income during retirement.

We then used a simplified model to derive an analytical expression for the value of liabilities and revenues due to the GWB option. Using these results, we analyzed the impact of various risk factors on the same for typical values of the problem parameters. This analysis showed:

- The value of the product is sensitive to the choice of the asset mix by the investor. The more aggressive (or volatile) the asset mix, the more expensive is the GWB guarantee.
- The value of GWB is also significantly dependent on the cohort age. In general, the product becomes more expensive for the company (and conversely more attractive to the investor), the shorter the pre-withdrawal period and longer the withdrawal period. The dependence of the break-even fee on investor char-

acteristics might lead to an adverse bias in company's customer profile, unless more price differentiation is provided.

- The product is negatively exposed to volatility in interest rates.
- The product value is quite sensitive to the client pool's mortality distribution. This is a challenge, because this risk can be difficult to hedge in practice.

In the next chapter, we seek to value the GWB in a more realistic setting, taking into account the various risk factors identified in this chapter.

Chapter 5

Guaranteed Lifetime Payouts: Further Analysis

5.1 Introduction

In Chapter 4, we introduced the Guaranteed Withdrawal Benefits (GWB) for life product that insurance companies have been offering off late as a retirement investment solution. We also presented in Section 4.4 of this chapter, the simple CBSME model for analysing the GWB. The various simplifying assumptions - notably Black Scholes asset price dynamics, exponentially distributed retirement lifespans and continuous step-ups helped us to derive an almost closed-form solution for the value of GWB liabilities and revenues.

We shift the focus of this chapter to modeling the GWB product in a more realistic setting. Inevitably, this will come at the cost of the tractability that we achieved with the CBSME model. We first remove the assumptions about continuous step-ups and an exponential distribution for the retirement lifespans. This gives us a base-line valuation of the product in the Black-Scholes model.

Ballotta [9], Lee and Stock [84], Wang, Gerrard and Haberman [115] have pointed out the importance of considering interest rate risks for VA products with embedded options. However, except for the GAO/GMIB type of products, which are primarily interest rate options, valuation impact due to stochastic interest rates is seldom considered. Since GWB for life has long durations, that can easily extend over 30 to 40 years, interest rates can fluctuate significantly during the product's lifetime. The analysis presented in Section 4.5 of Chapter 4 using the CBSME model also indicated a negative exposure to interest rate volatility. Apart from interest rate risks, as the historical analysis presented in Section 4.3 of Chapter 4 revealed, GWB is like a deep-out-of-the money option. Hence, sound modeling of the "tail events" in asset returns is crucial for GWB valuations. As we pointed out in Chapter 4, the GMWB/GWB family has been regarded in the literature as primarily an equity based product and has not really been analysed in the literature beyond the Black-Scholes model, which assumes constant interest rates and a normal distribution for (continuously compounded) asset returns.

In this chapter, we refine the basic Black-Scholes models to investigate the effect of stochasticity in interest-rates and "stochastic volatilities" on GWB pricing. Milevsky and Salisbury [95], Dai, Kwok and Zong [48] and Chen, Vetzal and Forsyth [34] show how fully rational or strategic investor behavior can impact GMWB costs in a Black-Scholes model. We find that even under "typical" investor behavior, accounting for stochasticity of interest rates as well as stochastic volatilities can significantly impact GWB valuation. We also consider an alternate withdrawal behavior for investors where they refuse step-ups in the Phase 2 of the product, but find it to have a limited negative impact on valuation. Like Chapter 4, we also conduct a sensitivity analysis GWB value with respect to investor characteristics such as age and asset-mix choice as well as mortality rates, under different models. We then examine some issues related to hedging the GWB.

5.1.1 Findings and Contributions

In this chapter, we seek to price the GWB in a realistic setting, incorporating the risk-factors identified in Chapter 4. We proceed by taking a typical GWB offering and pricing it under a series of valuation models. We start with the basic Black Scholes framework and then incrementally refine it to account for first, stochastic interest rates and then, stochastic volatilities. We ensure that the models are tuned so that they correspond to the same level of interest rates and option implied volatilities. We find:

- All models corroborate the key inferences made from the simple CBSME model, namely,
  - GWB value is quite sensitive to both the chosen asset mix as well as the age of the investor.
  - GWB has a high sensitivity to the client pool's mortality.

- Accounting for stochasticity in interest rates and return volatilities has a material downside effect on valuations.

- The value put to the GWB has a high dependence on the class of models used. Neither the GWB nor an equity-option with duration anywhere close to the GWB lifespan is actively traded in the market. Any valuation of a GWB product will involve assumptions about markets and investors, not all of which can be inferred from the market data. Hence companies issuing the GWB must value them using their own proprietary models[1]. Given its impressive sales, this "model risk" associated with the GWB implies that there is a high possibility of a significant mis-valuation of the insurance companies' liabilities due to model mis-specification.

- Investor behavior on the other hand, does not seem to impact the cost significantly. This, we believe, can be attributed to the step-up feature and the indexing of fees to the benefit base, which provide a natural hedge against dynamic withdrawals.

- Hedging GWB is challenging because its valuation itself is sensitive to the model family used. Further, it has a high exposure to risk factors such as mortality rates for which effective hedging instruments may not be available. Moreover, on the whole, though net-value seems less sensitive to investor behavior, hedging policy would depend on how customer behavior evolves. We also show that even simple delta-hedging poses challenges, as the step-up feature makes the delta discontinuous.

We find that on the whole, there is considerable ambiguity surrounding the true value of GWB for life. This and a high exposure to risk factors that are not actively traded implies that the GWB for life bears a substantial risk for its underwriters that will be very challenging to manage and hedge against.

5.1.2 Chapter Layout

In this chapter, we consider three arbitrage pricing models to value the GWB for life. These models differ in their assumptions about asset price dynamics in the risk-neutral world. In Section 5.2, we first outline a general pricing framework for the

[1] GWB for life should be a Level 3 asset in the parlance of the FAS 157 standard issued by the Financial Accounting Standards Board (FASB), [55].

GWB. We then describe in detail each of the three models that we use for pricing. In Section 5.3, we present an alternate and dynamic mode of withdrawals by investors and discuss how GWB value can be computed for this withdrawal strategy under the different models considered in this chapter. In Section 5.4, we present and analyze numerical results from valuations under different models and assumptions on withdrawal behavior. This is followed by a discussion on hedging and potential issues therein in Section 5.5. Finally, we summarize the findings and some directions for future research in Section 5.6.

5.2 Valuation Models

In this section, we describe the three asset dynamics models that we use to price the GWB guarantees. The CBSME model in Section 4.4 of Chapter 4 made two rather strong assumptions:

1. Step-ups and cash-flows happen continuously.
2. The residual life of investor in retirement is exponentially distributed.

Step-ups for a GWB type product typically happen annually. Continuous step-ups can significantly distort valuation. The 2^{nd} assumption is also inaccurate as mortality rates typically increase with age.

We relax both these assumptions for the analysis in this chapter.

This means we consider both the step-ups as well as cash-flows to happen discretely, more specifically, annually. The product that we analyze then corresponds exactly to the specifications given in Section 4.2 of Chapter 4. For the analysis in this section, we will assume that the investor withdraws the contractually guaranteed amount at each anniversary. We restate the dynamics in (4.1) of the GWB state variables under this assumption:

$$\begin{aligned} C_{n+1} &= (C_n \cdot R^s_{n+1} - q_{n+1} B_n)^+ \ ; \\ B_{n+1} &= \max(B_n, C_{n+1}) \ , \end{aligned} \tag{5.1}$$

where

$$q_n = \begin{cases} 0 \ , & \text{if } n \le T \ , \\ q \ , & \text{if } n > T \ . \end{cases} \tag{5.2}$$

As before $R^s_{n+1} = \frac{S_{n+1}}{S_n}$ denotes the return on the underlying VA fund for the period $(n, n+1]$.

In addition, we also now allow the investor's residual life, T_A, to have an arbitrary distribution, say $f_{T_A}(\cdot)$. The only assumption we make is that T_A is independent of the market factors.

When the mortality related risk is diversifiable or the insurance company is risk-neutral to such risk (as we assume for our valuation analysis), there are two equivalent ways to account for the randomness in T_A:

1. Find the value of the product as a function of T_A. The fair value of the product is obtained as simply an expectation of this value function under $f_{T_A}(\cdot)$.

2. Model the investor's death as the arrival time of a non-homogeneous Poisson process, i.e., one with time dependent intensity such that the implied distribution of the residual life T_A is $f_{T_A}(\cdot)$. Note for any distribution function $f_{T_A}(\cdot)$, such an intensity process (often known as a mortality or hazard rate) can always be defined[2]. We then find the value of GWB as a function of time using backward substitution.

The second method, as we shall illustrate shortly, has computational advantages. While we do not consider the case of stochastic mortalities here, this framework also allows one to incorporate the same, if desired, more easily (see Biffis [16]).

5.2.1 General Framework for pricing GWB products

Let $L_n(C_n, B_n, \omega)$ and $G_n(C_n, B_n, \omega)$ respectively denote the values of liabilities and revenue base (i.e., revenues normalized by the fee rate h), at time n, *excluding the cash-flows pertaining to the n^{th} anniversary* for a given sample path ω.

Let $r_n^f \triangleq \int_{n-1}^{n} r_s ds$ denote the continuously compounded risk-free rate for the interval $(n-1, n]$ and I_n be the indicator variable that the investor is alive at the end of year n. Then on a given sample path ω,

$$
\begin{aligned}
I_n(\omega)L_n(C_n, B_n, \omega) &= I_{n+1}(\omega) \cdot e^{-r_{n+1}^f} \left((q_{n+1}B_n - C_n R_{n+1}^s)^+ + L_{n+1}(C_{n+1}, B_{n+1}, \omega) \right) \\
I_n(\omega)G_n(C_n, B_n, \omega) &= I_{n+1}(\omega) \cdot e^{-r_{n+1}^f} \left(B_n + G_{n+1}(C_{n+1}, B_{n+1}, \omega) \right) \qquad (5.3)
\end{aligned}
$$

For an investor aged A, we define

$$
\lambda_n^A \triangleq \ln \left(\frac{\mathbb{P}(I_n = 1)}{\mathbb{P}(I_{n-1} = 1)} \right) .
$$

[2]To achieve this, hazard rate at time λ_t at time t is simply set a $\lambda_t \triangleq \frac{f_{T_A}(t)}{1-F_{T_A}(t)}$, with $F_{T_A}(\cdot)$, being the cumulative density function for T_A.

λ_n^A is thus a discrete time "hazard" or mortality rate. It is common to assume that the residual life spans corresponding to different cohorts have stationary distributions. This means that the hazard rate can be expressed as a function of the investor's current age. More specifically, if λ_n is the hazard rate of a population at birth then, $\lambda_n^A = \lambda_{A+n}$. Thus $\mathbb{P}(I_n = 1) = \exp(-\sum_{i=1}^n \lambda_i^A)$. We assume there exists a finite $\overline{N}$ such that $\mathbb{P}(I_{\overline{N}+1} = 1) = 0$, or equivalently, $\lambda_{\overline{N}+1} = \infty$.

We define product epochs as dates where contractual adjustments or cash-flows related to the product occur. From (5.1) and (5.3), it can be seen that all epochs related to the GWB, whether cash-flows in terms of withdrawals by the investor and payment of fees to the company or changes to state variables due to step-ups or withdrawals, happen on contract anniversaries. Hence, any arbitrage pricing model that gives the joint distribution of one year risk free rate r_n^f and R_n^s can be used to price the GWB. Most commonly used asset price and interest rate dynamics models are Markovian. If we are working with such a Markovian model with a state vector Y_t, which by assumption must be independent of mortality dynamics, then it must be possible to express liabilities and revenue stream values as some functions of time and the state variables C_n, B_n and Y_n. Let $L_n(C_n, B_n, Y_n)$ and $G(C_n, B_n, Y_n)$, respectively denote the fair value of liabilities and revenues conditional on the investor being alive at year n. These functions must satisfy the following recursive relations:

$$\begin{aligned} L_n(C_n, B_n, Y_n) &= e^{-\lambda_{n+1}^A} \mathbb{E}^{\mathbb{Q}} \Big[e^{-r_{n+1}^f} \Big\{ \left(q_{n+1} B_n - C_n R_{n+1}^s\right)^+ \\ &\quad + L_{n+1}\left((C_n R_{n+1}^s - q_{n+1} B_n)^+, \max(B_n, C_n R_{n+1}^s - q_{n+1} B_n), Y_{n+1}\right) \Big\} \mid Y_n \Big] , \\ G_n(C_n, B_n, Y_n) &= e^{-\lambda_{n+1}^A} \mathbb{E}^{\mathbb{Q}} \Big[e^{-r_{n+1}^f} \{ 1 \\ &\quad + G_{n+1}\left((C_n R_{n+1}^s - q_{n+1} B_n)^+, \max(B_n, C_n R_{n+1}^s - q_{n+1} B_n), Y_{n+1}\right) \Big\} \mid Y_n \Big] . \end{aligned} \tag{5.4}$$

For compliant investor withdrawals, we have a homogeneity relation along the lines of Proposition 4.1 in Chapter 4.

Proposition 5.1. *If investor withdraws at the contractual rate $q \cdot B_n$, then*

$$\begin{aligned} L_n(C_n, B_n, Y_n) &= B_n \cdot L_n\left(\frac{C_n}{B_n}, 1, Y_n\right) , \\ G_n(C_n, B_n, Y_n) &= B_n \cdot G_n\left(\frac{C_n}{B_n}, 1, Y_n\right) . \end{aligned}$$

Proof. We note that the dynamics of the state variables B_{n+1} and C_{n+1} as described in (5.1) are homogeneous in B_n and C_n. Also, all the cashflows that the insurance company incurs are homogeneous in C_n, B_n. The proof then follows along the same lines as Proposition 4.1 in Chapter 4. □

In light of Proposition 5.1, we define functions $l_n(x, Y_n)$ and $g_n(x, Y_n)$ for $0 \leq x \leq 1$. as follows:

$$l_n(x, Y_n) \triangleq L_n(x, 1, Y_n) \ ,$$
$$g_n(x, Y_n) \triangleq G_n(x, 1, Y_n) \ .$$

From Proposition 5.1 and (5.4), it follows that:

$$\begin{aligned} l_n(x, Y_n) &= e^{-\lambda^A_{n+1}} \cdot \Bigg\{ x \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \left(\frac{q_{n+1}}{x} - R^s_{n+1} \right)^+ \mid Y_n \right] \\ &\quad + \ \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \cdot l_{n+1} \left(x \cdot R^s_{n+1} - q_{n+1}, Y_{n+1} \right) \cdot \mathbf{1}_{\{ \frac{q_{n+1}}{x} < R^s_{n+1} \leq \frac{1+q_{n+1}}{x} \}} \mid Y_n \right] \\ &\quad + x \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \cdot \left(R^s_{n+1} - \frac{q_{n+1}}{x} \right) \cdot l_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{ R^s_{n+1} > \frac{1+q_{n+1}}{x} \}} \mid Y_n \right] \Bigg\} \ . \end{aligned} \tag{5.5}$$

$$\begin{aligned} g_n(x, Y_n) &= e^{-\lambda^A_{n+1}} \cdot \Bigg\{ \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \mid Y_n \right] \\ &\quad + \ \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \cdot g_{n+1} \left(x \cdot R^s_{n+1} - q_{n+1}, Y_{n+1} \right) \cdot \mathbf{1}_{\{ \frac{q_{n+1}}{x} < R^s_{n+1} \leq \frac{1+q_{n+1}}{x} \}} \mid Y_n \right] \\ &\quad + x \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r^f_{n+1}} \cdot \left(R^s_{n+1} - \frac{q_{n+1}}{x} \right) \cdot g_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{ R^s_{n+1} > \frac{1+q_{n+1}}{x} \}} \mid Y_n \right] \Bigg\} \ . \end{aligned} \tag{5.6}$$

The above model is quite generic and will hold for any Markovian asset returns process. If, the asset price returns and interest rates have no memory (such as is the case in the Black Scholes Model), or have dynamics such that the state vector Y_n equilibrates, i.e., reaches a steady state distribution in an interval corresponding to the epoch interval (which is 1 year in our case), then we may drop the conditioning and dependency on Y_n. This leads to a simpler model, which we call the independent period returns model. The simplifications can be explicitly written as:

$$
\begin{aligned}
l_n(x) &= e^{-\lambda_{n+1}^A} \cdot \Bigg\{ x \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} \cdot \left(\frac{q_{n+1}}{x} - R_{n+1}^s \right)^+ \right] \\
&\quad + \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} l_{n+1} \left(x \cdot R_{n+1}^s - q_{n+1} \right) \cdot \mathbf{1}_{\{\frac{q_{n+1}}{x}) < R_{n+1}^s \leq \frac{1+q_{n+1}}{x}\}} \right] \\
&\quad + x \cdot l_{n+1}(1) \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} \cdot \left(R_{n+1}^s - \frac{q_{n+1}}{x} \right)^+ \right] \Bigg\}
\end{aligned}
\tag{5.7}
$$

$$
\begin{aligned}
g_n(x) &= e^{-\lambda_{n+1}^A} \cdot \Bigg\{ \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} \right] \\
&\quad + \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} \cdot g_{n+1} \left(x \cdot R_{n+1}^s - q_{n+1} \right) \cdot \mathbf{1}_{\{\frac{q_{n+1}}{x}) < R_{n+1}^s \leq \frac{1+q_{n+1}}{x}\}} \right] \\
&\quad + x \cdot g_{n+1}(1) \cdot \mathbb{E}^{\mathbb{Q}} \left[e^{-r_{n+1}^f} \cdot \left(R_{n+1}^s - \frac{q_{n+1}}{x} \right)^+ \right] \Bigg\}
\end{aligned}
\tag{5.8}
$$

To use the model in equations (5.7) and (5.8), all one needs is a joint distribution of one period asset returns and interest rates. The Black-Scholes model is just one example of such asset return dynamics. We may use, for example, a more heavy tailed-distribution (to adjust for the observed smiles and skews in option volatilities) for the asset price returns.

In practice, most sophisticated asset return models are Markovian with state vectors that are mean reverting. For GWB, product epochs are contract anniversaries and are thus spaced apart by a year. If the state variables have mean reversion times that are small compared to this interval, then we may assume that the corresponding Markov process has equilibrated by the next epoch and thus approximate the returns over different periods as independent.

A drawback with this approach is that it cannot be used for mid-epoch-interval valuation, i.e., when we are trying to find the value of the product between epochs, especially when a contract anniversary is imminent. Mid-epoch-interval valuation is important if the GWB is to be hedged dynamically with frequent trading. However, this shortcoming can be easily dealt with by invoking the stationarity assumption only after a certain duration, e.g, after the second product epoch (or the one after the closest product epoch).

State variables for some models such as, equity stochastic volatility models sometimes have relatively fast mean reversion, (see for example, Dragulescu and Yakovenko [50]) compared to the one year time-frame. Unfortunately, however many interest rate models used in practice have dynamics that involve variations over long periods[3]. For example, multi-factor affine rate models have long equilibrating periods and hence state vectors associated with them would have an impact on product valuations. In this case, we need to preserve the state of the system.

5.2.2 Pricing Models for the GWB for life

We now consider three different models to price GWB for life. We assume that the underlying VA fund is invested in a mix of equity and bonds and is continuously rebalanced to maintain a fixed proportion by value between these two asset classes. The equity exposure is the key source of volatility in the account. We take the relatively less volatile bonds asset class to have excess returns (over the risk-free rates) that are log-normal and independent of the risk-free rate and equity returns in all the models considered. For the joint dynamics of the risk-free rate and excess equity returns, we consider three different models:

BSM model : In this we assume that the interest rates are non-stochastic and constant and the equity returns are lognormal. Since the bond returns are also log-normal and the VA fund is continuously rebalanced to hold a constant proportion by value of equities and bonds, it will also have log-normal returns in this model. Thus, we do not have any state variables in the BSM model.

SILN model : In this model, we allow for interest rates to be stochastic and use a two factor Vasicek model, which is an extension of the model proposed by Vasicek [113] for interest rates. We continue to assume log-normal excess returns for equities and by extension for the VA fund.

SISV model : Here, the interest rates are modeled by a two-factor Vasicek model as in the SILN model. In addition, we use the Heston model (see Heston [68]), which is a stochastic volatility model for excess equity returns.

We do not consider the effect of jumps in asset prices which is sometimes important for short-duration derivatives. However, we do not believe this to be significant for the GWB, as jump effects are likely to be diffused out over the interval between two epochs[4]. Also models that employ jumps are difficult to calibrate in practice.

[3]This may be an artifact of the fact that stochastic volatility models are calibrated using options, which have a much smaller duration, typically less than a year, while interest rate models are calibrated using instruments with substantially longer durations.

[4]Chen, Vetzal and Forsyth [34] report that jumps significantly increase the product price for the GMWB. However, they add jumps on top of the diffusion process and do not attribute how much of

The risk due to jumps in asset prices becomes significant when one is very close to a contract anniversary, especially if a step-up is imminent. However, the probability of such an event should be small. We refer to Bakshi, Cao and Chen [8], Biffis and Millossovich [17] for more information on stochastic jump based models.

We now describe each of the three models presented earlier in more detail. For notational convenience, we define an equity market index E_t.

Black Scholes Asset Dynamics - BSM Model

In this section, we consider the valuation of GWB in the Black-Scholes setting, i.e., the risk-free rates $r(t)$ and the equity market volatility $\sigma_e(t)$ are known as functions of time. Since the underlying VA fund is continuously rebalanced to hold a constant proportion of stocks and bonds, it will also follow a geometric Brownian motion. Further, the volatility of the returns in the VA fund will also be a known function of time. Thus S_t evolves as:

$$\frac{dS_t}{S_t} = r(t)dt + \sigma(t)dZ_t^Q \tag{5.9}$$

in the risk neutral world. For brevity, henceforth we shall refer to this model as the BSM model. As discussed earlier, the asset returns in Black-Scholes world are independent (in fact, returns over arbitrarily small periods that are disjoint are independent) and there is no state-vector associated with this model. Using (5.7) in this context, we get

$$\begin{aligned} l_n(x) \;=\; e^{-(\lambda^A_{n+1}+r^f_{n+1})}\cdot\Bigg\{ & x\cdot\int_{-\infty}^{\ln\left(\frac{q_{n+1}}{x}\right)}\left(\frac{q_{n+1}}{x}-e^z\right)\cdot\Phi\left(z;\mu_{n+1},\sigma_{n+1}\right)dz \\ & + \int_{\ln\left(\frac{q_{n+1}}{x}\right)}^{\ln\left(\frac{1+q_{n+1}}{x}\right)} l_{n+1}\left(xe^z-q_{n+1}\right)\cdot\Phi\left(z;\mu_{n+1},\sigma_{n+1}\right)dz \\ & + x\cdot l_{n+1}(1)\cdot\int_{\ln\left(\frac{1+q_{n+1}}{x}\right)}^{\infty}\left(e^z-\frac{q_{n+1}}{x}\right)\cdot\Phi\left(z;\mu_{n+1},\sigma_{n+1}\right)dz\Bigg\}\;. \end{aligned} \tag{5.10}$$

the increase can be attributed to an increase in the effective volatility of the diffusion process and how much is due to the jump effects.

Here, $\sigma^2_{n+1} \triangleq \int_n^{n+1} \sigma^2(s)ds$ and $\mu_{n+1} \triangleq r_{n+1} - \frac{1}{2}\sigma^2_{n+1}$ and $\Phi(z;\mu,\sigma)$ denotes the normal density function for mean μ and variance σ^2. For $g_n(x)$, we have a similar relation:

$$
\begin{aligned}
g_n(x) \quad = \quad & e^{-(\lambda^A_{n+1}+r^f_{n+1})} \cdot \left\{ 1 + \int_{\ln\left(\frac{q_{n+1}}{x}\right)}^{\ln\left(\frac{1+q_{n+1}}{x}\right)} g_{n+1}\left(xe^z - q_{n+1}\right) \cdot \Phi\left(z;\mu_{n+1},\sigma_{n+1}\right) dz \right. \\
& \left. + \, x \cdot g_{n+1}(1) \cdot \int_{\ln\left(\frac{1+q_{n+1}}{x}\right)}^{\infty} \left(e^z - \frac{q_{n+1}}{x}\right) \cdot \Phi\left(z;\mu_{n+1},\sigma_{n+1}\right) dz \right\} .
\end{aligned}
\tag{5.11}
$$

We also have the boundary conditions:

$$
\begin{aligned}
l_{\overline{N}} &= 0 , \\
g_{\overline{N}} &= 0 .
\end{aligned}
\tag{5.12}
$$

Using the integral equations in (5.10), (5.11) together with (5.12), one can recursively solve backwards for the function $l_n(\cdot)$, $g_n(\cdot)$. In fact, the first and the last expectation terms in (5.10) and the last term in (5.11), can be computed using the Black-Scholes option formulae for call, put and digital options. However, the central term in these equations seems to render an exact closed form solution difficult, though very good approximations can be computed quickly using interpolation.

In Section G.1 in Appendix G, we provide further details on how we use equations (5.10) and (5.10) for numerical computations.

Remark 5.1. *Note that if we had an exponential residual life distribution, i.e.,* λ^A_n *were constant for* $n \geq T$ *and the risk-free rates* r^f_n *and volatilities* σ^f_n *were also constants, then the functions* $l_n(\cdot)$ *and* $g_n(\cdot)$ *would be identical for* $n \geq T_0$ *and these could be solved using a recursive version of equations (5.10) and (5.11). Integral equations of these kind are known as Fernholz Integral Equations of the* 2^{nd} *kind. These equations are counterparts of the equations (4.16) and (4.17) in Chapter 4. The key change is that the boundary condition (4.11) that applied to the latter, was a consequence of continuous step-ups and does not hold in the current setting.*

Stochastic Interest Rates and Lognormal Equity Returns - SILN Model

Through this model, we consider the impact of stochasticity of interest rates of GWB valuation. For this, we use a two-factor Vasicek Model for the short-rate, as suggested in [39] for GAO pricing. A 2-factor model is preferred over the simpler one factor interest rate models, because in a one factor model, all forward interest rates become fully correlated. The two factor model is the simplest model, which allows for a relative movement between different points in the yield curve. The short rate model (under the risk-neutral measure) is given by

$$\begin{aligned} r_t &= x_t^1 + x_t^2 + b(t) \ ; \\ dx_t^1 &= -\kappa_1 x_t^1 + \sigma_1 dZ_t^1 \ ; \\ dx_t^2 &= -\kappa_2 x_t^2 + \sigma_1 dZ_t^2 \ ; \\ < dZ_t^1 \cdot dZ_t^2 > &= \rho dt \ . \end{aligned} \tag{5.13}$$

The function $b(t)$ is deterministic and chosen so as to agree with initially observed yield curve in the market. Our model for the equity returns is

$$\frac{dE_t}{E_t} = r_t dt + \sigma_e dZ_t^e \ . \tag{5.14}$$

Z_t^e is assumed to be independent of the factors driving the short rate. Thus, in essence we assume that the excess equity returns are independent of the short rate process and has a log-normal distribution. Again, as the excess bond returns are also assumed to be log-normal and as the VA fund is continuously rebalanced, it should also have log-normal excess returns. We will refer to this model as the Stochastic Interest Rate, Lognormal Excess returns Model or SILN model for short. For the SILN model, the state variable Y_n is the two dimensional vector $(x_{1,n}, x_{2,n})$.

The short rate model, considered here is a Gaussian model, and entails a small but non-zero probability of the short-rate becoming negative. A way to avoid this problem is to instead use the Cox-Ingersoll-Ross process for short rates, see Cox, Ingersoll and Ross [44] and Duffie [51].

The Gaussian model however has the advantage that the Stochastic Differential Equation in (5.13) has a solution, which makes valuation computations using simulations much faster. More details on how this can be used for faster GWB valuations, as well as on how $b(t)$ is adjusted to fit the initial yield curve are provided in Section G.1 in Appendix G.

Stochastic Interest Rates and Stochastic Volatility - SISV Model

In the SISV model, we take both interest rates and equity volatilities to be stochastic. For interest rates, we use the same 2-factor Gaussian Model as described by (5.13) for the SILN model. We model excess equity market retuns to follow the Heston process (see Heston [68]) under the risk neutral measure:

$$
\begin{aligned}
\frac{dE_t}{E_t} &= r_t dt + \sqrt{V_t} dZ_t^e \ ; \\
dV_t &= \kappa_e(\theta_s - V_t) + \xi \sqrt{V_t} dZ_t^\sigma \ ; \\
< dZ_t^Q \cdot dZ_t^\sigma > &= \rho_e dt \ .
\end{aligned} \tag{5.15}
$$

Z_t^e and Z_t^σ are assumed to be independent of the short rate factors. The Heston Model is a stochastic volatility model and can capture the skew and smile effects seen in the implied volatility curves. In addition, because of the square-root term in (5.15), it does not allow the variance to go negative. We re-emphasize that even for the SISV model, we assume that the excess bond returns are log-normal. Thus the Heston model is used only to model excess equity returns[5]. Section G.3 in Appendix G briefly outlines how we use the SISV model in numerical valuations of GWB.

5.3 Valuation under Alternate Withdrawal Strategy

So far, we assumed that the investor continues to withdraw anually at the product stipulated rate q, that will allow her to capture the step-ups. In reality investors do have some leeway in selecting their withdrawal patterns and this can impact the cost of the GWB to the insurance company. For example, Milevsky and Salisbury [95] analyze that for the related GMWB, strategic withdrawals by investors can increase the break-even fee by over two-folds over the one obtained assuming a passive withdrawal scheme. "Optimal" withdrawal schemes usually require the investor to take large excess withdrawals in many scenarios (see for example, Dai, Kwok and Zong [48]). As remarked in Section 4.2, Chapter 4, we believe this is likely to be sub-optimal in practice because large excess withdrawals incur for the investor not only an imposition of a moderate surrender charge or an excess withdrawal penalty by the insurance company but also possibly more severe indirect costs in terms of tax payments. Given that the only incentive for excess withdrawals will be to save on future GWB premium (which is typically small at few tens of basis points), it is unlikely that in the presence of reasonably high indirect costs due to taxes, a rational investor will make the large excess withdrawals that can hurt the insurance company

[5]Note that the returns on a fixed proportion by value portfolio of multiple assets do not follow a Heston process, even when returns for each individual asset are modeled to follow one.

severely.

Generally, excess withdrawal fees are not charged for withdrawing gains in the contract portfolio value. Considering this and the fact that withdrawing during a down market actually will cause the benefit payments to reset to a lower value and is likely to be sub-optimal, we consider the following alternative withdrawal strategy for the investor

- The investor does not take any withdrawals in Phase 1, i.e., upto the first T years.
- In Phase 2, i.e., after time T, the investor takes withdrawals in a way so as to sustain the guaranteed payment level set at time T.

The rationale behind this "excess withdrawals only during up-markets" is that by avoiding any excess withdrawals that will cause the benefit base to fall, the investor does not let the level of the guaranteed withdrawals to drop in retirement. On the other hand, by not allowing the same to rise in Phase 2 years, the investor does not allow the fees to rise in the up-market scenarios, where the GWB protection is unlikely to be exercised. This is because the average remaining duration of the GWB guarantee decreases every year, once in Phase 2 and the probability of a shortfall given an account value that fully capitalizes the benefit base is also declining as a result. The investor allows for step-ups to take place in the Phase 1 years, because when in Phase 1, the average duration of the guarantee does not decrease at the next anniversary.

The dynamics of the state variables C_n and B_n under this "dynamic withdrawal policy" are given by:

$$\begin{aligned} W_n &= \begin{cases} 0 & \dots \quad n \leq T \ , \\ \max(q \cdot B_{n-1}, C_{n-1}(R_n^s - 1)) & \dots \quad n > T \ . \end{cases} \qquad (5.16) \\ C_n &= (C_{n-1} \cdot R_{n-1}^s - W_{n-1})^+ \\ B_n &= \begin{cases} \max(B_{n-1}, C_{n-1} \cdot R_n^s) & \dots \quad n \leq T \ , \\ B_{n-1} & \dots \quad n > T \ . \end{cases} \end{aligned}$$

W_n represents the withdrawal made on the n^{th} anniversary.

Again, if one considers a Markovian asset dynamics model with the state variable Y_n, the value of future liabilities and revenue streams under this modified withdrawal policy at time n will be given by some functions $\tilde{L}(C_n, B_n, Y_n)$ and $\tilde{G}(C_n, B_n, Y_n)$ respectively.

Further, because the dynamics as given in (5.17) are again homogeneous in the state variables, it can be easily verified that Proposition 5.1 holds for $\tilde{L}(C_n, B_n, Y_n)$ and $\tilde{G}(C_n, B_n, Y_n)$ as well and

$$\begin{aligned}\tilde{L}(C_n, B_n, Y_n) &= B_n \cdot \tilde{L}\left(\frac{C_n}{B_n}, 1, Y_n\right) ; \\ \tilde{G}(C_n, B_n) &= B_n \cdot \tilde{G}\left(\frac{C_n}{B_n}, 1, Y_n\right) .\end{aligned}$$

Let $\tilde{l}_n(x) \triangleq L_n(x, 1, Y_n)$ and $\tilde{g}_n(x) \triangleq G_n(x, 1, Y_n)$. Then, we have:

$$\tilde{l}_n(x, Y_n) = \begin{cases} e^{-\lambda^A_{n+1}} \cdot \Big\{\mathbb{E}^{\mathbb{Q}}[e^{-r^f_{n+1}} \cdot \tilde{l}_{n+1}(xR^s_{n+1}, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} \le \frac{1}{x}\}}] \\ \qquad + x \cdot \mathbb{E}^{\mathbb{Q}}[e^{-r^f_{n+1}} \cdot R^s_{n+1} \cdot \tilde{l}_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} > \frac{1}{x}\}} |Y_n]\Big\} , & \text{if} \quad n \le T ; \\ \\ e^{-\lambda^A_{n+1}} \cdot \Big\{x \cdot \mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \left(\frac{q}{x} - R^s_{n+1}\right)^+ \mid Y_n\right] \\ \qquad + \mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \cdot \tilde{l}_{n+1}\left(xR^s_{n+1}, Y_{n+1}\right) \cdot \mathbf{1}_{\{\frac{q}{x} < R^s_{n+1} \le \frac{1+q}{x}\}} \mid Y_n\right] \\ \qquad + \mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \cdot \tilde{l}_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} > \frac{1+q}{x}\}} \mid Y_n\right]\Big\} , & \text{if} \quad n > T. \end{cases} \tag{5.17}$$

$$\tilde{g}_n(x, Y_n) = \begin{cases} e^{-\lambda^A_{n+1}} \cdot \Big\{\mathbb{E}^{\mathbb{Q}}[e^{-r^f_{n+1}} |Y_n] + \\ \qquad + \mathbb{E}^{\mathbb{Q}}[e^{-r^f_{n+1}} \cdot \tilde{g}_{n+1}(xR^s_{n+1}, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} \le \frac{1}{x}\}}] \\ \qquad + x \cdot \mathbb{E}^{\mathbb{Q}}[e^{-r^f_{n+1}} \cdot R^s_{n+1} \cdot \tilde{g}_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} > \frac{1}{x}\}} |Y_n]\Big\} , & \text{if} \quad n \le T; \\ \\ e^{-\lambda^A_{n+1}} \cdot \Big\{\mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \mid Y_n\right] \\ \qquad + \mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \cdot \tilde{g}_{n+1}\left(xR^s_{n+1} - q, Y_{n+1}\right) \cdot \mathbf{1}_{\{\frac{q}{x} < R^s_{n+1} \le \frac{1+q}{x}\}} \mid Y_n\right] \\ \qquad + \mathbb{E}^{\mathbb{Q}}\left[e^{-r^f_{n+1}} \cdot \tilde{g}_{n+1}(1, Y_{n+1}) \cdot \mathbf{1}_{\{R^s_{n+1} > \frac{1+q}{x}\}} \mid Y_n\right]\Big\} & \text{if} \quad n > T . \end{cases} \tag{5.18}$$

Relations in (5.17) and (5.18) can be used to obtain GWB valuations for any Markovian model for asset price dynamics, including the BSM, SILN and SISV models described earlier.

We now turn to a numerical analysis of GWB valuations and sensitivity analysis under different models and withdrawal strategies.

5.4 Numerical Results

We use the models described in Sections 5.2 to price the GWB numerically. The values of different parameters related to GWB terms and investor characteristics are the same as in Section 4.5 of Chapter 4 and are restated below:

Product Parameters

The guaranteed withdrawal rate q is set to 6%. The minimum waiting period W is taken to be 3 years, and the retirement age A_R to be 65 years. For computing net value of the GWB, we will assume that the fees are charged at the rate h= 0.65% of the benefit base per year and take a reference initial investment of $C_0 = 100$. Thus the net value figures that we obtain can be interpreted to have units of percentage of sales of VAs for which the GWB for life was opted.

Investors' Profile

We perform a valuation of GWB across a cross-section of investor ages and asset mix choices as in Section 4.5 of Chapter 4.

- We consider investor age at inception, i.e., A to vary in the range 50 to 70 years.
- Investors choose an asset-mix for investment by selecting a level for α, which is the exposure (by value) their VA account will have to equities. We consider four levels for α - $20\%, 40\%, 60\%$ and 80%. The balance of the portfolio will be invested in bonds.
- Finally, we consider the average value of GWB for the insurance company for sales across different investor cohorts and investment plans. For this, we assume that clients' age at inception, A, weighted by their initial investment amount is uniformly distributed in the range 50 to 70. Further, each investor chooses the values of α - $20\%, 40\%, 60\%$ and 80% with probabilities $0.1, 0.4, 0.4$ and 0.1 respectively. We refer to a portfolio of clients with the stated distribution of age and asset-mix selection as the "meta-portfolio".

Mortality Rates

We compute the mortality rates λ_n from the table published by the Pension Benefit Guaranty Corporation (PGBC) in [101]. Relevant values from this table are also listed in Appendix F for reference. As before we use de-Moivre's approximation to convert these mortality rates into hazard rates[6]. For sensitivity analysis with respect to mortality we consider two cases:

- Population mortality rates shift to the right by 1 year[7] to get the new mortality rates as $\lambda_n^- = \lambda_{n+1}$, with $\lambda_m^- = \infty$ for $m \geq \overline{N}$.
- Population mortality rates shifts to the left by 1 year to get the new mortality rates as $\lambda_n^+ = \lambda_{n-1}$.

Asset Dynamic Models

As indicated before, we use three different asset dynamic models for pricing GWB. For each of these models, we assume that excess returns for bonds are lognormally distributed and are independent of equity and interest rate dynamics. Also for all the models considered, we assume that excess equity returns are independent of interest rates. The structural parameters for various models are as follows:

BSM Model : We take the functions $r(t)$ and $\sigma(t)$ in (5.9) to be constant at 3.5% and 20% respectively. Both numbers are annualized values. In addition, we take the volatility of bond returns to be $\sigma_b = 2\%$. The effective volatility of the VA fund returns for a given level of α is then given by (4.30).

SILN Model : For the interest rate dynamics in (5.13), we use the same structural parameter values as specified in [39] i.e., $\kappa_1 = 0.77$, $\kappa_2 = 0.08$, $\sigma_1 = 2\%$, $\sigma_2 = 1\%$ and $\rho = -0.7$. The function $b(t)$ is deterministic and chosen so as to agree with the initially observed yield curve in the market. To make this comparable to the BSM model, unlike Chu and Kwok [39], who prefer to choose an upwardly sloping curve as the initial yield curve, we use a flat initial yield curve, $\gamma_0(t) = 3.5\%$. The parameter values chosen translate to about 0.63% annualized volatility in one year interest rates. This volatility would also contribute to the total volatility in equity and bond returns. Hence, for the accompanying equity rate dynamics as specified in (5.14), we choose a slightly smaller value of σ_e, to ensure that the one-year total volatility in equity returns is the same as that for the BSM model i.e., 20%. This value of σ_e turns out to be 19.99% and is not much different from the one used in the BSM model. We similarly adjust the volatility in excess bond returns down to $\sigma_b = 1.90\%$

[6]Because we effectively use a discrete time framework with one year time steps for valuing GWB, one year mortality rates and hazard rates can actually be used interchangeably.

[7]Note that this operation leads to a decrease in longevities.

so that the total one year return volatility matches that in the case of the BSM model.

SISV Model : Since the interest rate dynamics for the SISV model are the same as in SILN model, we continue to use the same values for parameters κ_1, κ_2, σ_1, σ_2, ρ and $b(t)$ as in the SILN nodel. For the structural parameters in the Heston model (5.15), we set their values close to those suggested by Bakshi, Cao and Chen in [7]. In particular, we take $\kappa_e = 1$, $\xi = 0.42$, $\rho_e = -0.75$. The values of ρ_e, ξ and κ_e basically give the volatility smile curve its shape, with ρ_e primarily influencing the the skew and ξ the curvature. Various empirical estimations of the Heston model from market data report parameter values for ρ_e and ξ close to the ones that we use, see Bakshi, Cao and Chen [7, 8], Moodley [96], Zhang and Shu [119]. There seems to be a large variation in reported values of κ_e, see Dragulescu and Yakovenko [50]. The wide range of values is indicative of the practical issues in estimating or calibrating parameter values for sophisticated models such as the Heston model. The value of κ_e, that we choose is at the lower end of this spectrum. After fixing these parameters, we impose the constraint $V_0 = \theta$, and simultaneously adjust them so that the Black-Scholes implied volatility[8] for a 1-year ATM option, assuming a fixed interest rate of 3.5% is 20%. This yields a value of $\theta = V_0 = 0.0496$. We again adjust the volatility in excess bond returns down to $\sigma_b = 1.90\%$ so that the total one year return volatility matches that in the case of the BSM model.

Additional details about computational procedures for each model are provided in Appendix G.

Remark 5.2. *In practice, estimation of model parameters for asset-dynamic models is a challenging problem in itself, especially for complex models such as the two-factor Gaussian short rate model and the Heston Model. These are usually computed by a so-called calibration of the models, so that the model implied prices agree with observed market prices of common liquid instruments. Since the number of instruments used for calibration is typically much greater than that of model parameters, "pure" models would rarely fit well to explain all the observed market prices. Moreover, the best-fitting parameter values for one day are unlikely to be same as the ones the next day. Hence model parameter values are often estimated by using both cross-sectional*

[8] We do not consider the effect of interest rate volatilities here, but as discussed in the context of SILN model, do not expect this to be significant.

data of market instruments and time-series data, see Bakshi, Cao and Chen [7]. This presents some difficulties, as while the model is indicative of asset dynamics under the risk-neutral measure, the time series data is representative of the real world dynamics.

For the case of the GWB, this task is especially challenging, as an ideal model should represent risk-neutral dynamics over a very long time-frame.

Finally, interest rate models and stochastic volatility models are likely to be calibrated independently in practice as they serve to price mostly non-overlapping universes of securities. An estimation of market implied correlation between the factors driving these models is in most cases difficult because of the lack of liquid securities having exposure to both. For our analysis here, we have ignored these complexities of model estimations, and chosen a set of "typical" values for the more structural parameters while choosing the values of other parameters so that the models agree on the annual volatility of the equity index and the initial yield curve.

5.4.1 Valuation and Impact of Model Selection

We see that the choice of model used to price the GWB guarantee has a substantial impact on the valuation. Table 5.1 shows the break-even fees while Table 5.2 shows the net value at 0.65% fees for various combinations of investor age at inception and asset-mix selection. As compared to the continuous step-ups and exponential residual mortality rates framework of Chapter 4, the break even fees are much smaller and the net value much higher. This can be primarily attributed to the discreteness of step-ups.

Accounting for interest rate stochasticity (SILN model) substantially reduces the valuations over the BSM model. The break-even fee for the meta portfolio increases from 47 basis points under the BSM model to 55 basis points under the SILN model, while the net value decreases from 3.75% of sales to 2.11% of sales. These decreases are driven by interest rate volatility effects. Figures 5-1 and 5-2 show how respectively the break even fee and net value for the meta-portfolio change with interest rate r for the BSM model. As in the case of the CBSME model, the relationship is convex and indicates that interest rate volatility will make the GWB offering more expensive for the insurance company.

Incorporating stochastic volatility through the SISV model leads to a further significant decrease in valuations over the SILN model. These reductions are again uniform across all cohorts and asset mix selections. Thus stochastic interest rates and stochastic volatilities are both significant effects in the case of GWB for life.

		Contract Specified Withdrawals			***Dynamic Withdrawals***		
Cohort Age	**Equity Exposure**	**BSM**	**SILN**	**SISV**	**BSM**	**SILN**	**SISV**
50	20%	0.16%	0.26%	0.27%	0.16%	0.26%	0.27%
	40%	0.29%	0.39%	0.42%	0.29%	0.39%	0.43%
	60%	0.49%	0.57%	0.62%	0.49%	0.58%	0.62%
	80%	0.71%	0.78%	0.83%	0.71%	0.79%	0.83%
55	20%	0.19%	0.30%	0.32%	0.19%	0.30%	0.32%
	40%	0.35%	0.45%	0.49%	0.35%	0.46%	0.49%
	60%	0.57%	0.66%	0.70%	0.58%	0.66%	0.70%
	80%	0.81%	0.89%	0.92%	0.82%	0.90%	0.93%
60	20%	0.24%	0.36%	0.37%	0.24%	0.36%	0.37%
	40%	0.43%	0.53%	0.55%	0.43%	0.53%	0.56%
	60%	0.67%	0.75%	0.78%	0.68%	0.76%	0.80%
	80%	0.93%	1.00%	1.02%	0.95%	1.02%	1.04%
65	20%	0.17%	0.26%	0.27%	0.17%	0.26%	0.28%
	40%	0.34%	0.41%	0.44%	0.34%	0.42%	0.45%
	60%	0.56%	0.62%	0.66%	0.57%	0.63%	0.67%
	80%	0.79%	0.85%	0.89%	0.81%	0.87%	0.91%
70	20%	0.06%	0.12%	0.13%	0.06%	0.12%	0.13%
	40%	0.17%	0.22%	0.25%	0.17%	0.22%	0.26%
	60%	0.32%	0.37%	0.43%	0.33%	0.38%	0.44%
	80%	0.51%	0.55%	0.62%	0.53%	0.57%	0.64%
Meta	Meta	0.47%	0.55%	0.58%	0.47%	0.55%	0.59%

Table 5.1: Break-even fees under different models.

		Contract Specified Withdrawals			***Dynamic Withdrawals***		
Cohort Age	**Equity Exposure**	**BSM**	**SILN**	**SISV**	**BSM**	**SILN**	**SISV**
50.00	0.20	13.47	10.90	10.52	13.40	10.74	10.37
	0.40	10.17	7.63	6.63	10.01	7.39	6.42
	0.60	4.86	2.42	0.97	4.75	2.24	0.82
	0.80	-2.03	-4.36	-5.88	-1.89	-4.30	-5.82
55.00	0.20	10.41	7.98	7.66	10.34	7.84	7.53
	0.40	7.05	4.78	3.98	6.89	4.57	3.79
	0.60	1.93	-0.19	-1.21	1.82	-0.33	-1.36
	0.80	-4.48	-6.49	-7.40	-4.33	-6.35	-7.34
60.00	0.20	7.41	5.35	5.15	7.33	5.23	5.04
	0.40	4.14	2.35	1.84	3.98	2.17	1.67
	0.60	-0.47	-2.09	-2.70	-0.59	-2.19	-2.83
	0.80	-6.01	-7.53	-7.99	-5.85	-7.33	-7.89
65.00	0.20	7.20	5.89	5.71	7.12	5.79	5.63
	0.40	4.90	3.74	3.26	4.69	3.53	3.08
	0.60	1.57	0.51	-0.16	1.32	0.29	-0.37
	0.80	-2.52	-3.52	-4.17	-2.61	-3.56	-4.26
70.00	0.20	7.49	6.88	6.74	7.42	6.79	6.67
	0.40	6.43	5.82	5.32	6.21	5.59	5.14
	0.60	4.57	3.96	3.13	4.23	3.64	2.87
	0.80	2.02	1.43	0.41	1.67	1.11	0.14
Meta	Meta	3.75	2.11	1.40	3.61	1.95	1.25

Table 5.2: Net value under different models.

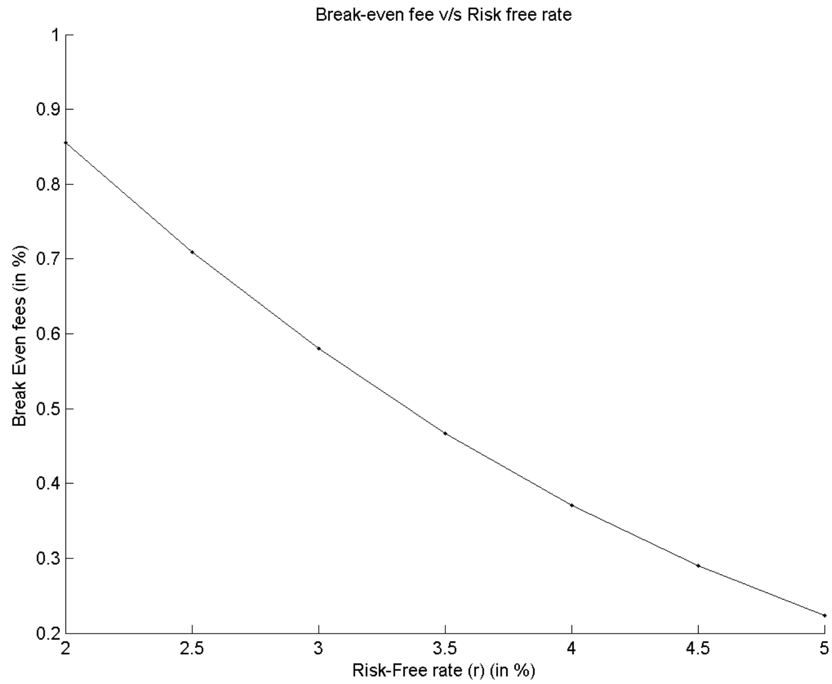

Figure 5-1: Break even fee for the meta portfolio as a function of risk-free rate for the BSM model.

On the other hand, if investors follow a dynamic withdrawal policy, there is a relatively small effect on valuations. The net values on the whole again decrease when compared to those obtained under a contractual withdrawal policy assumption.

We also see that changes in the break-even fees are somewhat more subdued as compared to the changes in net-value with a change in withdrawal policies. This is because changing the withdrawal behavior from contractual to a dynamic policy decreases both the payout liabilities and the revenue streams. Their net result in this case is typically a very small increase in break-even fees. These differences are revealed more clearly in Tables 5.3 and 5.4, which list the liabilities and revenue streams separately. Thus, although a dynamic withdrawal behavior might not impact net-valuation much, it still can have an important bearing on hedging. Tables 5.3 and 5.4 also show that, the revenue-stream valuations are considerably stable across different models (with a difference of less than 30 basis points on sales.) However, the value of the guarantee can swing by as much as 2.5 percentage points on sales. This suggests that model selection will have an effect not only on the valuations, but also on capital requirements for prudent risk and liquidity management.

Since, the dynamic withdrawal policy has on the whole a relatively small effect on

		Contract Specified Withdrawals			***Dynamic Withdrawals***		
Cohort Age	**Equity Exposure**	**BSM**	**SILN**	**SISV**	**BSM**	**SILN**	**SISV**
50	20%	4.25	7.07	7.56	4.21	7.02	7.52
	40%	8.43	11.25	12.43	8.23	11.05	12.27
	60%	15.02	17.74	19.35	14.48	17.20	18.90
	80%	23.43	26.03	27.62	22.32	24.94	26.67
55	20%	4.28	6.94	7.32	4.24	6.89	7.29
	40%	8.38	10.89	11.77	8.18	10.67	11.61
	60%	14.53	16.88	17.93	13.98	16.31	17.48
	80%	22.12	24.35	25.21	20.97	23.18	24.25
60	20%	4.35	6.59	6.80	4.31	6.54	6.77
	40%	8.20	10.18	10.67	7.99	9.94	10.50
	60%	13.58	15.37	15.91	12.99	14.74	15.41
	80%	19.99	21.66	21.97	18.72	20.37	20.90
65	20%	2.55	3.98	4.16	2.52	3.95	4.13
	40%	5.31	6.59	7.03	5.17	6.42	6.91
	60%	9.23	10.40	10.97	8.81	9.94	10.61
	80%	13.96	15.08	15.56	13.04	14.09	14.77
70	20%	0.79	1.49	1.63	0.79	1.48	1.62
	40%	2.22	2.91	3.38	2.17	2.85	3.33
	60%	4.55	5.23	5.98	4.38	5.04	5.82
	80%	7.60	8.26	9.15	7.18	7.81	8.77
Meta	Meta	9.58	11.40	12.09	9.21	11.00	11.78

Table 5.3: Value of insurance company liabilities from GWB for life under different models.

		Contract Specified Withdrawals			***Dynamic Withdrawals***		
Cohort Age	**Equity Exposure**	**BSM**	**SILN**	**SISV**	**BSM**	**SILN**	**SISV**
50	20%	17.72	17.97	18.07	17.61	17.76	17.89
	40%	18.60	18.88	19.06	18.24	18.44	18.69
	60%	19.88	20.16	20.32	19.23	19.44	19.71
	80%	21.40	21.67	21.74	20.43	20.64	20.85
55	20%	14.69	14.93	14.98	14.58	14.73	14.82
	40%	15.43	15.67	15.75	15.07	15.24	15.40
	60%	16.46	16.69	16.72	15.80	15.97	16.12
	80%	17.65	17.86	17.81	16.64	16.83	16.91
60	20%	11.75	11.94	11.95	11.64	11.77	11.81
	40%	12.34	12.52	12.50	11.96	12.10	12.17
	60%	13.11	13.28	13.21	12.40	12.54	12.58
	80%	13.98	14.13	13.98	12.88	13.03	13.02
65	20%	9.74	9.87	9.87	9.64	9.73	9.76
	40%	10.21	10.33	10.28	9.86	9.95	9.99
	60%	10.80	10.91	10.81	10.13	10.23	10.24
	80%	11.44	11.55	11.39	10.42	10.52	10.51
70	20%	8.29	8.37	8.37	8.21	8.27	8.29
	40%	8.65	8.73	8.70	8.39	8.45	8.48
	60%	9.11	9.19	9.11	8.61	8.68	8.69
	80%	9.62	9.69	9.56	8.85	8.92	8.91
Meta	Meta	13.33	13.51	13.50	12.82	12.95	13.03

Table 5.4: Value of revenues from GWB for life under different models.

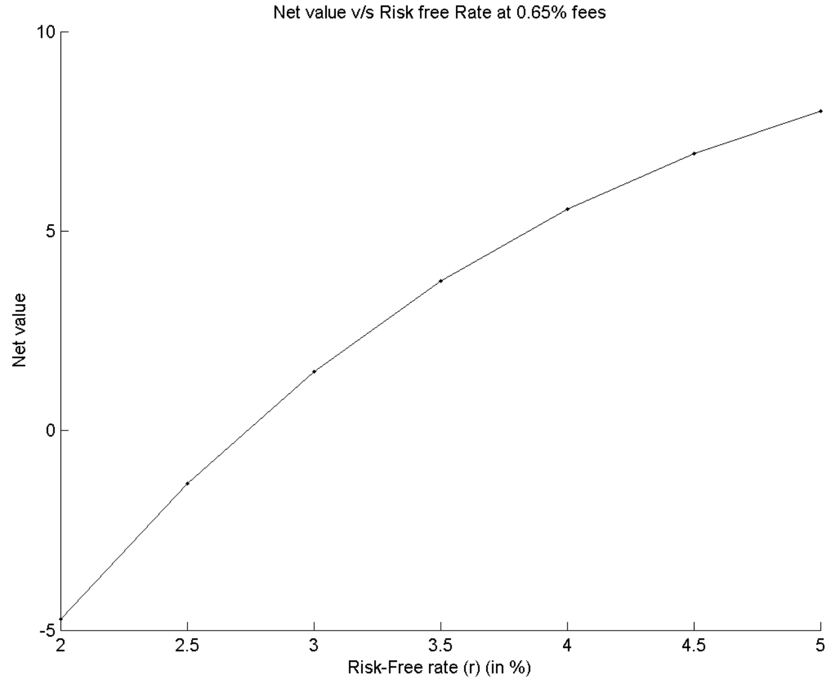

Figure 5-2: Net value of the meta portfolio as a function of risk-free rate for the BSM model.

valuations, henceforth we consider valuations based on contract-compliant withdrawal policy only.

5.4.2 Valuation Spreads Across Asset mixes and Cohort Ages

Figures 5-3 and 5-4 show that across all models, break-even fees and net-value numbers are significantly different for different asset mixes and cohort ages. This confirms the findings of Section 4.5 in Chapter 4 that charging a uniform premium across all cohorts and asset-mixes might lead to a selection bias in the investor pool. In general, investors choosing a more aggressive asset-mix get a “better deal” on the GWB. Also, the optimal age to select the GWB feature is around 62 years. This is the age that has the highest ratio of years during which withdrawals can be made to the overall GWB duration.

5.4.3 Sensitivity to Mortality rates

We next investigate how a shift in cohort mortality rates would impact GWB for life values. For this, GWB for life valuation was done for the cases when the mortality

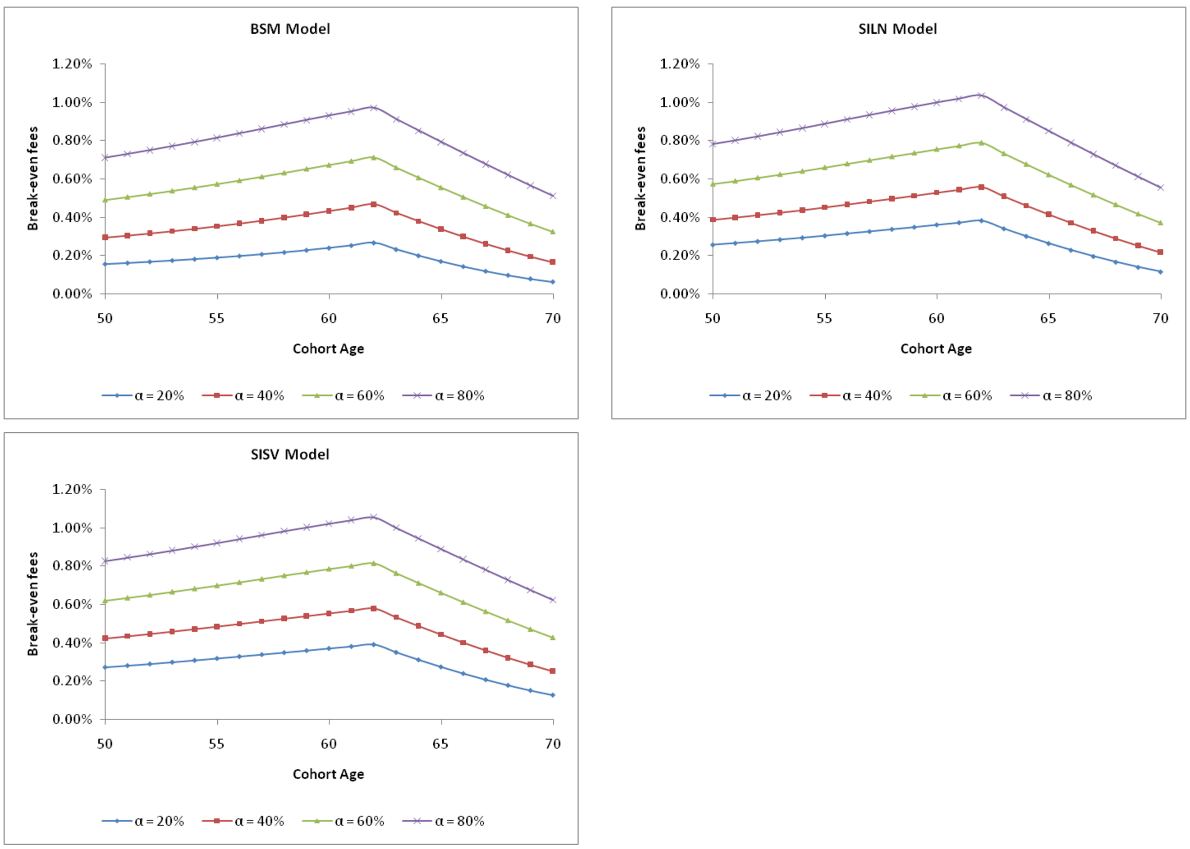

Figure 5-3: Break even fee as a function of Cohort Age for different models and asset mixes assuming contract compliant withdrawals.

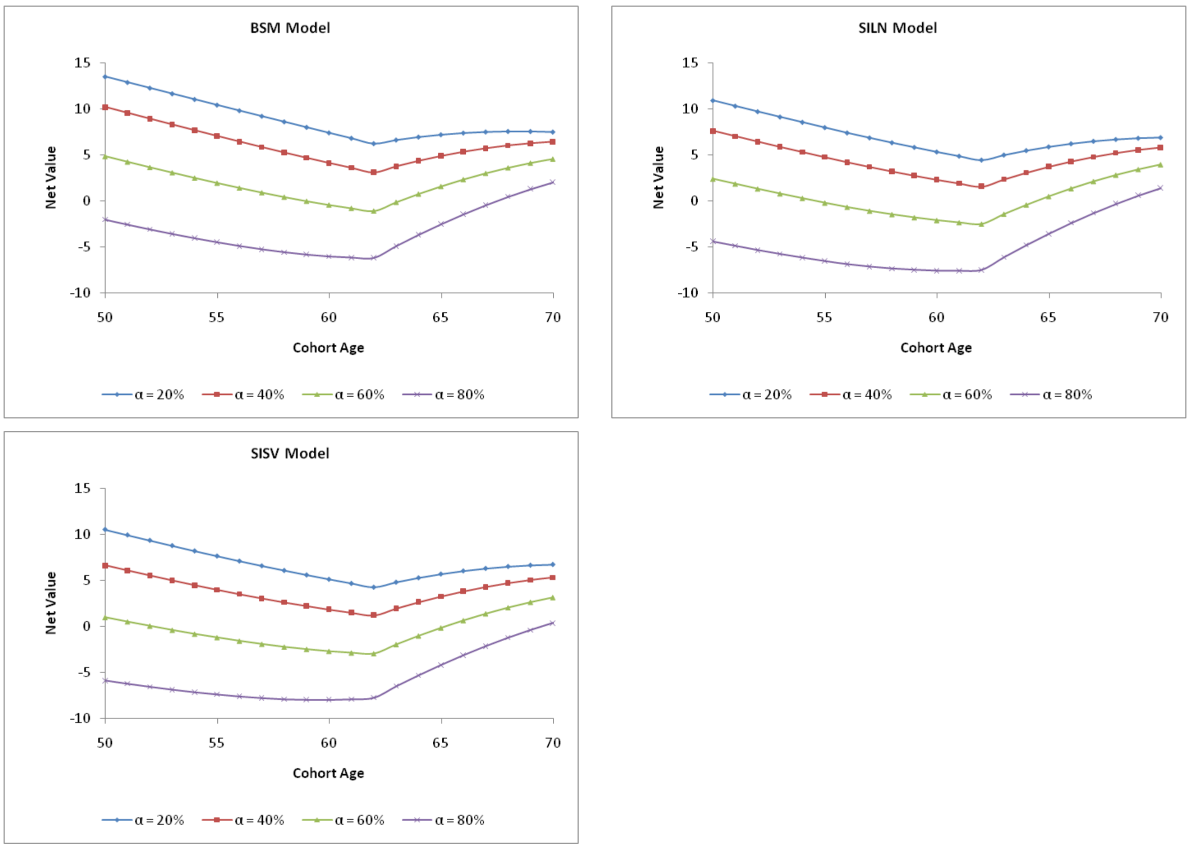

Figure 5-4: Net value as a function of Cohort Age for different models and asset mixes assuming contract compliant withdrawals.

rate curve was shifted in and out by 1 year to obtain the modified mortality rates λ_n^+ and λ_n^- respectively as described earlier in the section. Table 5.5 summarizes the impact on valuation of the meta-portfolio in these scenarios. The numbers again corroborate the analysis in Section 4.5 of Chapter 4. All models indicate a very high sensitivity to assumed mortality rates.

		Mortality rates		
	Model	λ_n	λ_n^+	λ_n^-
	BSM	0.47%	0.51%	0.43%
Break-even fees	SILN	0.55%	0.59%	0.50%
	SISV	0.58%	0.63%	0.54%
	BSM	3.75	2.96	4.47
Net Value	SILN	2.11	1.19	2.96
	SISV	1.40	0.52	2.23

Table 5.5: Sensitivity of the meta portfolio value to mortality rates under different models assuming contract compliant withdrawals.

5.5 Hedging Considerations for the GWB

Since GWB for life constitutes a financial option like guarantee whose risk cannot be diversified away, it must be "hedged" using offsetting derivatives or trading strategies.

The mechanism of delta-hedging derivatives is a well-understood method and follows the same principle as stated in their ground-breaking paper by Black and Scholes [18]. Naive delta-hedging however is unlikely to be effective for GWB for life.

Below we outline how one may hedge the GWB under the BSM model, which is the simplest pricing model. Hedging with more advanced models is similar in principle but will involve more instruments.

- Each pricing model identifies a *market factor of risk*. For the BSM model, the only market factor risk directly considered by the model is the underlying VA fund index S_t. This risk can be hedged against easily by taking an offsetting position in an instrument that has similar characteristics as the VA fund, for example a portfolio of equity index futures and other asset classes that the VA fund invests in. The magnitude of this position is given by the delta or the sensitivity of the GWB value to R_t^s, the instantaneous return in S_t. Delta-hedging provides a first-order insulation and can protect value only for small magnitudes of R_t^s. A higher level of protection can be achieved by gamma hedging, where one takes a position in another instrument so that the second order derivative of the overall portfolio value under the BSM model with respect to R_t^s is zero. Any instrument which has a non-linear dependence on R_t^s, for example an option on a proxy for the VA fund, can be used for the purpose.

- Model parameters are obtained by a calibration process using liquid market instruments that the model should be able to price. In addition to the model implied risk, valuations also face a *Model parameter risk* which arises due to errors in calibration of model parameters or the drift in their values over time as the model is recalibrated. The BSM model has two parameters, return volatilities and interest rates. The model can be calibrated by using, for example, long term options and bonds. Model parameter risk can be mitigated by taking offsetting positions in the calibrating instruments so that the overall portfolio value is stable under small changes in the values of model parameters. If the mis-specification in parameters, is large then this risk cannot be eliminated completely. Also, if the overall portfolio has a high convexity or concavity with respect to model parameters then its value will drift over time in a biased manner. [9]. There is significant risk due to this in the case of GWB, when the BSM model is used. For example, the relationship between GWB values and interest rate, as illustrated in Figure 5-2 is convex.

  In the case of GWB ideally instruments that have exposure to volatilities and interest rates over a long duration should be used for model parameter estimations. A key difficulty with this is that instruments that are sensitive to volatilities, such as options, are liquid only over short time horizons, typically less than 2 years.

- Finally, we have the *model specification risk*, where the model used for pricing does not take into account all possible risk factors or their interactions. This is substantially high for the GWB, where dynamics of the various risk factors such as interest rates, asset returns and their volatilities as well as the correlations between them are difficult to model accurately. This represents an ambiguity in model specification, often termed as model risk and is closely connected to hedging in incomplete markets. A robust towards hedging is required to control model related risk. Theoretical work in this field is still in its infancy, see Cont [42].

[9]This could happen if the assumed model family is not appropriate for the market factor dynamics.

Hedging GWB also requires insurance companies to manage the fluctuations in value due to changes in mortality trends. In general, there are few liquid instruments that are sensitive to mortalities. However, insurance companies also issue life insurance policies or features like GMDB with VAs, which bear an inverse relationship to longevities as compared to the GWB. Risk management can try to balance the issuance of various products so as to reduce the company's overall exposure to population mortalities. They can get even better hedging performance, if these securities with opposing mortality sensitivities can be marketed to the same client.

As analysis in Section 5.4 suggests, the product valuation itself is sensitive to the choice of pricing model. Hedging strategies recommended by these models will also be different. We also observed that the liabilities due to the GWB guarantees can vary depending on how investors take their withdrawals as well as realized mortality rates in the client population.

All these difficulties indicate that there is a significant level of "unhedgeable" risk involved in underwriting the GWB.

From an execution point of view, it is also important to consider the effect of non-linearities that arise due to the step-up feature. To illustrate this point, consider again the example of hedging the GWB under the BSM model described in Section 5.2. Figure 5.5 shows the variation of the delta of GWB value or its sensitivity to the return R^s_T over time on one particular sample path for a cohort aged 60 years at start and with 50% exposure to equities. Note that this delta is the same as the magnitude of the offsetting exposure that the company needs to take in a correlated index to hedge this risk.

We see large discontinuities near contract anniversaries. This is because close to the contract anniversary, the GWB acts like a knock-in put option if a step-up is imminent and has a "barrier" like feature.

Remark 5.3. *To understand this "barrier" nature of the GWB product first consider a simplified scenario, when the GWB has only two more years to go and we are close to the penultimate anniversary. In practice, the insurance company can never know this apriori, but we only seek to illustrate the problems with delta hedging here. We also assume that the contract and benefit values are such that a step-up is almost certain. Note that, after the last but one anniversary the GWB product will be like a (deep-out-of-the-money) put option. The strike of this put option will however be set on the coming anniversary in accordance with the step-up rule. Then, just before the anniversary, if the VA fund appreciates, the strike of the unset put option is pushed higher - this in effect increases the cost of the GWB product. Thus, just before the last*

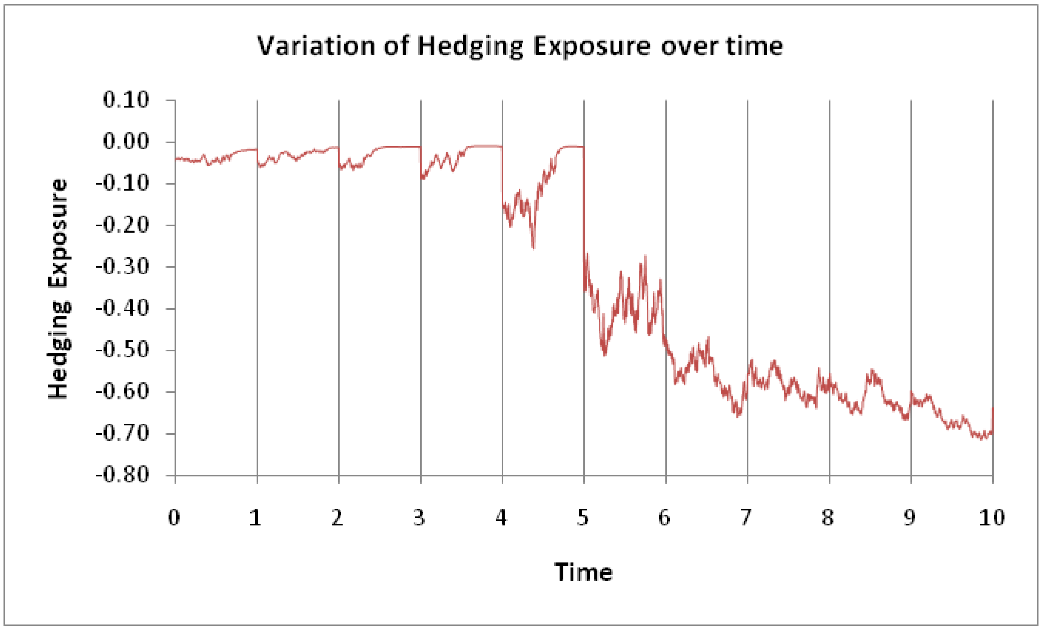

Figure 5-5: Variation of hedging exposure with time for delta hedging under the BSM model.

but one anniversary, the product will have a positive delta with respect to returns on S_t. However, once the anniversary has passed and the step-up registered, the product will behave like an ordinary put option and have a negative delta. As a result, the delta of the GWB is typically discontinuous near the product epochs when a step-up is imminent. GWB, in this respect, is not unlike the exotic barrier options, see Carr [30]. This discontinuity in the "delta" of GWB value poses a challenge for hedging.

The discontinuities at different points in time suggest that hedges are unlikely to be effective at these points. Also readjusting them will need flipping large positions and thus entail high transaction costs.

The insurance companies can mitigate this non-linearities at product-epochs to some extent by spreading out these epochs for different investors over different dates. This will also make the overall product less sensitive to jump risks or risks due to large market movements. Another alternative is to use options with maturities coinciding with GWB epochs as hedging instruments.

On the whole hedging is likely to be at best partially successful in the case of GWB for life.

5.6 Summary and Closing Remarks

In this chapter, we examined the GWB for life product in a realistic setting. While, such products are exciting innovations and can be very useful for individuals and households in planning their retirements, they pose substantial risks for the companies that offer them. We observed that

1. On the whole, a fee of 65 basis points appears enough to cover the cost of the guarantee, though this rests on assumptions about the distribution of investor profiles.

2. There is insufficient price discrimination and GWB for life can be priced at below par value for the riskiest of segments i.e., customers with imminent retirement and choosing an aggressive investment portfolio. The product has to be either re-engineered or priced differently for different age groups and investment styles to avoid the risk of adverse selection.

3. Accounting for interest rate volatility and using models that imply fatter tails for the fund returns can substantially increase the cost of the GWB for the company.

4. Valuation of GWB has a high level of dependence on model choice. All three models considered here - BSM, SILN and SISV can be considered in principle as reasonable models for pricing GWB, especially since there are no comparable securities of such duration traded in the markets. Choosing one model over other can change valuation by more than 2% of sales, or for a typical insurance company selling $10 billion of VAs annually, by about $200 million a year. Note that the actual value of the product will not be realized for decades. This pricing ambiguity poses a difficulty for investors. For regulators too, the magnitudes of financial risk suggested by different models can be substantially different and correspondingly lead to different requirements of ideal capitalization levels. Hence, we believe that it is important for regulatory and accounting bodies to recommend a standardized model for pricing GWB like guarantees.

5. Step-up feature and indexing of fees to the benefit base are effective in making the overall product value less sensitive to investor behavior.

6. Hedging GWB is challenging because of model risk and many assumptions that a valuation must make about investor behavior. Also, there are not enough

liquid instruments that can be used to hedge against all the risk factors that the GWB is exposed to.

A mitigating factor for insurance companies with respect to price differentiation is that individuals who choose less aggressive asset mixes are likely to be more risk averse and hence possibly also more willing to pay a risk-premium for the GWB like guarantees. The same argument may also be used for investors who buy the GWB for life guarantees much before their retirement and effectively pay a higher fee compared to those who buy the same later. However this argument does not apply to investors who enter such scheme at latter years and are seen to pay a fee much higher than the fair amount. Moreover, the insuree's willingness to pay is unlikely to be a determinant of prices in the long run as the VA product space is highly competitive. Also, we find that for the most aggressive investors, the "fair value" of the guarantees is in fact more than the price tag in most cases and this cannot be sustainable.

There are many directions of future research with respect to the GWB to address or quantify the issues that we described above. In particular, robust strategies to hedge GWB, in presence of an ambiguity over value will be an interesting line of research both theoretically as well as practically. Another interesting area of research will be to consider the impact of a business cycle risk, where high interest rates and bull markets are followed by bear markets and low interest rates for the GWB. Product design or restructuring so that it has a more manageable level of risk, while still being useful to the investors is another useful direction of investigation.

Chapter 6

Dynamic Consistency and Dynamic Risk and Asset Management

6.1 Introduction

Many decisions in finance involve solving Sequential Decision Problems (SDPs) under uncertainty and require contingency planning over a multi-period horizon. In such problems, the decision making agent is required to apply controls one at a time and at each step more information that is previously unknown is revealed. SDPs, in fact, also occur in a wide variety of applications outside the field of finance and are an important subject matter of the decision and control theory. Examples include inventory problems, adaptive control for air traffic etc.

A basic conceptual framework used to solve an SDP with uncertainty is the notion of contingent-planning. A contingent-planning framework has the decision maker (DM) deciding not only on the current course of action, but also on a "contingent plan" or a schedule of decisions that will be taken in response to uncertain[1] events or information. In fact, the assertion about the "optimality" of the DM's immediate action rests on the assumption of her carrying through the contingent plan. If it so turns out that as the events unfold, the DM would no longer find the initially planned course of action optimal (or sometimes even feasible), then this leads to a consistency issue. SDPs in which the DM can deviate from a planned course of action are commonly referred to as "dynamically inconsistent" or sometimes "time inconsistent".

[1]But not *unanticipated*. When the uncertainty is revealed, it should not come as a *surprise* to the DM. This means that the sample space of the uncertainty is known to the DM.

This is referred to as an "inconsistency" because although there is uncertainty in the system, there is no real "surprise" or unanticipated information. In that sense, the propensity to deviate is entirely foreseeable by the DM and could have been avoided. The issue is non-trivial because an arbitrary SDP need not be "dynamically consistent". In particular, in finance, as we shall discuss in Sections 6.2 and 6.4 commonly used financial metrics such as variance, Value at Risk (VaR), Conditional Value at Risk (CVaR), when employed in a multi-period setting without due-care, can result in dynamically inconsistent SDPs. For finance applications, as these metrics are often linked to material provisions such as capital or ratings, dynamic consistency has implications beyond normative or philosophical considerations.

6.1.1 Related Work

The problem of dynamic consistency has been studied in a variety of contexts. For example, in preference theory, that traditionally deals with an agent's preferences between lotteries over random outcomes, the property of dynamic consistency has been discussed extensively and a good review of some of the key concepts can be found in Lotito [88]. Strotz [112] gives one of the first examples of dynamic inconsistency. Axiomatic frameworks for dynamic consistency of preferences have been extensively studied (for example, Hammond [64], Weller [118], Cubitt [45], Karni and Schmeidler [79], Volij [114]). These frameworks reveal that dynamic consistency of preferences is closely connected to a Bayesian update of probabilities by the agent and the independence axiom. These properties, with few additional assumptions, imply that agent's preferences have a von-Neumann Morgenstern Expected Utility representation[2]. Machina [89] pointed out that, an implicit assumption in coming to these conclusions is that past events that did not happen do not affect the agent's future preferences. This assumption is known as "Consequentialism". An entirely new axiomatic framework, which showed that non-expected utility preferences can also be dynamically consistent in the presence of ambiguity was proposed by Epstein and Schneider, [53]. In this framework, the agent has multiple priors about the uncertainty in the system and seeks to maximize a robust (worst-case) measure of expected utility. If the priors satisfy a certain property[3], then the agent's preferences in this setting will also be dynamically consistent. Such preferences have typically

[2] See however Johnsen and Donaldson [76] for a different treatment.

[3] The set of priors that the agent considers should be closed under a "pasting" operation of marginal and conditional probability laws at any time step. This property has been variously referred to as stability, consistency, rectangularity etc.

been suggested as a model for ambiguity aversion, see Epstein and Schneider [53], Sarin and Wakker [107], Ozdenoren and Peck [99].

In game theory, the notion of a sub-game perfect equilibrium, also sometimes to referred as "credible threats" is somewhat analogous to dynamic consistency. A sub-game perfect equilibrium implies that players have no incentives to deviate from their respective optimal strategies. See Fudenberg and Tirol [58], and Gibbons [60] for further details.

Finally, in the context of finance, considerable work has recently been devoted to the subject matter of risk-measures. Artzner et al. [4], Follmer and Scheid [56] introduced a framework for respectively "coherent" and convex measures of risk, to characterize sound decision rules for deciding whether an uncertain financial position should be deemed acceptable or not. There has been substantial recent activity in the literature focused that seeks to extend this notion to a multi-period setting. See for example, Reidel, [102], Artzner et al. [5], Roorda and Schumacher [106], Frittelli and Scandolo [57], Wang [116]. Korkmaz [81] provides a survey of the work in this field. Since risk measures are in essence inverted preferences, it is not surprising that there are many parallels between the notions of dynamic consistency in the context of risk measures and the same in the context of preferences. There are also two bi-furcations in the literature in so far as the treatment of risk - measures is concerned. The more common "coherent risk measures", approach as followed by Reidel [102], Artzner et al. [5] define risk measures to be functionals on random variables or processes and then propose criteria for dynamic consistency in terms of these functionals. They show that the dynamically consistent risk measures are equivalent to a robust (worst-case) expectation under a family of priors that satisfy the same rectangularity condition as was stated in the context of preferences under ambiguity by Epstein and Schneider [53]. In contrast, Weber [117] proposes a new paradigm of "distribution-invariant risk measures" where risk measures are treated as functionals on distributions or probability laws. Notice the similarity with the preference theory which is concerned with ranking of "lotteries" or distributions. Weber [117] defines a notion of dynamic consistency in this context and shows that it is, under additional technical restrictions, equivalent to an expected utility based criterion.

6.1.2 Contributions

Although the notions of dynamic consistency in various contexts are analogous, they are not strictly identical. An interesting link that connects all the notions of dynamic consistencies, is that they in some sense legitimize the use of backward induction or the Bellman principle of optimality for dynamic programs, see Bertsekas, [14]. In this chapter, we examine the problem of dynamic consistency in the context of SDPs, i.e., we are interested in dynamic consistency of “strategies” or “plans”. See Sarin and Wakker, [107] and Hazen, [65] for a similar point of view. Moreover, we will motivate our discussion and framework in the context of applications in finance. In this chapter, we will assume, that there is no ambiguity, i.e., the uncertainty in the problem is anticipated and its distributional properties are known to the DM. Our key contributions are:

1. We show that many natural extensions of standard one-period decision problems in finance to a multi-period setting can lead to dynamically inconsistent SDPs. We also highlight how dynamic consistency for risk measures does not necessarily translate to dynamic consistency for the implied strategies, when they are used as strictly an acceptability criteria.

2. Applicability of Bellman's dynamic programming principle or backward induction is sometimes considered to be an equivalent assertion of Dynamic Consistency, see Boda and Filar [19], Sarin and Wakker[107]. We provide examples to illustrate that dynamic consistency is actually a weaker property. Also, even for SDPs that can be solved by backward induction, the objective V need not have a dynamic programming representation. However, we show that in this case, the resultant strategies can also be recovered as solutions to an SDP with a surrogate objective function $\tilde{V}$ that is sum-decomposable across time and mutually exclusive events and thus has a dynamic programming representation.

3. For financial applications, we suggest ways to address some of the dynamic inconsistency related issues. In particular, we propose a new dynamically consistent objective function for risk management. We also point out cases, with concrete examples, where using a pre-commitment solution (or taking a non-consequentialist approach) becomes appropriate.

6.1.3 Chapter Layout

In Section 6.2, we illustrate the issue of dynamic inconsistency in the context of finance problems. In Section 6.3, we describe a formal model for an SDP and explore the connection between the notions of dynamic consistency, backward induction and dynamic programming. We also provide an equivalent representation for dynamically consistent objectives. In Section 6.4, we further discuss how dynamic inconsistency can be mitigated for some SDPs that arise in finance. We conclude with a summary and suggestions for future work in Section 6.5.

We now motivate the issue of dynamic inconsistency in SDPs with specific examples of problems in the finance domain.

6.2 Examples of Dynamically Inconsistent Formulations

We first provide a non-technical definition of the notion of dynamic inconsistency. We say that the formulation is dynamically inconsistent if there exist some chain A of events with the following property. Suppose a DM solves an SDP at time 0 and as a result obtains a plan, which gives her amongst other things a course of action to be followed in the case event A happens, say at time $t > 0$. Now suppose at time $t > 0$, the event A actually happens. The DM updates her SDP to reflect this new information (e.g., by updating probabilities of future outcomes by their values conditional on A.) and resolves the updated SDP and finds that the course of action planned at time 0 to be taken in event A happening is no longer feasible or optimal for the updated SDP.

We will provide a more formal definition in Section 6.3, but for the current discussion and examples in this section, this intuitive notion of dynamic inconsistency should suffice.

Most financial problems involve investment and consumption strategies, subject to certain risk-averseness constraints. A typical static or one-period problem, in general may-be written as

$$\begin{aligned} \max_{y} \quad & f(W, \mathbb{P}) \; ; \\ \text{s.t.} \quad & W \;=\; y' \cdot z \; , \\ & \phi(W, \mathbb{P}) \;\leq\; \alpha \; , \\ & \text{s.t. } y \;\in\; \mathbb{Y} \; . \end{aligned}$$

Here, for example, z could denote returns in a stock (that are random) and $\mathbb{P}$, the distribution of these returns. y the amount that the investor can invest in the stock subject to fixed constraints represented by the set $\mathbb{Y}$. Thus, W is the resulting payoff of the strategy. The investor's goal is to optimize a "performance -metric" functional

$f(\cdot,\cdot)$ on the random variable W given the probability distribution $\mathbb{P}$ on z. The investor must take care that the functional $\phi(W,\mathbb{P})$ which represents a risk-metric does not exceed a pre-specified level α. For example, in the standard risk-reward trade-off proposed by Markowitz [90], we have $f(\cdot,\cdot) \equiv \mathbb{E}[W]$ and $\phi(W,\mathbb{P}) \equiv \text{var}(W)$, where both the mean and variance are taken with respect to the probability law $\mathbb{P}$ on random return z.

Now, a "natural" extension of this one-period problem, in a multi-period setting with T periods, where the DM is only interested in end-of-horizon wealth W_T is

$$\begin{aligned} \max_{\{y_s:0\le s<T\}} \quad & f(W,\mathbb{P})\ ; \\ \text{s.t.} \quad & W \ = \ \sum_{i=0}^{T-1} y_t' \cdot z_{t+1}\ , \\ & \phi(W,\mathbb{P}) \ \le \ \alpha, \\ & \{y_0, y_1, \ldots, y_T\} \ \in \ \mathbb{Y}\ , \\ & y_t \text{ is adapted.} \end{aligned} \tag{6.1}$$

$\mathbb{P}$ now denotes the probability law for the discrete time process $z_1, z_2, \ldots z_T$. Also, the optimal solution y_t is now a policy adapted to the filtration on the z_t process. (6.1) is representative of a typical SDP in finance.

For most common risk metrics, including *dynamic risk measures*, the formulation in (6.1) is dynamically inconsistent.

We illustrate this with the expectation operator under the probability law, a trivial risk measure, but one that is valid as both a coherent dynamic risk measure as well as a distribution invariant dynamic risk measure. For concreteness, we consider the following two-stage problem and suppose there are two possible outcomes for z_1: u and d with equal probabilities $\frac{1}{2}$ at $t = 1$.

$$\begin{aligned} \max_{\{y_o,y_1\}} \quad & \mathbb{E}[\exp(-W_2)] \\ \text{s.t.} \quad & W_2 \ = \ \sum_{i=0}^{1} y_t' \cdot z_{t+1}, \\ & \mathbb{E}[-W] \ \le \ \alpha, \\ & y_t \text{ is adapted.} \end{aligned} \tag{6.2}$$

Suppose, the optimal strategy S is such that the constraint $\mathbb{E}[-W] = \alpha$ is tight with $\mathbb{E}[-W|x_1 = u] = \alpha - \Delta$ and $\mathbb{E}[-W|x_1 = d] = \alpha + \Delta$, $\Delta > 0$. Then if the DM tries to solve (6.1) again, using the Bayes rule updated probabilities at $t = 1$ in the event $z_1 = d$, the she will find this strategy to be actually infeasible. Such inconsistency will arise, in general with any risk measure including variance, VaR, CVaR, other commonly used static risk measures. As, we elaborate in Section 6.3, the reason for this can be traced to a "coupling" of mutually exclusive sub-strategies that the risk-measure constraint causes.

Remark 6.1. *So-called dynamic risk measures can lead to dynamically inconsistent strategies for SDPs because for both coherent and distribution invariant risk-measures, dynamic consistency of risk measures is defined in a sense similar to dynamic consistency of preferences over time (see Reidel[102], Artzner et al [5] and Weber [117] for details). These conditions do not turn out to be strong enough for dynamic consistency in the context of SDP strategies. For example, Reidel [102] defines dynamic consistency for coherent risk measures so as to ensure acceptance and rejection consistency, i.e., if a random variable will be accepted (respectively rejected) at* $t + 1$ *in all possible states of the world, then it should also be acceptable at time* t*. Dynamic consistency requirements for distribution invariant risk measures as defined in [117] are also identical and are equivalent to the following property - if a family of distributions is acceptable(rejectable), then any convex combination of the same is also acceptable(rejectable). This is a version of the "sure thing" principle and ensures consistency only going back in time. For SDP formulations based on dynamic risk measure constraints, to ensure that an optimal strategy is at-least not rendered infeasible, one also needs to guarantee that a strategy that is acceptable at* t*, will also be "acceptable" at* $t + 1$*, in* all *states of the world.*

The inconsistency issues in the problem (6.1), might appear obvious, but nevertheless are important and often overlooked. For example, Basak and Shapiro [12] solve a version of the SDP (6.1), using VaR as the risk metric and draw conclusions based on the optimal strategy, without recognising the essential dynamic inconsistency of the original formulation. See Cuoco He and Issaenko [46] for further details. An implicit assumption that can justify solving such an SDP, is that the DM would pre-commit to the initially deviced strategy. As we discussion in Section 6.4, this assumption is justifiable only in specific contexts. In practice, as stressed in Cuoco, He and Issaenko [46] VaR and CVaR metrics are meant and used by institutions to manage risk in a static context. Nonetheless, it will be useful to have corresponding risk metrics in a dynamic setting.

A possible, resolution to the dynamic inconsistency is to avoid any risk-metric based constraints, and instead consider an alternate formulation where we add a penalty term involving the risk measure to the performance metric objective[4], i.e., we use a risk-adjusted performance metric as our objective functional.

$$
\begin{aligned}
\max_{\{y_s:1\le s\le T\}} \quad & f(W,\mathbb{P}) - \lambda\phi(W,\mathbb{P})\alpha \ ; \\
\text{s.t.} \quad & W \ = \ \sum_{i=0}^{T} y_t' \cdot z_{t+1} \ , \\
& \{y_1, y_2, \ldots, y_T\} \ \in \ \mathbb{Y} \ , \\
& y_t \text{ is adapted} \ .
\end{aligned} \tag{6.3}
$$

If one holds the level of penalty term constant over time, it can be verified that the version in (6.3) is dynamically consistent and the "Lagrangian" transformation does remove the inconsistency arising in the specific case of (6.2). Unfortunately, this approach does not work, if the performance metric is derived from any of the commonly used non-trivial risk-metrics such as variance, VaR and CVaR.

Consider, for example, the Lagrangian version of the multi-period mean-variance problem. Markowitz [90] proposed the one period mean-variance framework for portfolio selection in 1950. It provides an intuitive and tractable way for risk management. The multi-period version is easily formulated as

$$
\begin{aligned}
\max \quad & \mathbb{E}[W_T] - \frac{\lambda}{2}\text{var}[W_T] \ ; \\
\text{s.t.} \quad & W_T \ = \ \sum_{t=1}^{T} y_t' \cdot x_{t+1} \ .
\end{aligned} \tag{6.4}
$$

(6.4) can also be considered as a quadratic utility maximization problem.

This formulation is in general dynamically inconsistent. We illustrate this through a very simple two-period example.

Example 6.1. Dynamic inconsistency from Variance based objectives

Consider the following simple 2 *period mean- variance optimization problem. Suppose the investment opportunity set consists* 2 *assets and there are no short-selling constraints. Asset* 1 *is riskless and offers no returns. Asset* 2 *offers a return* r_1 *in period 1 and* r_2 *in period 2. We assume* z_1 *and* z_2 *are independent and are either* u *or* d*, each with equal probability* ($\frac{1}{2}$)*. Suppose* y_0^* *denotes the optimal time* 0 *investment in the risky asset while* $y^*(z_1)$ *denotes the same for time* 2*. Note that* $y^*(z_1)$ *can*

[4] This can also be viewed as a Lagrangian transformation of the original problem.

depend on the new information available at time 1, that is the risky asset return for the period. It should follow that period t wealth W_t is given by

$$\begin{aligned} W_1 &= W_0 + y_0^* \cdot z_1 \ ; \\ W_2 &= W_1 + y_1^*(z_1) \cdot z_2 \\ &= W_0 + y_0^* \cdot z_1 + y_1^*(z_1) \cdot z_2 \ . \end{aligned}$$

Thus, our problem is to

$$\max_{y_0^*, y_1^*} \mathbb{E}[W_2] - \frac{\lambda}{2}\mathrm{var}(W_2) \ \equiv \ \max_{y_0^*} \left(\max_{y_1^*} \mathbb{E}[W_2] - \frac{\lambda}{2}\mathrm{var}(W_2) \right) \ .$$

We note that

$$W_2 \ = \ W_0 + y_0^* \cdot (z_1) + y_1^*(u) \cdot 1_{\{z_1 = u\}} \cdot (y_2) + y_1^*(d) \cdot 1_{\{z_1 = d\}} \cdot (z_2) \ .$$

It can be verified that the optimal solution to the above problem is

$$\begin{aligned} y_0^* &= \frac{4(u^2+d^2)(u+d)}{\lambda(u-d)^4} \\ y_1^*(u) &= -\frac{4d(u+d)}{\lambda(u-d)^3} && (6.5) \\ y_1^*(d) &= \frac{4u(u+d)}{\lambda(u-d)^3} && (6.6) \end{aligned}$$

Note that, for the SDP to be consistent then at time step 1, we must choose control given by (6.5) if the previous period risky return was u and that by (6.6) otherwise.

Let us now solve directly for what optimal investment policy should be at time 1. Suppose W_1 is the wealth at time 1. Then,

$$W_2 \ = \ W_1 + y_1^* z_2 \ .$$

It can be easily verified that the objective $\mathbb{E}[W_2] - \frac{\lambda}{2}\mathrm{var}(W_2)$ is optimized by setting

$$y_1^* \ = \ \frac{2(u+d)}{\lambda(u-d)^2} \ . \qquad (6.7)$$

and is independent of the wealth W_1. The value given by (6.7) is different from that implied by (6.5) or (6.6). Thus, multi-period mean-variance optimization is dynamically inconsistent.

Unlike (6.2), dynamic inconsistency in the SDP in (6.4) occurs because the objective V_t itself becomes coupling due to the presence of the variance term. Note,

$$\begin{aligned} V_t &\triangleq \mathbb{E}[W_T|\mathcal{F}_t] - \frac{\lambda}{2}\text{var}(W_T|\mathcal{F}_t) \\ \text{Hence, } V_{t-1} &= \mathbb{E}[V_t|\mathcal{F}_{t-1}] - \frac{\lambda}{2}\text{var}(\mathbb{E}[W_T|\mathcal{F}_t]|\mathcal{F}_{t-1}) \end{aligned}$$

where we have used the standard notation $\mathcal{F}_t$, to denote information known at time t. The component $\text{var}(\mathbb{E}[W_T|\mathcal{F}_t]|\mathcal{F}_{t-1})$ couples together sub-strategies corresponding to mutually exclusive events, and as we discuss in Section 6.3 makes the problem dynamically inconsistent. It is interesting to note that versions of the basic multi-period mean variance problem in (6.4) have been widely considered and/or solved - for example, Chen, Jen and Zionts, [33], Li and Ng, [87], Leippold, Trojani and Vanini, [85], Bajeux-Besnainou and Portait, [6], Basak and Chabakauri, [11]. Only the recent paper by Basak and Chabakauri, [11], remarks on the dynamic inconsistency issue in the underlying problem and seeks to address it by recursively defining the mean-variance based objective function. However, the reformulated problem is not truly dynamically consistent, but rather forced to become so by an imposition of the requirement that optimal policies be obtained by backward induction[5].

Li and Ng [87] and Leippold, Trojani and Vanini [85] actually solve a surrogate problem, that is equivalent to the actual problem at time $t = 0$. The objective function

$$\mathbb{E}[W_T] - \frac{\lambda'}{2}\mathbb{E}[W_T^2]$$

in the surrogate problem is dynamically consistent. The assumption implicit in the solution so obtained is that of "Resolute Choice" or pre-commitment, also discussed in Section 6.4.

Apart from normative issues, a dynamically inconsistent SDP also poses computational challenges. This is because the backward induction method, a natural "divide and conquer" style algorithm for SDPs, cannot be used to solve a dynamically inconsistent SDPs. As a result even when the DM is willing to "pre-commit", the optimal strategy cannot be derived using backward induction.

We illustrate this link using another commonly used risk measure, the CVaR, as an objective. CVaR (also known as Tail Conditional Expectation, TCE) is a coherent risk measure and satisfies certain desirable properties for a static or one-period risk

[5] This is equivalent to the assumption of 'Sophisticated Choice', that we elaborate in Section 6.4.

metric, see Brown, [25]. CVaR is usually used in risk management in a static sense by enforcing a constraint of the form that CVaR at a given (say 5 %) level of a portfolio under management should stay above a recommended level. In a dynamic setting, using a constraint that involves CVaR in an SDP will lead to the same issues as we observed for (6.2). CVaR, however, cannot be incorporated as a penalty term in the SDP objective either without making it dynamaically inconsistent. We again illustrate the problems with using CVaR with a simple 2- period setting.

Example 6.2. Dynamic inconsitency from Variance based objectives *Consider a two period SDP with two assets one risk-free and the other risky. The agent is free to borrow or lend and the interest rate is zero. Returns on the risky asset in the two periods, denoted by* z_1 *and* z_2*, are independent and identically distributed with the following density function* $f(z)$ *:*

$$f(z) \;=\; \begin{cases} 4 & \ldots -\frac{1}{8} \le z \le 0 \ , \\ \frac{1}{4} & \ldots 0 < z \le 2 \end{cases} \quad .$$

The DM's objective is to choose y_0 *and* y_1*, amounts to be invested in the risky security in periods one and two so as to maximize* $\mathrm{CVaR}_{0.5}(W_2)$*. Her SDP is given by:*

$$\begin{aligned} \max_{y_0, y_1} \quad & \mathrm{CVaR}_{0.5}(W_2) \ ; \\ s.t. \quad & W_2 = y_0 \cdot z_1 + y_1 \cdot z_2 \ , \\ & |y_0|, |y_1| \ \le \ 2 \ . \end{aligned}$$

To make the problem bounded, we have imposed a limit on the size of the position in the risky asset that the DM can take in either period. Recall that $\mathrm{CVaR}_\alpha(W)$ *is defined as* $\mathbb{E}[W | W < F_W^{-1}(\alpha)]$*, where* $F_W(\cdot)$ *denotes the cumulative density function or CDF of* W*. In our example thus, the* $\mathrm{CVaR}_{0.5}$ *of one period return, for an investment* Y *is* $-\frac{Y}{8}$*, if* $0 < Y \le 2$ *and* $-|Y|$ *if* $-2 \le Y \le 0$*. Thus at time* $t = 1$*, as* $\mathrm{CVaR}_{0.5}$ *is linear, the optimal decision is to not invest irrespective of the outcome* z_1*. By backward induction then, the optimal decision at time* 0 *should also be not to invest at all in the risky security, Thus the optimal value of* $\mathrm{CVaR}_{0.5}$*, obtained by backward induction is* 0*. However, it is easy to construct a contingent plan that beats this* *objective. Consider the following strategy:*

$$\begin{aligned} y_0 &= 2 \\ y_1 &= \begin{cases} 0 & \ldots \quad z_1 \geq 0 \; , \\ 1 & \ldots \quad z_1 < 0 \; , \end{cases} \end{aligned}$$

Thus,

$$W_2 \;=\; 2 \cdot z_1 + \mathbf{1}_{\{z_1<0\}} \cdot z_2 \; .$$

For this strategy, after somewhat tedious calculations, it can be verified that $\mathbb{P}(W_2 \leq \frac{15}{16}) = \frac{1}{2}$ *and* $\mathrm{CVaR}_{0.5}(W_2) > 0$*. Thus backward induction fails to produce the optimal strategy for SDPs involving CVaR.*

Using similar arguments, it can be shown that Value at Risk (VaR), another risk-metric, which is more commonly used than CVaR, also suffers from the same problem in a dynamic setting. A somewhat complex example to illustrate this is given by Boda and Filar [19]. While, VaR and CVaR are mostly used by financial institutions to report short-term or myopic (typically daily) risk exposures and maybe appropriate in this context, it is clear that a long term risk-management strategy targeting VaR, CVaR or variance is fraught with fundamental conceptual difficulties.

To recap, so-called dynamic risk measures cannot be used as an acceptability criterion for strategies in an SDP context. Also, if the objective function comprises common risk-metrics such as variance, VaR and CVaR, the resultant SDP again becomes dynamically inconsistent. Our goal for the remainder of this chapter is to characterize objective functions that will lead to a consistent SDP. For this, we describe a formal framework for SDPs in the next section.

6.3 Conditions for Dynamic Consistency

Over the horizon of an SDP, the DM actually solves a series of sub-problems. If an SDP is dynamically inconsistent, then this inconsistency maybe traced to either

- An explicit change in the problem itself OR
- An implied change, through the update of probability laws.

We mainly focus on the case, where there is no explicit change to the problem, as this can also be interpreted as a change of taste. We find that when explicit change is ruled out, the property of dynamic consistency is strongly connected but not strictly equivalent to the SDP being a dynamic program. We first establish a formal framework to describe SDPs.

6.3.1 Framework

We consider a discrete time framework with periods $1, 2, \ldots, T$. Decisions are made at times $0, 1, \ldots, T$. Decision at time t is made *after* observing the information released at time t. Mathematically, the revelation of information is modeled by a random process adapted to a filtration $\mathcal{F}_t$. Although information is new, it is not "unanticipated", i.e., the sample-space or the set of all possible outcomes for the information process is known to the DM. Also, the DM has a probability distribution on the set of all possible outcomes. We now elaborate on each component of the formal framework.

- **Randomness**
  - Randomness or Information in the system is denoted by the discrete process z_t, with $1 \leq t \leq T$.
  - We use $Z_{s:t}$ to denote the sub-sequence $(z_s, z_{s+1}, \ldots z_t)$, with $Z = Z_{1:T}$.
  - $\mathcal{F}_t$ and Ω_t denote the filtrations and sample spaces associated with $Z_{1:t}$ respectively.
  - For simplicity, it is assumed that $\Omega \triangleq \Omega_T$ is finite and $|\Omega| = N$. Also, we enforce an order among outcomes to simplify notation. Thus Ω is an ordered set $\{Z^1, Z^2, \ldots, Z^N\}$. For any subset A of Ω, the elements in A are considered as per their relative order in Ω.
- **Probability law**
  - The probability law considered by the agent at time t is denoted by $\mathbb{P}_t$. Since Ω is assumed to finite, $\mathbb{P}_t$ is specified if $p_t^i \triangleq \mathbb{P}_t(Z = Z^i)$ is specified $\forall i : 1 \leq i \leq N$.
  - If A is any event then the conditional probability law, which we denote by ${}_A\mathbb{P}$ is defined in the standard way:

$$
{}_Ap^i = \begin{cases} \frac{p^i}{\mathbb{P}(A)} & , \quad Z^i \in A \\ 0 & , \quad Z^i \notin A \end{cases}
$$

Probability Update rules: Probability laws themselves evolve over time as they are updated as new information is revealed. However, an update rule that will yield $\mathbb{P}_t$, given $\mathbb{P}_{t-1}$ and z_t is specified a-priori. The most common way to do this is using Bayesian updates. This means

$$\begin{aligned} \mathbb{P}_{t+1}(Z^i) &= \frac{\mathbb{P}_t(Z^i, Z_{1:t})}{\sum_{j=1}^{N} \mathbb{P}_t(Z^j, Z_{1:t})} \\ &= \frac{\mathbb{P}_{t-1}(Z^i, Z_{1:t})}{\sum_{j=1}^{N} \mathbb{P}_{t-1}(Z^j, Z_{1:t})} \\ &= \frac{\mathbb{P}_0(Z^i, Z_{1:t})}{\sum_{j=1}^{N} \mathbb{P}_0(Z^j, Z_{1:t})} \ . \end{aligned}$$

Thus,

$$\mathbb{P}_t \equiv {}_{Z_{1:t}}\mathbb{P}_0 \tag{6.8}$$

Remark 6.2. *Bayesian updates implicitly assume the principle of Consequentialism (see Machina [89], Hammond [64]), which basically states that the outcomes that cannot occur should not influence decisions or preferences. In our setting this is equivalent to assigning zero probabilities to the events ruled out by z_t. An alternative is the non-consequentialist approach, that Machina [89] argues can be considered as appropriate in some settings. Under the non-consequentialist approach, probability laws are effectively never updated. As dynamic consistency follows trivially in this case, we consider only the Bayesian update rule for probabilities.*

- **Controls**
  - At time t a control y_t is applied after observing z_t.
  - y_t is chosen from (a possibly infinite) domain $\mathbb{Y}_t$ that is non-stochastic.
  - $Y_{s:t}$ denotes the sequence $(y_s, y_{s+1}, \ldots, y_t)$ with $Y \triangleq Y_{0:T}$.
  - y_t must be $\mathcal{F}_t$ measureable. For notational ease, we will define Y^i as the complete control sequence applied corresponding to the sample path Z^i.
- **Strategy**
  - A strategy S is a sequence of mappings from Ω_t to $\mathbb{Y}_t$ for $t = 0, 1, \ldots, T$. Thus, a strategy gives a recipe for applying controls adapted to $\mathcal{F}_t$.

- A strategy can be partitioned along both time and event space. Given indices j, k and an event $A \in \mathcal{F}_j$, a sub-strategy ${}_A S_{j:k}$ is a sequence of mappings from $\Omega_t \cap A$ to $\mathbb{Y}_t$ for $t = j, j+1, \ldots, k$ that agrees with S on its domain.
- Given individual strategies S^1, S^2, S^3, and an event $A \in \mathcal{F}_t$ and its complement A^c, one can construct a compound strategy $\{S^1_{0:t-1}, {}_A S^2_{t:T} \cup {}_{A^c} S^3_{t:T}\}$, in the obvious way.
- $S(Z_{1:t})$ denotes the entire control sequence response to $Z_{1:t}$, i.e., the sequence of controls $(S(\cdot), S(Z_{1:1}), \ldots, S(Z_{1:t}))$.
- $S_{i:j|A}$, will be used to denote a partial strategy, which is essentially a sequence of mappings from $\Omega_t \cap A$ to $\mathbb{Y}_t$ for $t = i, i+1 \ldots, j$. Again, the default value for j will be T and A must be $\mathcal{F}_i$ measureable.

- **Objective function:** At any time t, the agents objective function is given by

$$V(Z^1, Y^1, p^1_t, Z^2, Y^2, p^2_t, \ldots, Z^N, Y^N, p^N_t | Z_{1:t}, Y_{0:t-1}) \tag{6.9}$$

or more compactly denoted as

$$V(\Omega, S, \mathbb{P}_t | Z_{1:t}, Y_{0:t-1})$$

The "conditioning" notation in (6.9) and (6.10) is used as a short-hand to denote the constraint that $S_{0:t-1}$ is no longer under control and is pinned to $Y_{0:t-1}$ and that the law $\mathbb{P}_t$ has been updated to reflect the occurrence of $Z_{1:t}$. Also,

  - We constrain the Objective function V to remain the same in nature over time. This avoids the agent from having time-varying tastes, which can trivially lead to dynamic inconsistency.
  - Without loss of generality, we require that outcomes Z^i and corresponding decisions Y^i, that have been assigned $p^i_t = 0$ probability, cannot influence $V(\cdots)$ and may be dropped without any consequence.

Remark 6.3. *The model presented above has the following restrictions:*

- *The domain of allowed control variates does not "evolve" or change stochastically and is assumed to be fixed. (i.e., there are no stochastically changing constraints). A general SDP may include a constraint of the form* $f(Z, Y, \mathbb{P}_t) < b$.

We do not consider this case here as we seek to characterize structural properties of V for dynamic consistency. Also, as we saw in examples in Section 6.2, such constraints typically lead to dynamic inconsistency. Our-set up is actually general enough to include "deterministic" constraints of the nature $f(Z, Y) < b$, provided they are non-coupling or "rectangular" in nature. By this we mean that the same constraint cannot involve two controls that have a zero probability of being applied together. We use the more restricted setting to retain focus on the key properties.

- *The control variables y_t cannot influence the evolution of uncertainty Z in our model. This condition can be restrictive in some cases. In Dynamic Choice and Preference theory for example, the agent's decision problem is posed as that of a choice between two distributions of outcomes. The proposed framework, in general will be unsuitable to model problems of this nature*[6].

- *The reference to an explicit information process Z may appear superfluous as one can redefine the SDP with probability laws defined on a standardized sample space. However, we prefer this formulation involving Z to highlight that it has a physical significance in our setting and is used to refer to the finest measureable sample-space. Also, it serves to easily identify information known at time t.*

We now provide a formal definition of dynamic consistency in our setting.

Definition 6.1. *We say a solution strategy S^* is an* optimal dynamically consistent *policy if*

- *It optimizes $V(\Omega, S, \mathbb{P}_0)$.*

- *and remains optimal at all t, no matter what the outcome so far i.e., $Z_{1:t}$ has been and thus also maximizes $V(\Omega, S, \mathbb{P}_t | Z_{1:t}, S^*(Z_{1:t-1}))$ for all t and $Z_{1:t}$.*

We say that the SDP formulation is dynamically consistent *if all strategies that are optimal at time* 0 *are also dynamically consistent as defined above for any specification of initial probability law $\mathbb{P}_0$.*

This definition of dynamic consistency formalizes the notion that once the DM has solved her problem at time 0 and found an optimal strategy, she will not have

[6]Such a decision between lotteries maybe modeled by using an elevated uncertainty dimension. But, in order to ensure equivalence, certain additional linearity-like properties must be imposed on the pseudo-probabilities. See for example the Simple Reduction axiom in Lotito, [88].

any incentive to deviate from the same. This requirement is in a similar spirit as the notion of a sub-game perfect equilibrium in game theory. The following thought experiment makes the analogy clearer:

- Suppose, different DMs are in charge of exercising controls at different stages of the SDP. Figure 6-1 illustrates this schematic for a two period problem, where there are two possible random outcomes, u and d at each stage.

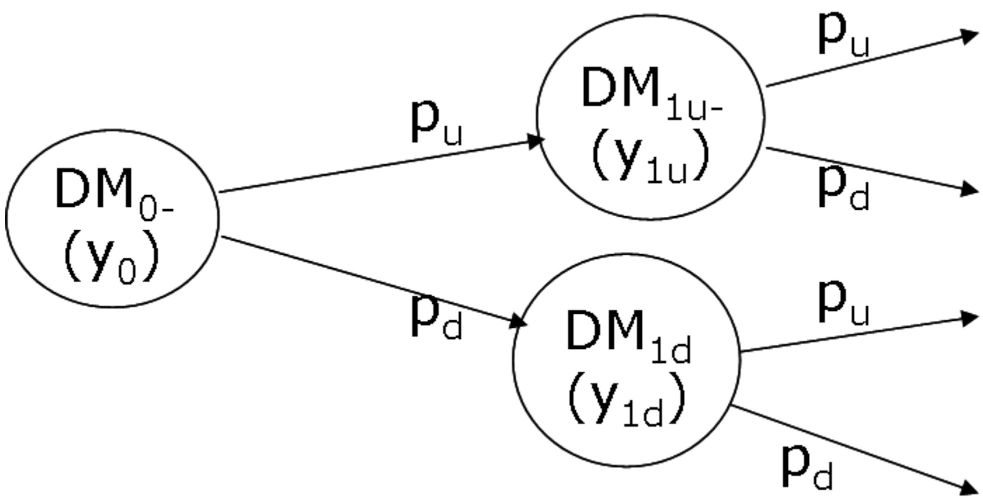

Figure 6-1: Illustration of Co-ordination between DMs for Dynamic Consistency

- We refer to the DM which exercises the control at time 0 as the principal DM or DM_0. DM_0 devices the optimal strategy or plan for the SDP. She then exercises the control y_0 to be applied at time 0 and instructs the DMs in charge of implementing the subsequent controls (DM_{1u} and DM_{1d}, in the example of Figure 6-1.), trusting them to follow her instructions.

- When new information is revealed, the DM that becomes active (DM_{1u} if $Z_1 = u$ and DM_{1d} if $Z_1 = d$ in the example) solves her own version of the SDP that is derived by appropriately updating the principal agent's SDP. Dynamic consistency then means that this DM will not have any incentive to deviate from the principal agent's instructions.

For dynamic consistency, then the objectives of different DMs must be aligned. Note that, in this version of dynamic consistency, the principal DM, DM_0 is allowed to communicate "instructions" to the future DMs and co-ordinate their actions. We can also define the following stronger notion of dynamic consistency, which does not allow for such co-ordination.

Definition 6.2. *The SDP formulation is "strongly dynamically consistent" if the following property holds: Let S^* be an optimal strategy for the problem at time* 0*. Let $S'_{t:T|Z_{1:t}}$, be an arbitrary optimal partial-strategy maximizing $V(\Omega, S, \mathbb{P}_t|Z_{1:t}, S^*(Z_{1:t-1}))$. Then the compound strategy $\{S^*_{0:t-1}, {}_{Z^c_{1:t}} S^*_{t:T} \cup S'_{t:T|Z_{1:t}}\}$ is also optimal for the problem at time* 0*. Morevoer this holds for any values of $t > 0$ and $Z_{1:t}$.*

Strong dynamic consistency implies (ordinary or regular) dynamic consistency. In the thought experiment, if the SDP is *strongly* dynamically consistent, then even when the principal DM, i.e., DM_0 cannot leave any instructions to the future DMs about what controls they should apply, these DMs would choose the right controls on their own accord, obtained by solving their version of the SDP. The following two-period example illustrates that one can have SDPs that are dynamically consistent without being strongly dynamically consistent.

Example 6.3. *Suppose DM at t for $t = 0, 1$ solves the SDP,*

$$V_t \quad = \quad -\mathbb{E}_t[(|Y_1| - Z_1)^2] + \text{var}_t(Y_1).$$

Where, the sample space for Z_1 is $\{1, -1\}$ and $\mathbb{Y}_0, \mathbb{Y}_1 = [-1, 1]$. This formulation, while dynamically consistent is not strongly dynamically consistent[7]*.*

Remark 6.4. *For a strongly dynamically consistent formulation, at any stage, it is sufficient to find* any *optimal sub-strategy and roll it back to construct an optimal dynamically consistent strategy. Thus backward induction works in a straightforward manner. For a formulation that is dynamically consistent but not strongly dynamically consistent, backward induction can be used to solve for the optimal strategy, provided one finds and considers* all *the optimal sub-strategies discovered during the process of rolling back. The distinction between the two becomes moot if uniqueness of optimal strategies can be guaranteed at all stages, i.e., there is no degeneracy.*

6.3.2 Examples of Dynamically Consistent Formulations

Before, formally proceeding with an investigation of properties for dynamic consistency, we examine the type of SDPs that are known to be actually dynamic consistent. Two commonly used formulations that lead to dynamically consistent plans are:

[7]Note that in this example, any optimal partial strategy at $t = 1$ is part of some optimal strategy. However not all combinations of optimal partial strategies are optimal.

- *Expected Utility formulation:* In the framework described here, the DM's objective function takes the form:

$$V(Z^1, Y^1, p_t^1, Z^2, Y^2, p_t^2, \ldots, Z^N, Y^N, p_t^N) = \sum_{i=1}^{N} p_t^i U(Z^i, Y^i) \quad (6.10)$$

  for some function $U(\cdot)$. Note that the Bayesian update of probabilities is crucial for the Expected Utility formulation to be consistent. Weller [118] formally shows this in the context of preference theory.

- *Worst Case/ Best Case formulation:* Here, the agent's objective function takes the form:

$$V(Z^1, Y^1, p_t^1, Z^2, Y^2, p_t^2, \ldots, Z^N, Y^N, p_t^N) = \min_{i:p_t^i>0} U(Z^i, Y^i) \; ; \quad (6.11)$$

$$V(Z^1, Y^1, p_t^1, Z^2, Y^2, p_t^2, \ldots, Z^N, Y^N, p_t^N) = \max_{i:p_t^i>0} U(Z^i, Y^i) \quad (6.12)$$

  These formulations will be in fact dynamically consistent for any probability update rule that continues to assign non-zero probabilities to events that are not ruled out by z_t.

  Interestingly, while worst-case (or best-case) objective based measures lead to dynamically consistent formulations, objectives which seek to maximize the 2^{nd} worst-case (or 2^{nd} best-case) do not. These formulations suffer from the same issues as problems with VaR based metrics in dynamic settings.

  In some sense, the worst-case or best-case formulations are not too different from the expected utility formulations and can be viewed as limiting cases of the same. For example, the worst-case formulation in (6.11) will lead the same actions as an SDP with the objective

$$V(Z^1, Y^1, p_t^1, Z^2, Y^2, p_t^2, \ldots, Z^N, Y^N, p_t^N) = \lim_{m \to -\infty} \left(\sum_{i=1}^{N} p_t^i (U(Z^i, Y^i))^m \right)^{\frac{1}{m}}$$

  which is similar to the formulation in (6.10).

6.3.3 Structural Properties of the Objective for Dynamic Consistency

As probability laws are updated according to the Bayes rule, it is the nature of the objective function V that will determine if the SDP is dynamically consistent. We find that the key property for V is **separability** which is defined as follows:

Definition 6.3. Strong Separability: *The SDP is strongly separable if the following condition if it satisfies the following condition: Fix $Z_{1:t}$ and $Y_{0:t-1}$. Then the optimizing partial sub-strategy $S^*_{t|Z_{1:t},Y_{0:t-1}}$ is independent of $p^i, Z^i, i \notin \{Z_{1:t}\}$.*

An SDP satisfying the strong separability condition can be solved by backward induction and hence will be strongly dynamically consistent. Thus strong separability is a sufficient condition for strong dynamic consistency. However, as the following example illustrates, it is not a necessary condition.

Example 6.4. *Suppose the DM seeks to optimize the following objective at time* 0.

$$V = -\mathbb{E}[(Y_1 - Z_1)^2] + |Y_0| \cdot \text{var}(Y_1) \ .$$

where, the sample space for Z_1 is $\{1, -1\}$ and $\mathbb{Y}_0, \mathbb{Y}_1 = [-1, 1]$ This formulation is not strongly separable as for $y_0 < 0$ the optimal control y_1 that should be applied in the event $Z_1 = 1$ is not independent of the control that should be applied when $Z_1 = -1$. The SDP however is strongly dynamically consistent.

Strong separability is a stronger condition than strong dynamic consistency as the latter will hold under the following weaker condition:
Suppose the partial strategy $S^\dagger_{0:t-1}$ is known to be optimal. Then optimizing sub-strategy $S^*_{t|Z_{1:t},S^\dagger_{0:t-1}(Z_{1:t})}$ is independent of $p^i, Z^i, i \notin \{Z_{1:t}\}$.
Thus for dynamic consistency, independence between optimizing sub-strategies corresponding to two mutually exclusive events need only hold for optimal control paths.

This makes a characterization of dynamically consistent SDPs difficult, as the above condition is difficult to verify in practice. However, as we now show, a converse characterization which shows that all strongly dynamically consistent SDPs are in some sense equivalent to SDPs that have strongly separable objective functions is possible. In fact, any dynamically consistent strategy can be thought of as arising from an SDP that can be solved by a dynamic program.

6.3.4 Dynamic Consistency and Decomposability

Generalization of the exact structure of the objective functions that lead to Dynamic consistency is difficult[8]. The following theorem however shows that nevertheless there exists a canonical representation for objective functions that lead to strongly dynamically consistent SDPs under Bayesian updates.

Theorem 6.1. *Let an SDP with the objective $V(\cdots)$ lead to strongly dynamically consistent SDPs under Bayes rules for probability updates. Then, the strategies generated by $V(\cdots)$ are identical to those generated by optimizing an SDP with the objective function $\tilde{V}$ which has the specific canonical form described below:*
Let $|\Omega_t| = N_t$ and $Z^i_{1:t}, 1 \leq i \leq N_t$ denote the singleton elements of Ω_t. Also define $\Omega^i_t \triangleq \Omega \cap \{Z^i_{1:t}\}$. Then,

$$V(\Omega, S, \mathbb{P}_0) \equiv \tilde{V}(\Omega, S, \mathbb{P}_0) = \sum_{t=0}^{T} \sum_{i=1}^{N_t} V_t(Z^i_{1:t}, Y^i_{0:t}, \Omega^i_t, {}_{\Omega^i_t}\mathbb{P}_0) \ . \tag{6.13}$$

Moreover,

$$\max_{S_{1:T}} \tilde{V}(\Omega, \{y_0, S_{1:T}\}) = \max_{S_{1:T}} V(\Omega, \{y_0, S_{1:T}\}) \ .$$

and hence the two SDPs attain the same optimal value.

Proof. Before we proceed with the proof, note that the above decomposable representation follows the familiar dynamic programming setting. $X_t = \{Z_{1:t}, Y_{0:t-1}\}$ can be interpreted as the "state" (or history) of the system at time t and $U_t = \{Z_{1:t} \cap \Omega, \ {}_{Z_{1:t}}\mathbb{P}_0\}$, as the future or prospects. Thus we have a recursive value function definition for the objective function as.

$$\begin{aligned} J_T(X_T, y_T) &= V_T(Z_{1:T}, Y_{1:T}, \mathbb{P}_0(Z_{1:T})) \\ J_t(X_t, U_t) &= \max_{y_t} \left(g_t(X_t, y_t, U_t) + \sum_{i=1}^{N_{t+1}} J_{t+1}(X^i_{t+1}, U^i_{t+1}) \right) \end{aligned} \tag{6.14}$$

for some functions $V_T(\cdot)$ and $g_t(\cdot)$.

We prove the theorem using a direct construction and induction on T. Recall ${}_{\Omega^j_t}S_{t:T}$ denotes a substrategy derived from S. It is easy to see that the theorem holds for $T = 1$, with $\tilde{V} = V$. Suppose that the theorem holds for $T = k$ for some $k > 1$.

[8] For example, any monotonic transformation of the objective function will preserve Dynamic Consistency.

We will show that it holds when $T = k+1$. Consider the functions $\bar{V}(\cdots)$ and $V^{\dagger}(\cdots)$ defined as

$$\begin{aligned}
V^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0) &\triangleq \max_{S'_T} V(\Omega, \{S_{0:T-1}\ S'_T\}, \mathbb{P}_0)\ ; \\
\bar{V}(\Omega, S, \mathbb{P}_0) &\triangleq V^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0)\ -\ \sum_{i=1}^{N_T} \max_{S'_T} V(\Omega_T^i, \{S_{0:T-1}\ S'_T\},\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&\quad +\ \sum_{i=1}^{N_T} V(\Omega_T^i, S,\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&= V^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0)\ -\ \sum_{i=1}^{N_T} \max_{S'_{T|Z^i}} V(\Omega_T^i, \{S_{0:T-1}(Z_{0:T-1}^i)\ S'_{T|Z^i}\},\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&\quad +\ \sum_{i=1}^{N_T} V(\Omega_T^i, \{S_{0:T-1}(Z_{0:T-1}^i)\ S_{T|Z^i}\},\ {}_{\Omega_T^i}\mathbb{P}_0)\ . \qquad (6.15)
\end{aligned}$$

It then follows that

$$\begin{aligned}
\max_{\{S'_T\}} \bar{V}(\Omega, \{S_{0:T-1}\ S'_T\}, \mathbb{P}_0) &= \max_{\{S'_T\}} V(\Omega, \{S_{0:T-1}\ S'_T\}, \mathbb{P}_0) \\
&= V^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0)\ . \qquad (6.16)
\end{aligned}$$

Hence, if one uses the objective $\bar{V}$ in the SDP instead of V, then the set of optimal partial strategies $S^*_{0:T-1}$ does not change. If $S^*_{0:T-1}$ denotes any such partial strategy, then, from (6.15), $\{S^*_{0:T-1}\ S^A_T\}$ is optimal for the modified SDP if and only if S^A_T solves

$$\max_{S_T} \quad V(\Omega_T^i, \{S^*_{0:T-1} S_T\},\ {}_{\Omega_T^i}\mathbb{P}_0)$$

for all i. Because of the Bayesian update rule (6.8), and as outcomes with zero probabilities do not impact V, this is exactly the same condition as each partition component ${}_{\Omega_T^i}S^A_T$ of S^A_T solving

$$\max_{S_{T|Z^i}} \quad V(\Omega, \{S^*_{0:T-1}\ S_{T|Z^i}\}, \mathbb{P}_T | Z^i_{1:T}, S^*_{0:T-1}(Z^i_{1:T-1}))$$

for all i. Hence the composite strategy $\{S^*_{0:T-1}\ S^A_T\}$ is also optimal for the original SDP because of strong dynamic consistency.

Thus SDPs with objectives V and $\bar{V}$ have the same set of optimal policies and take identical values at optimality.

Now if in (6.15), we replace $V^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0)$ by another function, say, $V_1^{\dagger}(\Omega, S_{0:T-1}, \mathbb{P}_0)$

such that they have common optimal partial strategies $S_{0:T-1}$ and the same optimal value, then it follows that the SDP obtained with the objective function

$$\tilde{V}(\Omega, S, \mathbb{P}_0) \triangleq V_1^\dagger(\Omega, S_{0:T-1}, \mathbb{P}_0) - \sum_{i=1}^{N_T} \max_{S'_T} V(\Omega_T^i, \{S_{0:T-1}\ S'_T\},\ {}_{\Omega_T^i}\mathbb{P}_0) + \sum_{i=1}^{N_T} V(\Omega_T^i, S,\ {}_{\Omega_T^i}\mathbb{P}_0) \tag{6.17}$$

will also have the same set of optimal strategies and the same optimal value.

Now, note that the SDP with objective function

$$V^\dagger(\Omega, S_{0:T-1}, \mathbb{P}_0) = \max_{S'_T} V(\Omega, \{S_{0:T-1}\ S'_T\}, \mathbb{P}_0)$$

is itself a strongly dynamically consistent SDP with $T-1 = k$ stages. From the induction hypothesis, then there exists a function $V_1^\dagger(\cdots)$ such that $V_1^\dagger(\cdots)$ and $V^\dagger(\cdots)$ have the same set of optimal strategies and the same optimality value. Further, $\tilde{V}_1(\cdot)$ has the sum decomposable representation

$$V_1^\dagger(\Omega, S_{0:T-1}, \mathbb{P}_0) = \sum_{t=0}^{T-1} \sum_{i=1}^{N_t} V_t^\dagger(Z_{1:t}^i, Y_{0:t}^i, \Omega_t^i, {}_{\Omega_t^i}\mathbb{P}_0) ;$$

and

$$\max_{\{S_{1:T-1}\}} V_1^\dagger(\Omega, \{y_0\ S_{1:T-1}\}, \mathbb{P}_0) = \max_{\{S_{1:T-1}\}} V(\Omega, \{y_0\ S_{1:T-1}\}, \mathbb{P}_0) . \tag{6.18}$$

Using $V_1^\dagger(\cdots)$ in (6.17), we get the following equivalent objective for the objective $\tilde{V}$ of the modified SDP

$$\begin{aligned}
\tilde{V}(\Omega, S, \mathbb{P}_0) &= \sum_{t=0}^{T-1} \sum_{i=1}^{N_t} V_t^\dagger(Z_{1:t}^i, Y_{0:t}^i, \Omega_t^i,\ {}_{\Omega_t^i}\mathbb{P}_0) - \sum_{i=1}^{N_T} \max_{\{S'_T\}} V(\Omega_T^i, \{S_{0:T-1}\ S'_T\},\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&\quad + \sum_{i=1}^{N_T} V(\Omega_T^i, \{S_{0:T-1} S'_T\},\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&= \sum_{t=0}^{T-1} \sum_{i=1}^{N_t} V_t^\dagger(Z_{1:t}^i, Y_{0:t}^i, \Omega_t^i, {}_{\Omega_t^i}\mathbb{P}_0) - \sum_{i=1}^{N_T} \max_{Y_T^i} V(Z_{1:T}^i, Y_{1:T-1}^i, Y_T^i,\ {}_{\Omega_T^i}\mathbb{P}_0) \\
&\quad + \sum_{i=1}^{N_T} V(Z_{1:T}^i, Y_{1:T}^i,\ {}_{\Omega_T^i}\mathbb{P}_0) ,
\end{aligned}$$

which is of the desired form.

Also, from (6.16), (6.17) and (6.18), it follows that

$$\max_{S_{1:T}} \tilde{V}(\Omega, \{y_0, S_{1:T}\}) \quad = \quad \max_{S_{1:T}} V(\Omega, \{y_0, S_{1:T}\}\ ).$$

The formulation in (6.14) is strongly separable. □

We have thus established that every dynamically consistent optimal strategy can be thought of as to arise from a strongly separable formulation. Moreover, the formulation has a sum decomposable representation as described in the statement of the theorem.

The above result shows that there exists an interesting relationship between strongly dynamically consistent formulations and formulations that can be solved using Bellman's principle and dynamic programming. Dynamic programming is basically an algorithmic trick that is used in deterministic problems as well, such as the shortest path problem when the problem is decomposable.

The essential requirement for strong dynamic consistency over time is then that optimal controls corresponding to mutually exclusive events can be obtained independently of each other, thus negating the need for a "co-ordination" amongst DMs at different stages. Note that in all the examples presented in Section 6.2 which suffered from dynamic inconsistency, the principal DM needed to co-ordinate strategies of the future DMs, even when these strategies would never be simultaneously executed on any sample path. This happened because either the objective or some of the constraints were "coupling".

6.4 Dynamic Consistency and Dynamic Risk and Asset Management

In the last section, we examined the problem of dynamic consistency for a general SDP and the relationship between dynamic consistency and objective functions. The dynamic inconsistency of commonly used financial metrics - variance, VaR, CVaR, Sharpe Ratio, when used as or part of objective functions can thus be easily traced to the fact that they are "non-separable" across disjoint events. Distribution invariant dynamic Risk-measures are representable as expected value of shortfall risk and these lead to dynamically consistent objectives. Dynamic coherent risk measures that are not of the expected value type have been proposed by Hardy and Wirch [103], Roorda and Schumacher [105], Artzner et. al [5]. An example is the iterated CVaR (iCVaR). These risk-measures do not have a simple representation however and can only be defined recursively. The definition thus depends on the number of time-steps or discretization level used.Note that the iterated VaR measure proposed in Cheridito and Stadje [35] has similar drawbacks besides being not coherent. Expected utility or expected shortfall then appear to be the most reasonable choices of performance metrics if dynamic consistency of strategies is desired.

We propose another alternative - Use expectations with respect to a probability law distorted in a dynamically consistent fashion e.g.,

$$V = \sum_{i=1}^{N} (p^i)^\alpha U(Z^i, Y^i).$$

If $\alpha < 1$, then effectively, this distorted measure will re-emphasize tail or extreme events. The function $U(\cdot)$ maybe chosen so as to ensure that tail events with negative risk are de-emphasized, e.g, by making it a shortfall function. It can be verified relatively easily that this objective will lead to dynamically consistent SDPs with Bayesian updates. Note, however this objective function is not “distribution-invariant” w.r.t. to $U(Z^i, Y^i)$ treated as a random variable. This does not create any difficulty in our setting as each state has a distinct identity. Such a risk measure may also be defined for an intermediate evaluation, on partial sequences $Z_{1:t}$ and $Y_{0:t}$. Note the similarity with and differences from the distortion idea introduced by Choquet, [37].

Practically speaking, however, that use of VaR and variance and Sharpe-ratio in or as performance metrics is widespread in practice. We now consider how dynamic inconsistency issues that may result from their use maybe dealt with. One approach would be to use them in a static sense as indicated in Cuoco, He and Issaenko [46]. However, this would not help in modeling risks related to terminal positions.

We again use the DMs at different time and stages viewpoint in the following discussion to emphasize the compromises that need to be made when faced with dynamically inconsistent SDPs. Dynamic inconsistency essentially means that it is impossible to satisfy all these DMs simultaneously and in all eventualities.

There are three possible courses of action when dealing with dynamic inconsistency, (see Lotito, [88], Machina, [89], McClennen, [92]).

Myopic Choice: Ignore the fact that the SDP is dynamically inconsistent. Re-solve the problem every time period and implement only the current strategy found to be optimal for the period of interest. This is essentially the idea behind

Rolling Horizon optimization and is often used in engineering. This is probably most undesirable as the resulting strategy will be sub-optimal for all the DMs in our setting.

Sophisticated Choice: Anticipate the dynamic inconsistency and implement current controls in a way that compensates for future dynamic inconsistency that is foreseeable. This effectively amounts to using Backward Induction to solve the Sequential Decision Problem. In our setting, this is a game-theoretic or a Nash-equilibrium solution between different DMs at different stages.

Resolute Choice: Forego solving the problem every period and commit to implement the optimal strategy obtained initially in all periods without resolving the problem in subsequent periods. This is tantamount to taking a non-consequentialist view point and to never update the pseudo-probability weights initially assigned even in face of new information. See how γ-people behave as in [89]. In this case the objective of DM at time 0, the principal agent takes priority over all other DMs.

In the context of preference theory, Machina [89] presents normative arguments and examples supporting the rationale for an agent to adopt a non-consequentialist approach arguing that events that did not or cannot happen represent risk that has been borne and must continue to be accounted for. We give a concrete and real life example here adoption of such a non-consequentialist approach can be objectively justified.

When is ignoring New Information appropriate?

The question whether to use a consequentialist(Bayesian) or non-consequentialist(none) approach for probability updates depends on the actual objective that the SDP is a model of. To illustrate let us take the example of an asset manager with a high turnover portfolio. The manager's performance will be eventually judged by analysts by a measurement of the Sharpe-ratio of her returns over several well-defined trading horizons. The manager, who is aware of this criterion then should naturally seek to optimize this Sharpe-ratio and can device a dynamic asset allocation strategy with 'Sharpe-Ratio' as the objective at the beginning of a trading horizon. While a legitimate modeling objective, it is relatively easy to check that the Sharpe-ratio which is the ratio of expected annual return and its standard deviation is a non-separable objective and hence will lead to dynamic inconsistency with Bayesian updates. The manager can use a surrogate objective that is dynamically consistent, but that will be sub-optimal.

Suppose we also assume that the manager does not perceive liquidity or bankruptcies related risk to be significant and thus can expect to be in business for a large period and that the market performance over different trading horizons is independent and identically distributed. We argue that in this case, it is appropriate for the manager to take up as her objective, Sharpe-ratio of the return at the end of the horizon, with a non-consequential approach for updating probabilities. Informally, this resolute choice strategy chooses that distribution of end-of-period returns that has the maximum Sharpe-ratio amongst all the feasible distributions. With large number of independent samples, this will give the manager the optimal performance metric.[9]

In this context, both Sophisticated Choice and Myopic Choice Strategies, insofar as they are distinct from Resolute Choice will be "sub-optimal". While the non-consequentialist approach used here may also be interpreted as accounting for risk already borne in the sense of [89], in essence what justifies this policy is the fact that

1. the objective is measured over well-defined time horizons and
2. there will be repeated trials of the same strategy

We believe that these two factors are key to validating a non-consequentialist or disciplined approach of being committed to the original plan even in light of new information, as the strategy is targeting a distribution profile.

Note that if returns corresponding to different trading periods were used to measure the manager's performance then the measured Sharpe-ratio is likely to be much worse than the blocks that were optimized for. [10] Also, we have only resolved the issue of what problem the asset manager should seek to solve. Solving for optimal dynamic strategy for a non-separable objective is a computationally challenging task but not the subject-matter of discussion here.

Such a non-consequentialist approach cannot be considered appropriate however for all dynamically inconsistent SDPs. Consider the modeling objective for risk management, which is to guard a firm against events that are likely to be catastrophic or of end-game nature. Through risk management while recognizing that eliminating a

[9] Strictly speaking, the manager should seek to maximize the asymptotic statistical estimator of Sharpe-ratio.

[10] This would mean that for an investor seeking to invest in the firm, the time of entry or exit into the fund may now become relevant to its performance! This is more a criticism of the performance metric used.

severe risk scenario might be too expensive or impossible, the DM seeks to limit the possibility of the same to a reasonable level. If such an undesirable event does occur, it is likely to change the playing field for the DM completely. Thus the criterion of repeated trials of a strategy fall through.

To illustrate this idea in the context of dynamic risk management, using VaR limits, consider a hypothetical case of a day-trader devicing a strategy to maximize returns over a horizon of T periods. The risk that the manager wants to guard herself against is that of a large downward swing in the end-of-the day net value[11]. To keep things simple, we assume that the manager can trade only once during a period in N different assets. The period t gross return vector is denoted by R_t. Let W_t denote the net value of the funds holdings at time t, while X_t the holdings in individual assets. We assume no transaction costs. This problem is typically posed as

$$\begin{aligned} \max_{X_t:0\leq t\leq T-1} \mathbb{E}[W_T] \; ; & \\ \text{s.t.} \qquad W_t &= X_{t-1}' R_t + W_{t-1}, \quad 1 \leq t \leq T \;, \\ \text{VaR}_\alpha(W_t) &> W_0(1-\beta) \;. \end{aligned} \tag{6.19}$$

Here $\text{VaR}_\alpha(X)$ denotes the α percentile in the distribution of X, i.e., if $F_X(\cdot)$ denotes the CDF of X then $\text{VaR}_\alpha(X) = \inf_x\{x : F_X(x) \geq \alpha\}$. β denotes the maximum draw-down and α a level of risk the manager is comfortable taking and is related to her risk-averseness. The VaR related constraint, as it is coupling, will lead to dynamically inconsistent solutions under Bayesian updates. Further, as we argued earlier taking a non-consequentialist approach in this scenario is inappropriate. Also, as VaR, is dynamically inconsistent as an objective, dualizing the constraint or penalizing it will not solve the issue as well. A possible way to tackle this problem is to consider the equivalent problem (6.21)

$$\begin{aligned} \max_{X_t:1\leq t\leq T} \quad & \mathbb{E}[W_T] \; ; \\ \text{s.t.} \quad & W_t = X_{t-1}' R_t + W_{t-1}, \quad 1 \leq t \leq T \;, \\ & \Pr(W_T < W_0(1-\beta)) < \alpha \;. \end{aligned} \tag{6.20}$$

and converting this constraint to a penalty as in a Lagrangian approach. The new

[11]For example, this might lead to a margin call

problem can be posed as

$$\begin{aligned} \max_{X_t:1\le t\le T} \quad & \mathbb{E}[W_T] - \lambda \mathbb{E}[\mathbf{1}_{W_T < W_0(1-\beta)}] \\ \text{s.t.} \quad & W_t = X'_{t-1} R_t + W_{t-1}, \quad 1 \le t \le T \; . \end{aligned} \tag{6.21}$$

This formulation though dynamically consistent is not equivalent to the original, but poses the manager's problem much more transparently. The formulation raises the question about what an appropriate value of multiplier λ should be. In principle, answering this question should be no more difficult than choosing "sensible" values of α. Indeed, if one had a target level of α available, then λ can be appropriately tuned to attain those or better risks of draw-downs at optimum[12]. The penalty based formulation also offers a much more straightforward interpretation. In this case it should be simply interpreted as a trade-off between cost of excessive large draw-down e.g., say a margin call or risk to credit perception, for a marginal increase in expected gains. Note that, if risk related to intermediate positions is of concern, then it can be addressed in a similar way.

Ideally, one would like to switch to a Dynamically Consistent Objective and formulation whenever possible. Managers or DMs often do not control the performance metrics on which they will be measured. It is in general difficult to promote or evangelize a new industry standard because of legacy reasons. For example, expected utility framework has been around for years but has never been a popular measure to rate managers on asset or risk management.

6.5 Conclusion

In this chapter, we studied the issue of dynamic inconsistency for general Sequential Decision Problems(SDPs). The issue is of particular importance in finance as many of the commonly used performance as well as acceptability criteria, can lead to dynamic inconsistency unless employed carefully. We provided several simple but illustrative examples for the same.

We then studied the general SDP framework and investigated the conditions that an objective function must satisfy in order to lead to a dynamically consistent formulation. In this context, we noted that dynamic consistency is almost equivalent to correctness of dynamic programming implied solutions. We also showed that any

[12] The strategy implemented in that case will be the same as the optimal pre-commitment strategy for (6.19).

dynamically consistent strategy can be thought of as arising from an objective that is sum-decomposable over time and mutually exclusive events. Based on these insights, we also proposed an alternate class of performance metrics based on shortfall expectation with respect to probabilites distorted by a suitable power factor. These can be used to emphasize tail events, the same rationale behind the VaR and CVaR metrics, while being dynamically consistent.

We then discussed how a DM faced with dynamically inconsistent SDPs may cope with this inconsistency. In certain cases, it may be reasonable for the DM to take a non-consequentialist approach and follow a pre-committed strategy. Some dynamically inconsistent formulations can be made consistent by rephrasing.

Many interesting directions are possible in this line of research. On the framework side, corresponding conditions for dynamic consistency for an infinite sample space would be an important extension of results. Another extension possible will be to consider the effect of ambiguity in probability distributions over the set of states, as in practice the probability distribution is almost never known precisely and the SDP goal is to find that strategy that works reasonably well for all classes of probability distributions in a family. If dynamic consistency, is a desirable property in this setting, then it will be interesting to see, what additional conditions on the SDP framing (besides separability across mutually exclusive events) are needed to guarantee dynamic consistency.

Chapter 7

Summary and Closing Remarks

We examined three different topics in financial modeling. The topics, though diverse, together illustrate the power as well as limitations of modeling, which is at the heart of Operations Research and separates it from pure mathematics and computer science.

For Employee Stock Options, we noted that, by not considering an employee's stock option portfolio holistically, traditional models can leave out a significant determinant of employee's exercise policies and thus indirectly its cost to the employer. Augmenting the models to incorporate this effect allows us to get an insight on the nature of this impact. As we studied in Chapter 2, the impact is in general to make the cost of a portfolio smaller than the sum of its parts. The models also show that issuance of new ESOs, if unanticipated by the employee, can have the surprising effect of changing the costs associated with the unexercised ESOs that the employee possesses. Though useful to generate these insights, the model is not very useful from an implementation viewpoint.

A different approach, based on risk-management and portfolio optimization presented in Chapter 3, allowed us to jointly model exercise behavior of multiple ESOs for an employee, while being amenable to computation. The risk-management based model agrees at a basic level with traditional models in prioritizing options for exercises and is based on an intuitive criterion of ranking options on the basis of a pseudo Sharpe ratio. This suggests another related model for exercise behavior, which while making option exercises independent, allows us to get closed-form bounds on their cost to the company.

An interesting and useful direction of research in this context will be to explore how exercise behavior predicted by these models compares to empirically observed behavior and quantifying their impact on pricing. Further refinement of models and strategies to hedge out the cost of ESO at the onset are other possible research directions.

We then analysed the GWB for life options in Chapters 4 and 5. Though, in wide use in the Variable Annuity (VA) space, a systematic understanding of this implied feature has only recently begun in the literature. We proposed and analysed a continuous time version of this product, with simplifying assumptions on residual mortality rates of the investor population in Chapter 4. The model, though cannot be used for pricing such schemes in real-life, provided useful insights in the key determinants and risk-factors for valuations. What is troubling for the GWB series of products, is that the models typically used for pricing securities would tend to disagree significantly on valuations as we saw in Chapter 5. The combination of this price ambiguity as well as several un-hedgeable factors such as investor behavior, imply that even with "hedging", GWB for life products will entail significant risk. Also, we observed that the product has insufficient price discrimination and is subject to adverse selection.

Our findings suggest that regulators and investors must lay down standards for valuation of complex securities like GWB. It will be useful to have a sound and practically useable framework to describe the risk in option pricing due to model ambiguity. More specific to the GWB, interesting directions of research include how one can possibly re-design the product to make it less prone to adverse selection as well as adverse dynamic behavior by investors so that it is not "gamed". Another possible direction is applying revenue management ideas to this product and find the right pricing premium for GWB given an individuals risk-averseness to financial and longevity risks.

The issue of dynamic inconsistency, as discussed in Chapter 6 seems almost pervasive in finance. It appears that criteria beyond expected utilities, will either be intractable or run into issues related to dynamic inconsistencies. We presented an alternative family of criteria in terms of expectations with respect to distorted state probabilities. These objectives emphasize extreme events while preserving dynamic consistency. The dynamic consistency issue is more severe when dealing with "chance guarantees". In essence, chance guarantees are by and large inconsistent in a dynamic setting, unless accompanied by a description of circumstances in which they will hold or fail.

It will be interesting to extend and examine the results presented to settings with infinite and uncountable sample-space. Effective risk and performance metrics in a dynamic setting that are easy to use and implement in practice will be a great tool in financial modeling but seem difficult to find.

From a broader perspective, although a general evaluation of modeling frameworks is difficult, the problems considered in this thesis highlight the importance of "commonly accepted good modeling principles" - they should incorporate key determinants of the problem, be solveable for real life use and be theoretically sound. Often, for complex examples, it is difficult to find a single model that will account for all nuances and complexities of the real-life problems and be still computationally solveable. We believe a robust framework that will allow a decision maker to simultaneously consider actionable recommendations from a family of models and choose the most desireable one, so as to minimize unpleasant surprises, will be useful in most application settings. Development and analysis of the efficacy of such a "Robust Model Optimization" framework in concrete settings will be an exciting and useful research avenue in the field.

Appendix A

Relative Order of ESO Exercises

In this appendix, we illustrate a simple example, where the order of exercise as stated in Lemma 2.5 can be violated for a ‘consumption’ type utility model. The Employee’s exercise problem is formulated in the same spirit as the model suggested in [75]. More concretely, the exercise policy is obtained as a solution to the following optimization problem.

$$
\begin{aligned}
\max \quad & \mathbb{E}\left[\sum_{t=1}^{T} U(t, P_t)\right] \\
P_t \quad &= \sum_{i=1}^{N} x_{i,t}(S_t - K_i)^+ \qquad \text{(A.1)} \\
& x_{i,t} \text{ is } \mathcal{F}_t - \text{measurable.} \\
\sum_{t=0}^{T} x_{i,t} &= \alpha_i \ldots 1 \leq i \leq N \\
x_{i,t} &= 0 \text{ if } t < V_i \text{ or } t > T_i \qquad \text{(A.2)}
\end{aligned}
$$

Now consider, the following simple instance of this problem, where the employee has two types of ESOs, i.e., $N = 2$, with strikes $K_1 = 70, K_2 = 90$.Both options are already vested and have common expiry $T = 1$. The grant size α_1, α_2 are each 100. The current stock price $S_0 = 100$ and the employee believes that the stock price at $T = 1$ could be either 80 or 120 with equal probability $\frac{1}{2}$. Also, we take $U(0, P) = U(1, P) = \ln(10 + P)$. The employee’s exercise problem for this case can be written as

$$\max \ln(30x_1 + 10x_2 + 10) + \frac{1}{2}\ln(10 + (100 - x_1)10) + \frac{1}{2}\ln(10 + (100 - x_1)50 + (100 - x_2)30)$$

The optimal solution to this problem is $x_1 = 33.47$, $x_2 = 100$. Thus, in this case all options with the highere strike i.e., strike 90, are in fact exercised before the options with a lower strike 70 are exhausted.

Appendix B

Supplementary Results for Chapter 3

We will find it useful to employ the following property of normal distributions. The property can be verified numerically, but for completion we provide a proof.

Lemma B.1. *The function* $g(x) \triangleq \frac{\Phi(x)}{N(x)} + x$ *is increasing in* x*.*

Proof. Consider

$$
\begin{aligned}
g'(x) &= 1 - \left(\frac{\Phi(x)}{N(x)}\right)^2 - \left(\frac{\Phi(x)}{N(x)}\right) x\Phi(x) \\
&= \frac{1}{N^2(x)} \left(N^2(x) - xN(x)\Phi(x) - \Phi^2(x)\right) \\
&= \frac{1}{N^2(x)} p(x) \text{ where,} \\
p(x) &\triangleq N^2(x) - xN(x)\Phi(x) - \Phi^2(x) \ . && \text{(B.1)} \\
\Rightarrow p'(x) &= N(x)\Phi(x) + x\Phi^2(x) + x^2 N(x)\Phi(x) && \text{(B.2)} \\
&= \Phi(x)h(x), \text{ where,} \\
h(x) &\triangleq N(x) + x\Phi(x) + x^2 N(x) \ . && \text{(B.3)} \\
\Rightarrow h'(x) &= 2\Phi(x) + 2xN(x) && \text{(B.4)} \\
\Rightarrow h''(x) &= 2N(x) > 0 \ . && \text{(B.5)}
\end{aligned}
$$

Now from (B.4) and (B.3) respectively, $\lim_{x \to -\infty} h'(x) = 0$ and $\lim_{x \to -\infty} h(x) = 0$. Hence $h(x) \geq 0$ and consequently $p'(x) \geq 0$ using the definition in (B.3). Once again we have $\lim_{x \to -\infty} p(x) = 0$. Thus $p(x) \geq 0$ or equivalently $g'(x) \geq 0$ $\forall x$ using (B.1). This completes the proof. □

Lemma B.2. *The function $\frac{C}{\sigma S\delta}$ is increasing in S.*

Proof. Note that,

$$\begin{aligned}\frac{C}{S\delta} &= 1 - \frac{Ke^{-rT}N(d-\sigma\sqrt{T})}{SN(d)} \\ &= 1 - e^{\frac{1}{2}\sigma^2 T}\frac{e^{-\sigma\sqrt{T}d}N(d-\sigma\sqrt{T})}{N(d)}\end{aligned}$$

where d as before is given by $d = \frac{\ln\frac{S}{K}}{\sigma\sqrt{T}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T}$. Since d is an increasing function of S, it suffices to show that $f(d) = \frac{e^{-\sigma\sqrt{T}d}N(d-\sigma\sqrt{T})}{N(d)}$ is decreasing in d to prove the lemma. Now

$$\begin{aligned}f'(d) &= \frac{e^{-\sigma\sqrt{T}d}}{(N(d))^2}\Big(N(d)\Phi(d-\sigma\sqrt{T}) - \sigma\sqrt{T}N(d)N(d-\sigma\sqrt{T}) - N(d-\sigma\sqrt{T})\Phi(d)\Big) \\ &= \frac{e^{-\sigma\sqrt{T}d}N(d-\sigma\sqrt{T})}{N(d)}\left(\left(\frac{\Phi(d-\sigma\sqrt{T})}{N(d-\sigma\sqrt{T})} + (d-\sigma\sqrt{T})\right) - \left(\frac{\Phi(d)}{N(d)} + d\right)\right)\end{aligned}$$

Here $\Phi(x)$ denotes the normal density function. Hence, to show that $f(d)$ is decreasing it suffices to show that the function $g(x) \triangleq \frac{\Phi(x)}{N(x)} + x$ is increasing in x. This was proved in Lemma B.1 in Appendix B

□

Corollary B.1. *The delta-barrier function $B(S, K, \tau) = \frac{C-E}{\sigma S\delta}$ achieves its peak at $S = K$.*

Proof. From Lemma 3.2, we know that $B(S, K, \tau)$ is decreasing in S for $S > K$. For $S < K$, $B(S, K, \tau) = \frac{C}{\sigma S\delta}$ and this function was shown to be increasing in S in Lemma B.2. Hence it follows that the delta-barrier achieves its maxima at $S = K$. □

Corollary B.2. *The lower bound function to C^{ν}, $C^{\nu-} = (C - \nu\sigma S\delta)^+$ is increasing in S.*

Proof. We only consider the case when S is such that $C - \nu\sigma S\delta \geq 0$, as the other case is trivial. Consider $S' > S$. By Lemma B.2

$$\begin{aligned}
\frac{C(S',K,\tau)}{\sigma S'\delta(S',K,\tau)} &\geq \frac{C(S,K,\tau)}{\sigma S\delta(S,K,\tau)} \\
\Rightarrow \frac{C(S',K,\tau) - \nu\sigma S'\delta(S',K,\tau)}{\sigma S'\delta(S',K,\tau)} &\geq \frac{C(S,K,\tau) - \nu\sigma S\delta(S,K,\tau)}{\sigma S\delta(S,K,\tau)} \\
\Rightarrow C(S',K,\tau) - \nu\sigma S'\delta(S',K,\tau) &> C(S,K,\tau) - \nu\sigma S\delta(S,K,\tau) \\
&\ldots\ \text{since, } (\ C(S,K,\tau) - \nu\sigma S\delta(S,K,\tau) > 0 \\
&\quad \text{and } S'\delta(S',K,\tau) > S\delta(S,K,\tau) > 0\ )
\end{aligned}$$

□

Corollary B.3. *At $T' = T - T_N$, where T' is defined as in the proof of Lemma 3.4,*

$$C - \nu\sigma S\delta > 0 \iff S > K\ .$$

Proof. The property follows directly from Lemma B.2 and the fact that at T', $\frac{C}{\sigma S\delta} = \nu$ for $S = K$. □

Lemma B.3. *For $T' = T - T_N$, where T' is as defined in the proof of Lemma 3.4,*

$$\begin{aligned}
& \mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \\
= \quad & C_0 - \nu\sigma S_0\delta_0 \\
& + Ke^{-rT} N_2\left(d(S_0,K,T) - \sigma\sqrt{T}, -d(S_0,K,T') + \sigma\sqrt{T'}, -\sqrt{\frac{T'}{T}}\right) \\
& -(1-\nu)S_0 N_2\left(d(S_0,K,T), -d(S_0,K,T'), -\sqrt{\frac{T'}{T}}\right).
\end{aligned}$$

where, $d(S,K,T) = \frac{\ln(\frac{S}{K})}{\sigma\sqrt{T}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T}$ and $N_2(x,y,\rho) = \Pr(X \leq x, Y \leq y)$, where X, Y are two jointly normal random variables, with variances 1 and correlation ρ.

Proof. We know by definition that at , $C(K,K,T-T') = \nu\sigma K\delta$, at $S = K$. Using the result proved in Corollary B.3 then,

$$\begin{aligned}
\mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] &= \mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'}); S_{T'} \geq K] \\
&= \mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})] + \mathbb{E}[e^{-rT'}(\nu\sigma S_{T'}\delta_{T'} - C_{T'}); S_{T'} < K] \\
&= C_0 - \nu\sigma S_0\delta_0 + \mathbb{E}[e^{-rT'}e^{-rT_N}KN(d - \sigma\sqrt{T_N}); S_{T'} < K] \\
&\quad -(1-\nu\sigma)\mathbb{E}[e^{-rT'}S_{T'}N(d); S_{T'} < K] \\
&= C_0 - \nu\sigma S_0\delta_0 + Ke^{-rT}\mathbb{E}[N(d - \sigma\sqrt{T_N}); S_{T'} < K] \\
&\quad -(1-\nu\sigma)\mathbb{E}[e^{-rT'}S_{T'}N(d); S_{T'} < K] \qquad \text{(B.6)}
\end{aligned}$$

where, we use $d = d(S_{T'}, K, T_N)$.Now,

Let us define z_0, z_1 as follows,

$$\begin{aligned}
z_0 &\triangleq \frac{\ln(\frac{S_{T'}}{S_0})}{\sigma\sqrt{T'}} - (\frac{r}{\sigma} - \frac{\sigma}{2})\sqrt{T'} \\
z_1 &\triangleq \frac{\ln(\frac{S_T}{S_{T'}})}{\sigma\sqrt{T_N}} - (\frac{r}{\sigma} - \frac{\sigma}{2})\sqrt{T_N}
\end{aligned}$$

Thus,

$$S_T = S_0 e^{(r-\frac{1}{2}\sigma^2)T} e^{\sigma\sqrt{T'}z_0 + \sigma\sqrt{T_N}z_1}.$$

$$\begin{aligned}
&\mathbb{E}[N(d - \sigma\sqrt{T_N}); S_{T'} < K] \\
=\ &\Pr(S_T > K, S_{T'} < K) \\
=\ &\Pr(\sigma\sqrt{T'}z_0 + \sigma\sqrt{T_N}z_1 > -\ln\frac{S_0}{K} - (r - \frac{1}{2}\sigma^2)T, z_0 \leq -d(S_0, K, T') + \sigma\sqrt{T'}) \\
=\ &\Pr\left(-\sqrt{\frac{T'}{T}}z_0 - \sqrt{\frac{T_N}{T}}z_1 < d(S_0, K, T) - \sigma\sqrt{T}, z_0 \leq -d(S_0, K, T') + \sigma\sqrt{T'}\right) \\
=\ &N_2\left(d(S_0, K, T) - \sigma\sqrt{T}, -d(S_0, K, T') + \sigma\sqrt{T'}, -\sqrt{\frac{T'}{T}}\right)
\end{aligned}$$

Now,

$$
\begin{aligned}
& \mathbb{E}[e^{-rT'}S_{T'}N(d); S_{T'} < K] \\
= & \int_{-\infty}^{-d(S_0,K,T')+\sigma\sqrt{T'}} e^{-rT'}S_0 e^{(r-\frac{1}{2}\sigma^2)T'+\sigma\sqrt{T'}z_0} \times \\
& N\left(\frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T_N}} + \frac{\sqrt{T'}}{\sqrt{T_N}}z_0 + (\frac{r}{\sigma} - \frac{\sigma}{2})\frac{T'}{\sqrt{T_N}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T_N}\right)\Phi(z_0)dz_0 \\
= & S_0\int_{-\infty}^{-d(S_0,K,T')+\sigma\sqrt{T'}} N\left(\frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T'}} + \frac{\sqrt{T'}}{\sqrt{T_N}}z_0 + (\frac{r}{\sigma} - \frac{\sigma}{2})\frac{T'}{\sqrt{T_N}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T_N}\right)\Phi(z_0 - \sigma\sqrt{T'})dz_0 \\
= & S_0\int_{-\infty}^{-d(S_0,K,T')} N\left(\frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T_N}} + \frac{\sqrt{T'}}{\sqrt{T_N}}z + (\frac{r}{\sigma} + \frac{\sigma}{2})\frac{T'}{\sqrt{T_N}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T_N}\right)\Phi(z)dz \\
= & S_0\int_{-\infty}^{-d(S_0,K,T')} N\left(\frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T_N}} + \frac{\sqrt{T'}}{\sqrt{T_N}}z + (\frac{r}{\sigma} + \frac{\sigma}{2})\frac{T}{\sqrt{T_N}}\right)\Phi(z)dz \\
= & S_0\int_{-\infty}^{-d(S_0,K,T')} \Pr\left(Z - \frac{\sqrt{T'}}{\sqrt{T_N}}z \leq \frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T_N}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\frac{T}{\sqrt{T_N}}\right)\Phi(z)dz \\
= & S_0\int_{-\infty}^{-d(S_0,K,T')} \Pr\left(\frac{\sqrt{T_N}}{\sqrt{T}}Z - \frac{\sqrt{T'}}{\sqrt{T}}z \leq \frac{\ln\frac{S_0}{K}}{\sigma\sqrt{T}} + (\frac{r}{\sigma} + \frac{\sigma}{2})\sqrt{T}\right)\Phi(z)dz \\
= & S_0 N_2\left(d(S_0,K,T), -d(S_0,K,T'), -\sqrt{\frac{T'}{T}}\right)
\end{aligned}
$$

Hence, from (B.6),

$$
\begin{aligned}
& \mathbb{E}[e^{-rT'}(C_{T'} - \nu\sigma S_{T'}\delta_{T'})^+] \\
= & C_0 - \nu\sigma S_0\delta_0 \\
& +Ke^{-rT}N_2\left(d(S_0,K,T) - \sigma\sqrt{T}, -d(S_0,K,T') + \sigma\sqrt{T'}, -\sqrt{\frac{T'}{T}}\right) \\
& -(1-\nu\sigma)S_0N_2\left(d(S_0,K,T), -d(S_0,K,T'), -\sqrt{\frac{T'}{T}}\right). \qquad \text{(B.7)}
\end{aligned}
$$

□

Appendix C

Illustration of GWB Product Evolution

Table 4.5 shows the evolution of a GWB contract over the life of an individual in a hypothetical scenario for the model described in Section 4.2 of Chapter 4. We take the GWB parameters as follows - guaranteed withdrawal rate $q = 6\%$, minimum waiting period $W = 3$ years, retirement age $A_R = 65$ years, and fee rate $h = 0.65\%$. In this example, the investor opened a VA account with GWB for life feature at the age of 60 with an investment of 100,000. The investor could not take any withdrawals before reaching age 65 or for a waiting period of 3 years whichever is later, which is the first 5 years in this case. At each anniversary, the benefit base is stepped up to the contract value if it falls below the same. The investor is guaranteed to be able to withdraw upto 6% every year beginning the 6^{th} anniversary. Withdrawals are deducted from her VA account until it drops to zero, after which, the shortfall is met by the company. If the withdrawals during any year exceed this amount then the benefit base is reset. Fees amounting to 65 basis points of the benefit base are paid (separately) to the insurance company by the investor every year.

In our example, the investor's benefit base stepped up at contract anniversaries 1, 2, 3, 6, 7 and 9. During year 8, the investor took a withdrawal that exceeded the contract limit and this caused the benefit base to reset. The withdrawal was less than the year's gain in contract value and hence did not incur a surrender charge. The investor's contract value dropped to 0 at the 25^{th} anniversary. The insurance company bore the shortfall in the guaranteed withdrawal level, which was 513 in year 25 and 6921 in year 26. The investor died during year 27. At the end of this year, the company would have returned the residual contract value (NIL in our case) to the investor's beneficiaries. Figure C-1 depicts how the Contract Value C_n and Benefit Base B_n evolve with time for this particular example.

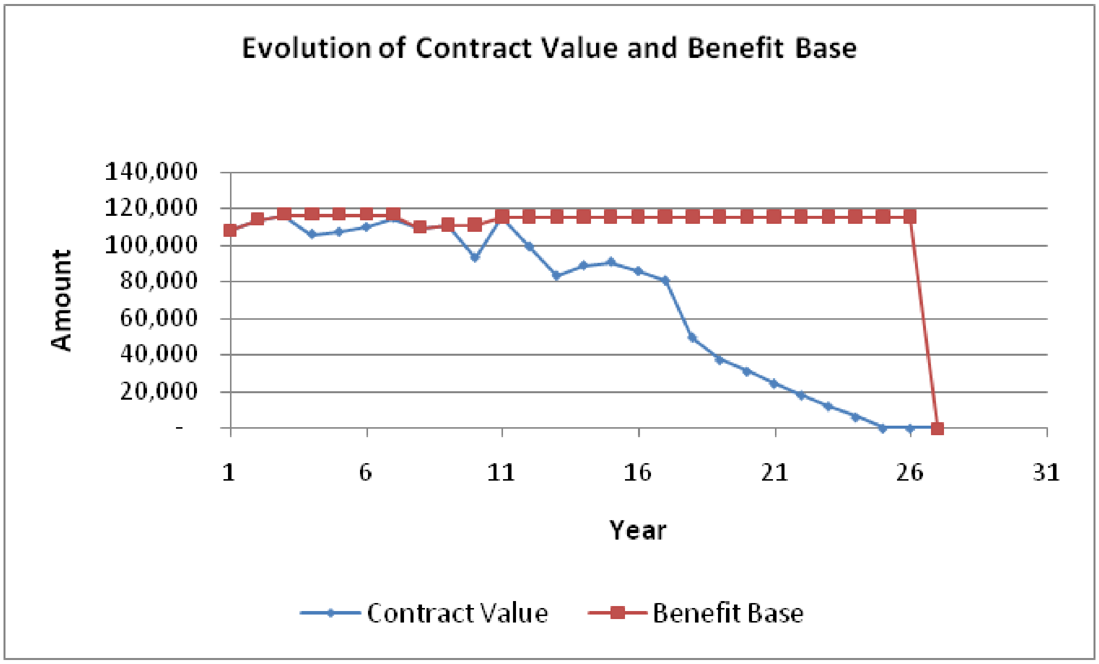

Figure C-1: Evolution of the GWB state variables with time for the example in Table C.1

Anniversary n	***Investor Age***	***VA Fund Return***	***Contract Value before withdrawals***	***Withdrawal Limit***	***Withdrawal Taken***	***Contract Value after withdrawals*** C_n	***Benefit Base*** B_n	***Benefit Base Operation***	***Fees Charged***	***Shortfall financed by the company***
0	60	-	-	-	-	100,000	-	Inception	0	-
1	61	8.00%	108,000	-	-	108,000	108,000	Step-up	702	-
2	62	5.50%	113,940	-	-	113,940	113,940	Step-up	741	-
3	63	2.20%	116,447	-	-	116,447	116,447	Step-up	757	-
4	64	-9.00%	105,966	-	-	105,966	116,447		757	-
5	65	1.12%	107,153	-	-	107,153	116,447		757	-
6	66	9.00%	116,797	6,987	6,987	109,810	116,447	Step-up	757	-
7	67	11.00%	121,889	6,987	6,987	114,903	116,447	Step-up	757	-
8	68	4.00%	119,499	6,987	10,000	109,499	109,499	Reset	712	-
9	69	7.25%	117,437	6,570	6,570	110,867	110,867	Step-up	721	-
10	70	-10.00%	99,781	6,652	6,652	93,129	110,867		721	-
11	71	31.00%	121,999	6,652	6,652	115,347	115,347	Step-up	750	-
12	72	-8.00%	106,119	6,921	6,921	99,198	115,347		750	-
13	73	-9.00%	90,270	6,921	6,921	83,349	115,347		750	-
14	74	15.00%	95,852	6,921	6,921	88,931	115,347		750	-
15	75	10.00%	97,824	6,921	6,921	90,903	115,347		750	-
16	76	2.00%	92,721	6,921	6,921	85,801	115,347		750	-
17	77	2.12%	87,620	6,921	6,921	80,699	115,347		750	-
18	78	-30.00%	56,489	6,921	6,921	49,568	115,347		750	-
19	79	-11.00%	44,116	6,921	6,921	37,195	115,347		750	-
20	80	2.00%	37,939	6,921	6,921	31,018	115,347		750	-
21	81	1.00%	31,328	6,921	6,921	24,408	115,347		750	-
22	82	2.00%	24,896	6,921	6,921	17,975	115,347		750	-
23	83	5.00%	18,874	6,921	6,921	11,953	115,347		750	-
24	84	8.00%	12,909	6,921	6,921	5,988	115,347		750	-
25	85	7.00%	6,407	6,921	6,921	-	115,347		750	513
26	86	-6.00%	-	6,921	6,921	-	115,347		750	6,921
27	Death	5.00%	-	-	-	-	-		0	-

Table C.1: Numerical illustration of evolution of the GWB product related variables and cash-flows in a hypothetical scenario.

Appendix D

Whittaker Functions and some basic properties

The function $\mathrm{WhM}(k,m,z)$ can be defined in terms of hypergeometric function (see Mathworld [91]) as

$$\mathrm{WhM}(k,m,z) \;=\; \exp(-\frac{z}{2})z^{m+\frac{1}{2}}{}_1F_1(\frac{1}{2}-m+k,1+2m,z)\ .$$

${}_1F_1(a,b,z)$ denotes a confluent hypergeometric function of the first kind. It has a power series representation

$$ {}_1F_1(a,b,z) \;=\; 1+\frac{a}{b}z+\frac{a(a+1)}{b(b+1)}\frac{z^2}{2!}+\ldots \qquad \text{(D.1)}$$

and an integral representation

$$ {}_1F_1(a,b,z) \;=\; \frac{\Gamma(b)}{\Gamma(b-a)\Gamma(a)}\int_0^1 \exp(z.t)t^{a-1}(1-t)^{b-a-1}dt\ .$$

Note that,

$$\begin{aligned}\frac{d{}_1F_1(a,b,z)}{dz} &= \frac{a}{b}+\frac{a(a+1)}{b(b+1)}z+\frac{a(a+1)(a+2)}{b(b+1)(b+2)}\frac{z^2}{2!}+\ldots\\ &= \frac{a}{b}{}_1F_1(a+1,b+1,z)\ .\end{aligned}$$

Proposition D.1. *Let*

$$f_1(x) \;\triangleq\; x^{-k}e^{-\frac{a}{2x}}\mathrm{WhM}\left(k,m,\frac{a}{x}\right) \;,$$

where $a > 0$. Then,

$$\lim_{x\to 0^+} f_1(x) \;=\; a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)} \;; \tag{D.2}$$

$$f_1'(x) \;=\; -\left(\frac{1}{2}+k+m\right)x^{-(k+1)}e^{-\frac{a}{2x}}\mathrm{WhM}\left(k+1,m,\frac{a}{x}\right) \;. \tag{D.3}$$

Proof.

$$\begin{aligned}
\lim_{x\to 0^+} f_1(x) &= \lim_{x\to 0^+} a^{-k}\left(\frac{a}{x}\right)^{\frac{1}{2}+k+m} e^{-\frac{a}{x}}{}_1F_1(\frac{1}{2}-k+m,1+2m,\frac{a}{x}) \\
&= a^{-k}\lim_{z\to\infty} z^{\frac{1}{2}+k+m}e^{-z}{}_1F_1(\frac{1}{2}-k+m,1+2m,z) \\
&= a^{-k}\lim_{z\to\infty} z^{\frac{1}{2}+m+k}e^{z}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)\Gamma(\frac{1}{2}+k+m)}\cdot\int_0^1 e^{zt}t^{-\frac{1}{2}-k+m}(1-t)^{-\frac{1}{2}+k+m}dt \\
&= a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)\Gamma(\frac{1}{2}+k+m)}\cdot \\
&\quad \lim_{z\to\infty}\int_0^1 e^{-z(1-t)}t^{-\frac{1}{2}-k+m}(1-t)^{-\frac{1}{2}+k+m}z^{k+\frac{1}{2}+m}dt \\
&= a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)\Gamma(\frac{1}{2}+k+m)}\lim_{z\to\infty}\int_0^1 e^{-zt}t^{-\frac{1}{2}+k+m}(1-t)^{-\frac{1}{2}-k+m}z^{k+\frac{1}{2}+m}dt \\
&= a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)\Gamma(\frac{1}{2}+k+m)}\lim_{z\to\infty}\int_0^z e^{-y}y^{-\frac{1}{2}+k+m}(1-\frac{y}{z})^{-\frac{1}{2}-k+m}dy \\
&= a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)\Gamma(\frac{1}{2}+k+m)}\lim_{z\to\infty}\int_0^z e^{-y}y^{-\frac{1}{2}+k+m}dy \\
&= a^{-k}\frac{\Gamma(1+2m)}{\Gamma(\frac{1}{2}-k+m)} \;.
\end{aligned}$$

The property (D.3) was obtained by symbolic differentiation using Maple 9. □

The function $\mathrm{WhW}(k,m,z)$ has the following integral representation (see Mathworld [91]):

$$\mathrm{WhW}(k,m,z) \;=\; \frac{e^{-\frac{z}{2}}}{\Gamma(\frac{1}{2}-k+m)}\int_0^\infty t^{-\frac{1}{2}-k+m}\left(1+\frac{t}{z}\right)^{-\frac{1}{2}+k+m}e^{-t}dt \;. \tag{D.4}$$

Proposition D.2. *Let*

$$f_2(x) \quad \triangleq \quad x^{-k} e^{-\frac{a}{2x}} \mathrm{WhW}\left(k, m, \frac{a}{x}\right) \ ,$$

where $a > 0$. *Then,*

$$\lim_{x\to 0^+} f_2(x) \quad = \quad 0 \ ; \tag{D.5}$$

$$f_2'(x) \quad = \quad x^{-(k+1)} e^{-\frac{a}{2x}} \mathrm{WhW}\left(k+1, m, \frac{a}{x}\right) \ . \tag{D.6}$$

Proof.

$$\begin{aligned}
\lim_{x\to 0^+} f_2(x) &= \lim_{x\to 0^+} a^{-k} \frac{(\frac{a}{x})^{2k} e^{-\frac{a}{x}}}{\Gamma(\frac{1}{2}-k+m)} \int_0^\infty t^{-\frac{1}{2}-k+m} \left(1+\frac{tx}{a}\right)^{-\frac{1}{2}+k+m} e^{-t} dt \\
&= \frac{a^{-k}}{\Gamma(\frac{1}{2}-k+m)} \lim_{z\to\infty} z^{2k} e^{-z} \int_0^\infty t^{-\frac{1}{2}-k+m} \left(1+\frac{t}{z}\right)^{-\frac{1}{2}+k+m} e^{-t} dt \\
&= \frac{a^{-k}}{\Gamma(\frac{1}{2}-k+m)} \left(\lim_{z\to\infty} z^{2k} e^{-z}\right) \cdot \left(\lim_{z\to\infty} \int_0^\infty t^{-\frac{1}{2}-k+m} \left(1+\frac{t}{z}\right)^{-\frac{1}{2}+k+m} e^{-t} dt\right) \\
&= 0 \ .
\end{aligned}$$

The property (D.6) was again obtained by symbolic differentiation using Maple 9. □

Appendix E

Joint distribution of B_t and C_t

Consider the following normalizations of the processs c_t and b_t defined in (4.26) and (4.27):

$$\begin{aligned} y_t &= \frac{c_t}{\sigma} = \frac{1}{\sigma}\ln(C_t) \ ; \\ u_t &= \frac{b_t}{\sigma} = \frac{1}{\sigma}\ln(B_t) \ . \end{aligned}$$

Applying Ito's lemma to (4.3), we get

$$\begin{aligned} dy_t &= \left(\frac{r}{\sigma} - \frac{1}{2}\sigma\right) dt + dZ_t^Q \ , \\ \text{i.e.,}\, y_t &= \eta t + Z_t^Q \ , \\ \text{where } \eta &= \frac{r}{\sigma} - \frac{1}{2}\sigma \ . \end{aligned}$$

Then, by Girsanov's theorem, y_t is a Brownian Motion under the measure R whose Radon-Nikodym derivative is given by

$$\begin{aligned} \frac{dR}{dQ} &= \exp\left(-\frac{1}{2}\eta^2 t - \eta Z_t^Q\right) \ , \\ \text{i.e.,}\, \frac{dQ}{dR} &= \exp(\eta y_t - \frac{1}{2}\eta^2 t) \ . \end{aligned} \tag{E.1}$$

R is absolutely continuous with Q.

Since y_t, is a Brownian motion under R and u_t its supremum, using the well-known refelction principle (see Karatzas and Shreve [78]), we get:

$$\mathbb{P}^R(y_t \le u-a, u_t \ge u) = \begin{cases} \mathbb{P}^R(y_t \ge u+a) & \ldots, \quad z \ge 0 \; ; \\ \mathbb{P}^(u_t \ge u) - \mathbb{P}^R(y_t > u-a, u_t \ge u) & \\ \qquad = \mathbb{P}^R(u_t \ge u) - \mathbb{P}^R(y_t > u-a) & \ldots, \quad z < 0 \; . \end{cases}$$

$$\text{i.e., } \mathbb{P}^R(y_t \le y, u_t \ge u) = \begin{cases} \mathbb{P}^R(y_t \ge 2u-y) & \\ \qquad = \mathrm{N}(\frac{2u-y}{\sqrt{t}}) & \ldots, \quad y \le u \; ; \\ \mathbb{P}^(u_t \ge u) - \mathbb{P}^R(y_t > y, u_t \ge u) & \\ \qquad = \mathbb{P}^R(u_t \ge u) - (1 - \mathrm{N}\left(\frac{y}{\sqrt{t}}\right)) & \ldots, \quad y > u \; . \end{cases}$$

Taking derivatives, we get:

$$f^R_{y_t|u_t \ge u}(y)\mathbb{P}^R(u_t \ge u) = \begin{cases} \frac{1}{\sqrt{t}}\Phi(\frac{2u-y}{\sqrt{t}}) & \ldots, \quad y \le u \; , \\ \frac{1}{\sqrt{t}}\Phi(\frac{y}{\sqrt{t}}) & \ldots, \quad y \ge u \; ; \end{cases} \tag{E.2}$$

$$\text{and } f^R_{y_t,u_t}(y,u) = \begin{cases} \frac{2(2u-y)}{t}\frac{1}{\sqrt{t}}\Phi(\frac{2u-y}{\sqrt{t}}) & \ldots, \quad y \le u \; , \\ 0 & \ldots, \quad y \ge u \; . \end{cases} \tag{E.3}$$

Then using the Radon-Nikodym derivative defined in (E.1) and doing a change of measure to **Q**, we get:

$$f^Q_{y_t|u_t \ge u}(y)\mathbb{P}^Q(u_t \ge u) = \begin{cases} \frac{1}{\sqrt{t}}\Phi\left(\frac{2u-y}{\sqrt{t}}\right)\exp\left(\eta y - \frac{1}{2}\eta^2 t\right) & \ldots, \quad y \le u \; , \\ \frac{1}{\sqrt{t}}\Phi\left(\frac{y}{\sqrt{t}}\right)\exp\left(\eta y - \frac{1}{2}\eta^2 t\right) & \ldots, \quad y > u \; ; \end{cases} \tag{E.4}$$

$$\text{and } f^Q_{y_t,u_t}(y,u) = \begin{cases} \frac{2(2u-y)}{t}\frac{1}{\sqrt{t}}\Phi\left(\frac{2u-y}{\sqrt{t}}\right)\exp\left(\eta y - \frac{1}{2}\eta^2 t\right) & \ldots, \quad y \le u \; , \\ 0 & \ldots, \quad y > u \; . \end{cases} \tag{E.5}$$

Using (E.5), we directly get,

$$f^Q_{c_t,b_t}(z,m) = \begin{cases} \frac{2(2m-z)}{\sigma^2 t}\frac{1}{\sigma\sqrt{t}}\Phi\left(\frac{2m-z}{\sigma\sqrt{t}}\right)\cdot\exp\left(\frac{\eta}{\sigma}z - \frac{1}{2}\eta^2 t\right) & \ldots, \quad z \le m \; ; \\ 0 & \ldots, \quad z > m \; . \end{cases}$$

This is the relation in (4.28), noting $\nu = \sigma\eta$.

Also, using (E.4),

$$
\begin{aligned}
\mathbb{P}^Q(u_t \geq u) &= \int_{-\infty}^{u} \frac{1}{\sqrt{t}}\Phi\left(\frac{2u-y}{\sqrt{t}}\right)\exp\left(\eta y - \frac{1}{2}\eta^2 t\right)dy \\
&\quad + \int_{u}^{\infty} \frac{1}{\sqrt{t}}\Phi\left(\frac{y}{\sqrt{t}}\right)\exp\left(\eta y - \frac{1}{2}\eta^2 t\right)dy \\
&= \exp(2\eta u)\int_{-\infty}^{u} \frac{1}{\sqrt{t}}\Phi\left(\frac{y-2u-\eta t}{\sqrt{t}}\right)dy + \int_{u}^{\infty}\frac{1}{\sqrt{t}}\Phi\left(\frac{y-\eta t}{\sqrt{t}}\right)dy \\
&= \exp(2\eta u)\mathrm{N}\left(\frac{-u-\eta t}{\sqrt{t}}\right) + 1 - \mathrm{N}\left(\frac{u-\eta t}{\sqrt{t}}\right) \ . \\
\text{Hence, } f_{u_t}^Q(u) &= \frac{1}{\sqrt{t}}\Phi\left(\frac{u-\eta t}{\sqrt{t}}\right) + \exp(2\eta u)\cdot\Phi\left(\frac{-u-\eta t}{\sqrt{t}}\right) - 2\eta\exp(2\eta u)\mathrm{N}\left(\frac{-u-\eta t}{\sqrt{t}}\right) \\
&= 2\cdot\frac{1}{\sqrt{t}}\Phi\left(\frac{u-\eta t}{\sqrt{t}}\right) - 2\eta\exp(2\eta u)\cdot\mathrm{N}\left(\frac{-u-\eta t}{\sqrt{t}}\right) \ . \qquad \text{(E.6)}
\end{aligned}
$$

Hence,

$$
f_{b_t}^Q(m) = 2\cdot\frac{1}{\sigma\sqrt{t}}\Phi\left(\frac{m-\sigma\eta t}{\sigma\sqrt{t}}\right) - \frac{2\eta}{\sigma}\exp\left(\frac{2\eta m}{\sigma}\right)\mathrm{N}\left(\frac{-m-\sigma\eta t}{\sigma\sqrt{t}}\right) \ .
$$

This is the same as in (4.29), as $\nu = \eta\sigma$.

Appendix F

Mortality Table

Table F.1: ERISA Section 4050 mortality rates for ages 49 and above (year 2008). Average residual life for each age is computed using the De-Moivre's approximation.

Age	*Mortality Rate (from table)*	*Implied Hazard Rate*	*Average Residual Life (De-Moivre's approximation)*
49	0.0013	0.0013	34.85
50	0.0014	0.0014	33.89
51	0.0015	0.0015	32.94
52	0.0017	0.0017	31.99
53	0.0020	0.0020	31.04
54	0.0022	0.0022	30.10
55	0.0025	0.0025	29.17
56	0.0029	0.0029	28.24
57	0.0034	0.0034	27.32
58	0.0039	0.0039	26.41
59	0.0044	0.0044	25.51
60	0.0050	0.0050	24.63
61	0.0058	0.0058	23.75
62	0.0066	0.0066	22.88
63	0.0076	0.0076	22.03
64	0.0086	0.0086	21.20
65	0.0097	0.0097	20.38
66	0.0110	0.0111	19.57
67	0.0122	0.0123	18.78
68	0.0132	0.0133	18.01
69	0.0144	0.0145	17.24
70	0.0154	0.0155	16.49
71	0.0167	0.0168	15.74
72	0.0183	0.0185	15.00
73	0.0200	0.0202	14.27
74	0.0220	0.0222	13.55
75	0.0243	0.0246	12.84
76	0.0269	0.0273	12.15
77	0.0306	0.0311	11.47
78	0.0346	0.0352	10.82

Continued on next page

Table F.1 – continued from previous page			
Age	*Mortality Rate (from table)*	*Implied Hazard Rate*	*Average Residual Life (De-Moivre's approximation)*
79	0.0391	0.0399	10.19
80	0.0441	0.0451	9.58
81	0.0497	0.0509	9.00
82	0.0558	0.0574	8.44
83	0.0615	0.0635	7.91
84	0.0684	0.0709	7.40
85	0.0757	0.0787	6.91
86	0.0840	0.0877	6.43
87	0.0949	0.0997	5.98
88	0.1063	0.1124	5.55
89	0.1188	0.1265	5.16
90	0.1328	0.1425	4.78
91	0.1461	0.1579	4.44
92	0.1622	0.1770	4.12
93	0.1791	0.1974	3.82
94	0.1954	0.2174	3.55
95	0.2151	0.2422	3.30
96	0.2327	0.2648	3.07
97	0.2530	0.2917	2.85
98	0.2745	0.3209	2.66
99	0.2929	0.3466	2.49
100	0.3116	0.3734	2.32
101	0.3388	0.4136	2.16
102	0.3588	0.4445	2.02
103	0.3807	0.4792	1.90
104	0.4044	0.5182	1.78
105	0.4279	0.5584	1.68
106	0.4491	0.5962	1.60
107	0.4660	0.6274	1.55
108	0.4786	0.6512	1.50
109	0.4881	0.6697	1.47
110	0.4948	0.6828	1.45
111	0.4987	0.6906	1.44
112	0.5000	0.6931	1.44
113	0.5000	0.6931	1.43
114	0.5000	0.6931	1.42
115	0.5000	0.6931	1.40
116	0.5000	0.6931	1.35
117	0.5000	0.6931	1.26
118	0.5000	0.6931	1.08
119	0.5000	0.6931	0.72
120	1.0000	Inf	0.00

Appendix G

Computational Methods for Chapter 5

In this Appendix, we provide a brief description of the computational methods used to value GWB under the different asset return models specified in Section 5.2 of Chapter 5.

G.1 BSM Model

For numerical computations, for each n, we evaluate the values of $l_n(x)$ and $g_n(x)$ as defined in (5.7) and (5.8) respectively, at $M+1$ evenly spaced points in the interval $[0,1]$ with $M = 300$. We then evaluate the integrals in (5.10) and (5.11) using simple linear interpolation. Because asset returns have lognormal distributions, these integrals can be evaluated easily. We describe in detail the procedure for evaluating the function $l_n(\cdot)$. An almost identical procedure with appropriate modifications based on (5.11) is used for evaluating $g_n(\cdot)$.

Let x_i, $0 \le i \le M$ denote the $M+1$ points on the grid, with $x_0 = 0$ and $x_M = 1$. For $0 \le i \le M+1$, let $l_{i,n} \triangleq l_n(x_i)$ denote the values that the function $l_n(\cdot)$ takes on the grid-points. Suppose these values are available for some n. We first find the linear interpolation coefficients $A^l_{i,n}, B^l_{i,n}$; $0 \le i \le M$ as follows:

$$
\begin{aligned}
A^l_{i,n} &= \frac{l_{i+1,n} - l_{i,n}}{x_{i+1} - x_i} \ ; \\
B^l_{i,n} &= l_{i,n} \ .
\end{aligned}
$$

We use the following linear approximation for evaluating $l_n(x)$ at an arbitrary point

x:

$$l_n(x) \approx A^l_{j,n}(x - x_j) + B^l_{j,n} \ , \tag{G.1}$$

where j is such that $x_j \leq x < x_{j+1}$.

Using the approximation in (G.1) in (5.10), we get

$$\begin{aligned}
l_{n-1}(x) &= e^{-(\lambda_n^A + r_n)} \cdot \left\{ x \cdot \int_{-\infty}^{\ln\left(\frac{q_n}{x}\right)} \left(\frac{q_n}{x} - e^z\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \right. \\
&\quad + \int_{\ln\left(\frac{q_n}{x}\right)}^{\ln\left(\frac{1+q_n}{x}\right)} l_n\left(xe^z - q_n\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \\
&\quad \left. + x \cdot l_n(1) \cdot \int_{\ln\left(\frac{1+q_n}{x}\right)}^{\infty} \left(e^z - \frac{q_n}{x}\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \right\} \\
&\approx e^{-(\lambda_n^A + r_n)} \cdot \left\{ \int_{-\infty}^{\ln\left(\frac{q_n}{x}\right)} \left(q_n - xe^z\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \right. \\
&\quad + \sum_{i=1}^{M} \int_{\ln\left(\frac{x_{i-1}+q_n}{x}\right)}^{\ln\left(\frac{x_i+q_n}{x}\right)} \left(A^l_{i-1}(xe^z - q_n) + B^l_{i-1}\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \\
&\quad \left. + l_n(1) \cdot \int_{\ln\left(\frac{1+q_n}{x}\right)}^{\infty} \left(xe^z - q_n\right) \cdot \Phi\left(z; \mu_n, \sigma_n\right) dz \right\} \ .
\end{aligned} \tag{G.2}$$

Now,

$$\begin{aligned}
\int_{-\infty}^{y} \Phi(z; \mu_n, \sigma_n) dz &= \mathrm{N}\left(\frac{y - \mu_n}{\sigma_n}\right) \\
&= \mathrm{N}\left(\frac{y}{\sigma_n} - \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right) \quad \text{and} \\
\int_{y}^{\infty} \exp(z)\Phi(z; \mu_n, \sigma_n) dz &= \exp\left(\mu_n + \frac{1}{2}\sigma_n^2\right) \mathrm{N}\left(\frac{-y}{\sigma_n} + \frac{r}{\sigma_n} + \frac{\sigma_n}{2}\right) \\
&= e^{r_n} \cdot \mathrm{N}\left(\frac{-y}{\sigma_n} + \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right) \ .
\end{aligned}$$

Then, from (G.2),

$$
\begin{aligned}
l_{n-1}(x) \quad \approx \quad & e^{-\lambda_n^A} \cdot \Bigg\{ q_n e^{-r_n} \cdot \mathrm{N}\left(\frac{-\ln \frac{x}{q_n}}{\sigma_n} - \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right) - x \cdot \mathrm{N}\left(\frac{-\ln \frac{x}{q_n}}{\sigma_n} - \frac{r_n}{\sigma_n} - \frac{\sigma_n}{2}\right) \\
& + x \cdot \sum_{i=1}^{M} A_{i-1}^l \cdot \left(\mathrm{N}\left(\frac{\ln \frac{x}{q_n + x_{i-1}}}{\sigma_n} + \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right) - \mathrm{N}\left(\frac{\ln \frac{x}{q_n + x_i}}{\sigma_n} + \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right)\right) \\
& + e^{-r_n} \sum_{i=1}^{M} (B_{i-1}^l - A_{i-1}^l q_n) \cdot \left(\mathrm{N}\left(\frac{\ln \frac{x}{q_n + x_{i-1}}}{\sigma_n} + \frac{r_n}{\sigma_n} - \frac{\sigma_n}{2}\right) - \mathrm{N}\left(\frac{\ln \frac{x}{q_n + x_i}}{\sigma_n} + \frac{r_n}{\sigma_n} - \frac{\sigma_n}{2}\right)\right) \\
& + l_{M+1,n} \cdot \left(x \cdot \mathrm{N}\left(\frac{\ln \frac{x}{q_n + 1}}{\sigma_n} + \frac{r_n}{\sigma_n} + \frac{\sigma_n}{2}\right) - q_n e^{-r_n} \cdot \mathrm{N}\left(\frac{\ln \frac{x}{q_n + 1}}{\sigma_n} + \frac{r_n}{\sigma_n} - \frac{\sigma_n}{2}\right)\right) \Bigg\} .
\end{aligned}
\tag{G.3}
$$

From the boundary conditions in (5.12), the values of A_n^l and B_n^l for $n = \overline{N}$ are known to be 0. Thus (G.3) can be used to successively evaluate $l_n(x)$ by backward substitution.

For computational work, we use the BSM model with constant values of r_n and σ_n in (5.10) (and (5.11)). Also, the values of x at which (G.3) is evaluated are fixed as the points x_i, $0 \leq i \leq M$ on the grid. This allows us some further computational speed-ups as many quantities in (G.3) can then be computed only once and stored for use in successive iterations.

A similar formula can be obtained for evaluating $g_n(x)$. The method can also be appropriately modified for valuing GWB under the alternate withdrawal strategy of Section 5.3.

Finally, note that because moments of the normal distribution are available in a closed form, in general, we can use any polynomial interpolation method for $l_n(\cdot)$ and $g_n(\cdot)$ and still avoid evaluating any integrals numerically.

G.2 SILN Model

For the SILN model, we use the product dynamics as given by (5.1) (and (5.17) in case of the alternate withdrawal strategy of Section 5.3) and Monte Carlo simulations (with 10000 sample paths) for valuing GWB. We simulate only the randomness due to the market risk factors. Mortality factors are accounted for directly in the spirit of (5.4).

For simulations, we need to generate the risk-free rate r_n and the excess return of the VA fund over r_n for each year. The excess return being log-normal and independent of r_n in the SILN model can be generated in a straightforward way.

The generation of risk-free rate samples, from the two factor Gaussian process is more involved. From the specification of the short-rate model as given in (5.13), it is clear that we have two state variables at any time t, $x_{1,t}$ and $x_{2,t}$. Under the Gaussian model, as we shall show shortly, the state variables $x_{1,T}$, $x_{2,T}$ and the factor $y_T \triangleq \int_0^T r_s ds$, which corresponds to the effective interest rate for the period $(0, T]$ are jointly normal. This then allows us to discretely sample the interest rate process and use a time step as large as the epoch interval in our simulations and speeds up the computations considerably. It is well known (see, for example, Oksendal [98]) that the stochastic differential equation in (5.13) has the solution:

$$
\begin{aligned}
x_{1,t} &= x_{1,0}\exp(-\kappa_1 t) + \sigma_1 \int_0^t \exp(-\kappa_1(t-s))dZ_s^1 \ ; \\
x_{2,t} &= x_{2,0}\exp(-\kappa_2 t) + \sigma_2 \int_0^t \exp(-\kappa_2(t-s))dZ_s^2 \ . \qquad \text{(G.4)}
\end{aligned}
$$

It then follows that,

$$
\begin{aligned}
y_T = \int_0^T r_t dt &= \int_{0^T} x_{1,t}dt + \int_{0^T} x_{2,t}dt + \int_0^T b(t)dt \\
&= x_{1,0}\int_0^T \exp(-\kappa_1 t)dt + \sigma_1 \int_0^T \int_0^t \exp(-\kappa_1(t-s))dZ_s^1 dt \\
&\quad + x_{2,0}\int_0^T \exp(-\kappa_2 t)dt + \sigma_2 \int_0^T \int_0^t \exp(-\kappa_2(t-s))dZ_s^2 dt + B(T) \\
&= x_{1,0}f(\kappa_1, T) + \sigma_1 \int_0^T f(\kappa_1, T-s)dZ_s^1 \\
&\quad + x_{2,0}f(\kappa_2, T) + \sigma_2 \int_0^T f(\kappa_2, T-s)dZ_s^2 + B(T) \ ; \qquad \text{(G.5)}
\end{aligned}
$$

where,

$$
\begin{aligned}
B(T) &= \int_0^T b(t)dt \ ; \\
f(\kappa, t) &\triangleq \frac{1-\exp(-\kappa t)}{\kappa} \ .
\end{aligned}
$$

From (G.4) and (G.5), it follows that $x_{1,T}$, $x_{2,T}$ and y_T should be jointly normal. Further, using Ito isommetry,

$$
\begin{aligned}
\mathrm{var}(x_{1,T}) &= \sigma_1^2 \int_0^t \exp(-2\kappa_1(t-s))ds \\
&= \sigma_1^2 f(2\kappa_1, T) \ ; \\
\mathrm{var}(x_{1,T}) &= \sigma_1^2 f(2\kappa_2, T) \ ; \\
\mathrm{var}(y_T) &= \sigma_1^2 \int_0^T \left(f(\kappa_1, T-s)\right)^2 ds + \sigma_2^2 \int_0^T \left(f(\kappa_2, T-s)\right)^2 ds \\
&\quad + 2\sigma_1\sigma_2\rho \int_0^T f(\kappa_1, T-s) \cdot f(\kappa_2, T-s) ds \\
&= \frac{\sigma_1^2}{\kappa_1^2}\left(T - 2f(\kappa_1, T) + f(2\kappa_1, T)\right) + \frac{\sigma_2^2}{\kappa_2^2}\left(T - 2f(\kappa_2, T) + f(2\kappa_2, T)\right) \\
&\quad + 2\frac{\sigma_1\sigma_2\rho}{\kappa_1\kappa_2}\left(T - f(\kappa_1, T) - f(\kappa_2, T) + f(\kappa_1 + \kappa_2, T)\right) \ .
\end{aligned}
$$

We can similarly find the covariances between the state variables and the effective discount term as

$$
\begin{aligned}
\mathrm{cov}(x_{1,T}, x_{2,T}) &= \sigma_1\sigma_2\rho f(\kappa_1 + \kappa_2, T) \ ; \\
\mathrm{cov}(x_{1,T}, y_T) &= \frac{\sigma_1^2}{\kappa_1}\left(f(\kappa_1, T) - f(2\kappa_1, T)\right) + \frac{\sigma_1\sigma_2\rho}{\kappa_2}\left(f(\kappa_1, T) - f(\kappa_1 + \kappa_2, T)\right) \ ; \\
\mathrm{cov}(x_{2,T}, y_T) &= \frac{\sigma_1\sigma_2\rho}{\kappa_1}\left(f(\kappa_2, T) - f(\kappa_1 + \kappa_2, T)\right) + \frac{\sigma_2^2}{\kappa_2}\left(f(\kappa_2, T) - f(2\kappa_2, T)\right) \ .
\end{aligned}
$$

The expressions in (G.6) and (G.6) can be used to directly generate a sample of the state variable $(x_{1,n+1}, x_{2,n+1})$ and the one year risk-free rate r_{n+1} from $(x_{1,n}, x_{2,n})$.

Finally, the function $b(t)$ is adjusted so that the forward rates implied by the short rate model match the market forward rates. Let $\gamma_0(t)$ denote the time t market forward rate. Then, we must have

$$
\exp\left(-\int_0^t \gamma_0(s)ds\right) = \mathbb{E}^{\mathbb{Q}}[\exp(-y_t)] \ .
$$

Using (G.5) and (G.6) and the fact that y_T is Gaussian, we get

$$
\begin{aligned}
\exp\left(-\int_0^t \gamma_0(s)ds\right) &= \exp\left(-x_{1,0}f(\kappa_1,t) - x_{2,0}f(\kappa_2,t) - B(t) + \frac{1}{2}\text{var}(y_t)\right) \; . \\
\text{Hence, } b(t) &= \gamma_0(t) + \frac{d}{dt}\left(-x_{1,0}f(\kappa_1,t) - x_{2,0}f(\kappa_2,T) + \frac{1}{2}\text{var}(y_t)\right) \\
&= \gamma_0(t) - x_{1,0}\text{exp}(-\kappa_1 t) - x_{2,0}\text{exp}(-\kappa_2 t) \\
&\quad + \frac{1}{2}\left(\sigma_1^2 \cdot (f(\kappa_1,t))^2 + \sigma_2^2 \cdot (f(\kappa_2,t))^2 + 2\sigma_1\sigma_2\rho \cdot f(\kappa_1,t)f(\kappa_2,t)\right) \; .
\end{aligned}
$$

G.3 SISV Model

Unlike the BSM and SILN models, the SISV model does not offer computational short-cuts. Like the SILN model, we use Monte Carlo simulations, again with 10,000 sample paths, to value GWB liabilities and revenue streams. We again simulate only the randomness coming from market factors, as the mortality related randomness can be directly incorporated in the pricing formulae. We discretely sample the interest rate process using the procedure described in Section G.2 for the SILN model.

To generate the excess VA fund returns, we first generate excess equity returns for one year using a discretized version of (5.15) with a small time step $\Delta = \frac{1}{250}$ years. This also gives us a sample of the other state variable in the system, i.e., V_n, or the instantaneous variance at the n^{th} anniversary. The one year excess bond return is obtained by directly sampling from a log-normal distribution with volatility σ_b. The one year excess portfolio return, i.e., $\ln(R_n^s) - r_n$ is then obtained by taking a weighted combination of the excess equity and the excess bond returns, with the weights selected in accordance with the composition chosen by the investor.

Bibliography

[1] http://www.econ.yale.edu/ shiller/data.htm.

[2] D. Aboody. Market valuation of Employee Stock Options. *Journal of Accounting and Economics*, 22:357–391, 1996.

[3] V. V. Acharya, K. John, and R. K. Sundaram. On the optimality of resetting executive stock options. *Journal of Financial Economics*, 57(1):65–101, July 2000.

[4] P. Artzner, F. Delbaen, J. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, July 1999.

[5] P. Artzner, F. Delbaen, J. Eber, D. Heath, and H. Ku. Coherent multiperiod risk adjusted values and Bellmans principle. *Annals of Operations Research*, 152(1):5–22, July 2007.

[6] I. Bajeux-Besnainou and R. Portait. Dynamic Asset Allocation in a Mean-Variance Framework. *Management Science*, 44(11):79–95, November 1998.

[7] G. Bakshi, C. Cao, and Z. Chen. Empirical Performance of Alternative Option Pricing Models. *The Journal of Finance*, 52(5):2003–2049, December 1997.

[8] G. Bakshi, C. Cao, and Z. Chen. Pricing and hedging long-term options. *Journal of Econometrics*, 94:277–318, 2000.

[9] L. Ballotta, G. Esposito, and S. Haberman. Modelling the Fair Value of Annuities Contracts: The Impact of Interest Rate Risk and Mortality Risk . Available at SSRN: http://ssrn.com/abstract=992211, December 2006.

[10] L. Ballotta and S. Haberman. Valuation of guaranteed annuity conversion options. *Insurance: Mathematics and Economics*, 33(1):87–108, August 2003.

[11] S. Basak and G. Chabakauri. Dynamic Mean-Variance Asset Allocation. Available at http://ssrn.com/abstract=965926, March 2008.

[12] S. Basak and A. Shapiro. Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices. *Review of Financial Studies*, 14(2):371–405, Summer 2001.

[13] D. Bauer, A. Kling, and J. Russ. A Universal Pricing Framework for Guaranteed Minimum Benefits in Variable Annuities. Available at: http://www.actuaries.org/AFIR/Colloquia/Stockholm/Bauer.pdf, 2006.

[14] R. E. Bellman and S. E. Dreyfus. *Applied Dynamic Programming* . Princeton University Press, Princeton, 1st edition, 1962.

[15] J. C. Bettis, J. M. Bizjak, and M. L. Lemmon. Exercise behavior, valuation, and the incentive effects of employee stock options. *Journal of Financial Economics*, 76(2):445–470, May 2005.

[16] E. Biffis. Affine processes for dynamic mortality and actuarial valuations. *Insurance: Mathematics and Economics*, 37(3):443–468, December 2005.

[17] E. Biffis and P. Millossovich. The fair value of Guaranteed Annuity Options. *Scandinavian Actuarial Journal*, 1:23–41, 2006.

[18] F. Black and M. Scholes. The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3):637–654, May-June 1973.

[19] K. Boda and J. A. Filar. Time consistent dynamic risk measures. *Mathematical Methods of Operations Research*, 63(1):169–186, January 2006.

[20] Z. Bodie. Common Stocks as a Hedge Against Inflation. *The Journal of Finance*, 31(2):459–470, May 1975.

[21] Z. Bodie, D. Ruffino, and J. Treussard. Contingent Claims Analysis and Life-Cycle Finance. *American Economic Review*, 98(2):291–296, May 2008.

[22] P. P. Boyle and M. Hardy. Guaranteed Annuity Options. *ASTIN Bulletin*, 33(2):125–152, January 2003.

[23] P.P. Boyle and E. Schwartz. Equilibrium prices of guarantees under equity-linked contracts. *Journal of Risk and Insurance*, 44(4):639–680, December 1977.

[24] M.J. Brennan and E. Schwartz. The pricing of equity-linked life insurance policies with an asset value guarantee. *Journal of Financial Economics*, 3(3):195–213, June 1976.

[25] D. Brown. *Risk and Robust Optimization*. PhD thesis, Massachusetts Institute of Technology, June 2006.

[26] G. W. Buetow. Ratchet Options. *Journal of Financial and Strategic Decisions*, 12(2):17–30, Fall 1999.

[27] J. Bulow and J. Shoven. Accounting for Stock Options. *Journal of Economic Perspectives*, 19(4):115–134, Fall 2005.

[28] A.J.G. Cairns, D. Blake, and K. Dowd. A two-factor model for Stochastic Mortaity with Parameter Uncertainty: Theory and Calibration. *The Journal of Risk and Insurance*, 73(4):687–718, 2006.

[29] J. N. Carpenter. The Exercise and Valuation of Executive Stock Options. *Journal of Financial Economics*, 48(2):127–158, 1998.

[30] P. Carr. Hedging Complex Barrier Options. Master's thesis, New York Univeristy, Courant Institute, New York, NY 10012, June 2002.

[31] P. Carr and V. Linetsky. The Valuation of Executive Stock Options in an Intensity- Based Framework. *European Finance Review* , 4(3):211–230, 2000.

[32] D. M. Chance. Expensing Executive Stock Options: Sorting out the Issues. Available at SSRN: http://ssrn.com/abstract=590324 or DOI: 10.2139/ssrn.590324, September 2004.

[33] A. H. Y. Chen, F. C. Jen, and S. Zionts. The Optimal Portfolio Revision Policy. *The Journal of Business*, 44(1):51–61, January 1971.

[34] Z. Chen, K. Vetzal, and P. A. Forsyth. The Effect of Modelling Parameters on the Value of GMWB Guarantees. Forthcoming, Insurance: Mathematics and Economics, December 2007.

[35] P. Cheridito and M. Stadje. Time-Inconsistency of VaR and Time-Consistent Alternatives . Available at SSRN: http://ssrn.com/abstract=1098863, November 2007.

[36] K. C. Cheung and H. Yang. Optimal stopping behavior of equity-linked investment products with regime switching. *Insurance: Mathematics and Economics*, 37(3):599–614, December 2005.

[37] G. Choquet. Theory of Capacities. *Annales de l'Institut Fourier*, 5:131–295, 1953.

[38] C. C. Chu and Y. K. Kwok. Reset and withdrawal rights in dynamic fund protection. *Insurance: Mathematics and Economics*, 34(2):273–295, 2004.

[39] C. C. Chu and Y. K. Kwok. Valuation of guaranteed annuity options in affine term structure models. Available at SSRN: http://ssrn.com/abstract=644224, January 2005.

[40] T. F. Coleman, Y. Kim, Y. Li, and M. Patron. Robustly Hedging Variable Annuities with Guarantees Under Jump and Volatility Risks. *Journal of Risk & Insurance*, 74(2):347–376, June 2007.

[41] T. F. Coleman, Y. Li, and M. Patron. Hedging guarantees in variable annuities under both equity and interest rate risks. *Insurance: Mathematics and Economics*, 38(2):215–228, April 2006.

[42] R. Cont. Model uncertainty and its impact on the pricing of derivative instruments. *Mathematical Finance*, 16(3):519–547, July 2006.

[43] J. Core and W. Guay. Estimating the Value of Employee Stock Option Portfolios and Their Sensitivities to Price and Volatility. *Journal of Accounting Research*, 40(3):613–630, June 2002.

[44] J. C. Cox, J. E. Ingersoll, and S. A. Ross. A Theory of the Term Structure of Interest Rates. *Econometrica*, 53(2):385–407, March 1985.

[45] R. P. Cubitt. Rational Dynamic Choice and Expected Utility Theory. *Oxford Economic Papers*, 48(1):1–19, January 1996.

[46] D. Cuoco, H. He, and S. Issaenko. Optimal Dynamic Trading Strategies with Risk Limits. Available at SSRN: http://ssrn.com/abstract=563901, April 2001.

[47] J. Cvitanic, Z. Wiener, and F. Zapatero. Analytic Pricing of Employee Stock Options. *Review of Financial Studeis*, 21(2):683–724, 2008.

[48] M. Dai, Y. K. Kwok, and J. Zong. Guaranteed Minimum Withdrawal Benefit in Variable Annuities. Available at SSRN: http://ssrn.com/abstract=964083, February 2007.

[49] J. Detemple and S. Sundaresan. Nontraded Asset Valuation with Portfolio Constraints: a Binomial Approach. *Review of Financial Studies*, 12(4):835–872, 1999.

[50] A. A. Dragulescu and V. M. Yakovenko. Probability distribution of returns in the Heston model with stochastic volatility. *Quantitative Finance*, 2:443–453, December 2002.

[51] D. Duffie. *Dynamic Asset pricing Theory*. Princeton University Press, third edition, 2001.

[52] P. H. Dybvig and M. Lowenstein. Employee Reaload Options: Pricing, Hedging and Optimal Exercise. *Review of Financial Studies* , 16(1):145–171, 2003.

[53] L. G. Epstein and M. Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, November 2003.

[54] Financial Accounting Standards Board. Proposed Statement of Financial Accounting Standards. Share-Based Payment: an Amendment of FASB Statements 123 and 95, March 2004. Available at: http://www.fasb.org/draft/ed_intropg_share-based_payment.shtml.

[55] Financial Accounting Standards Board. Statement of Financial Accounting Standards No. 157 - Fair Value Measurements, September 2006. Available at: http://www.fasb.org/pdf/fas157.pdf.

[56] H. Follmer and A. Schied. Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447, 2002.

[57] M. Frittelli and G. Scandolo. Risk Measures And Capital Requirements For Processes. *Mathematical Finance*, 16(4):589–612, October 2006.

[58] D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1st edition, 1991.

[59] P. Gaillardetz and X. S. Lin. Valuation of Equity-Linked Insurance and Annuity Products with Binomial Models. *North American Actuarial Journal*, 10(4):117–144, October 2006.

[60] R. Gibbons. *Game Theory for Applied Economists*. Princeton University Press, 1st edition, 1992.

[61] M. R. Grasselli. Nonlinearity, correlation and the valuation of employee stock options. Available at: arXiv:math/0511234v1, November 2005.

[62] B. J. Hall and K. J. Murphy. Stock Options for Undiversified Executives. *Journal of Accounting and Economics*, 33(1):3–42, February 2002.

[63] B. J. Hall and K. J. Murphy. The Trouble with Stock Options. *Journal of Economic Perspectives*, 17(3):49–70, Summer 2003.

[64] P. J. Hammond. Consequentialist Foundations for Expected Utility. *Theory and Decision*, 25(1):25–78, July 1988.

[65] G. B. Hazen. Does Rolling Back Decision Trees Really Require the Independence Axiom? *Management Science*, 33(6):807–809, June 1987.

[66] T. Hemmer, Matsunaga S., and T. T. Shevlin. Optimal exercise and the cost of granting employee stock options with a reload provision. *Journal of Accounting Research*, 36(2):231–255, Autumn 1998.

[67] V. Henderson. The impact of the market portfolio on the valuation, incentives and optimality of executive stock options. *Quantitative Finance*, 5(1):1–13, February 2005.

[68] S. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6(2):327–343, Summer 1993.

[69] D. Holz, A. Kling, and J. Rub. GMWB For Life: An analysis of Lifelong Withdrawal Guarantees. Presented at American Risk and Insurance Association 2007 Annual Meeting. Available at http://www.aria.org/meetings/2007papers/IE2007.

[70] S. Huddart. Employee Stock Options. *Journal of Accounting and Economics*, 18(2):207–231, September 1994.

[71] S. Huddart and M. Lang. Employee Stock Option Exercises: An Empirical Analysis. *Journal of Accounting and Economics*, 21:5–43, 1996.

[72] J. Hull and A. White. How to Value Employee Stock Options. *Financial Analysts Journal*, 60(1):114–119, January/Februay 2004.

[73] J. E. Ingersoll. The Subjective and Objective Evaluation of Incentive Stock Options. *The Journal of Business*, 79(2):453–487, 2006.

[74] International Accounting Standards Board . International Financial Reporting Standards 2, "'Share-Based Payment", February 2004.

[75] A. Jain and A. Subramanian. The Inter-temporal Exercise and Valuation of Employee Options. *The Accounting Review*, 79(3):705–743, July 2004.

[76] T. H. Johnsen and J. B. Donaldson. The Structure of Intertemporal Preferences under Uncertainty and Time Consistent Plans. *Econometrica*, 53(6):1451–1458, November 1985.

[77] A. Kadam, P. Lakner, and A. Srinivasan. Executive Stock Options: Value to the Executive and Cost to the Firm. Available at SSRN: http://ssrn.com/abstract=353422 or DOI: 10.2139/ssrn.353422, January 2007.

[78] I. Karatzas and S.E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer, 2^{nd} edition, 1991.

[79] E. Karni and D. Schmeidler. Atemporal dynamic consistency and expected utility theory. *Journal of Economic Theory*, 54:401–408, 1991.

[80] M. Kijima and T. Wong. Pricing of Ratchet Equity-Indexed Annuities under Stochastic Interest Rates. *Insurance: Mathematics and Economics*, 41(3):317–338, November 2007.

[81] S. Korkmaz. Dynamic Coherent Risk Measures. Master's thesis, The Middle Esat Technical University, August 2006.

[82] N. Kulatilaka and A. J. Marcus. Valuing Employee Stock Options. *Financial Analysts Journal*, 50(6):46–56, November-December 1994.

[83] R. A. Lambert, D. F. Larcker, and R. E. Verrechia. Portfolio Considerations in Valuing Executive Compensation. *Journal of Accounting Research*, 29(1):129–149, Spring 1991.

[84] J. H. Lee and D. R. Stock. Embedded options and interest rate risk for insurance companies, banks and other financial institutions. *The Quarterly Review of Economics and Finance*, 40(2):169–187, Summer 2000.

[85] M. Leippold, F. Trojani, and P. Vanini. A Geometric Approach To Multiperiod Mean Variance Optimization of Assets and Liabilities. *Journal of Economic Dynamics and Control* , 28(6):387–406, April 2004.

[86] T. Leung and R. Sircar. Accounting for Risk Aversion, Vesting, Job Termination Risk and Multiple Exercises in Valuation of Employee Stock Options. To appear in Mathematical Finance, August 2007.

[87] D. Li and W. Ng. Optimal dynamic portfolio selection: multi-period mean-variance formulation. *Mathematical Finance*, 10(3):387–406, 2000.

[88] G. Lotito. Dynamic inconsistency and different models of dynamic choice - a review. Available online at: http://polis.unipmn.it/pubbl/index.php?paper=1848, December 2006.

[89] M. J. Machina. Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, December 1989.

[90] H. Markowitz. Portfolio Selection. *Journal of Finance*, 7:77–91, 1952.

[91] Mathworld. http://mathworld.wolfram.com/WhittakerDifferentialEquation.html. Available at mathworld.wolfram.com.

[92] E. F. McClennen. *Rationality and Dynamic Choice.* Cambridge University Press., 1st edition, 1990.

[93] M. A. Milevsky and S. E. Posner. The Titanic option: Valuation of the guaranteed minimum death benefit in variable annuities and mutual funds. *Journal of Risk and Insurance*, 68(1):93–128, 2001.

[94] M. A. Milevsky, S. D. Promislow, and V. R. Young. Killing the Law of Large Numbers: Mortality Risk Premiums and th Sharpe Ratio. *The Journal of Risk and Insurance*, 73(4):673–686, 2006.

[95] M. A. Milevsky and T. S. Salisbury. Financial valuation of guaranteed minimum withdrawal benefits. *Insurance: Mathematics and Economics*, 38(1):21–38, 2006.

[96] N. Moodley. The heston model: A practical approach with matlab code. Master's thesis, University of the Witwatersrand, Johannesburg, South Africa, 2005.

[97] B. Mudavanhu and J. Zhuo. Valuing guaranteed minimum death benefits in variable annuities and the option to lapse. http://www.haas.berkeley.edu/MFE/download/student_papers/mfe02_mudavanhu-variable_annuities.pdf, 2007.

[98] B. K. Oksendal. *Stochastic Differential Equations: An Introduction with Applications*. Springer, fifth edition, 1998.

[99] E. Ozdenoren and J. Peck. Ambiguity aversion, games against nature, and dynamic consistency. *Games and Economic Behavior*, 62(1):106–115, January 2008.

[100] A. Pelsser. Pricing and hedging guaranteed annuity options via static option replication. *Insurance: Mathematics and Economics*, 33(2):283–296, 2003.

[101] Pension Benefit Guaranty Corporation (PGBC). Erisa section 4050 mortality rates (2008). http://www.pbgc.gov/practitioners/index.html.

[102] F. Reidel. Dynamic coherent risk measures. *Stochastic Processes and their Applications*, 112(2):185–200, August 2004.

[103] M. R.Hardy and J. L. Wirch. The Iterated CTE: A Dynamic Risk Measure. *North American Actuarial Journal*, 8(4):62–75, October 2004.

[104] L. C. G. Rogers and J. A.. Scheinkman. Optimal exercise of executive stock options. *Finance and Stochastics*, 11(3):357–372, 2007.

[105] B. Roorda and J.M. Schumacher. Time consistency conditions for acceptability measures, with an application to Tail Value at Risk. *Insurance: Mathematics and Economics*, 40(2):209–230, March 2007.

[106] B. Roorda, J.M. Schumacher, and J. Engwerda. Coherent Acceptability Measures in Multiperiod Models. *Mathematical Finance*, 15(4):589–612, October 2005.

[107] R. Sarin and P. P. Wakker. Dynamic choice and nonexpected utility. *Journal of Risk and Uncertainty*, 17(2):87–120, November 1998.

[108] Securities and Exchange Commission. Staff Accounting Bulletin No. 107, March 2005. Available at: http://www.sec.gov/interps/account/sab107.pdf.

[109] R. Sircar and W. Xiong. A general framework for valuating Incentive Options. *Journal of Economic Dynamics & Control*, 31(7):2317–2349, July 2007.

[110] T. K. Siu. Financial fair valuation of participating policies with surrender options and regime switching. *Insurance: Mathematics and Economics*, 37(3):533–552, December 2005.

[111] J. W. Snow. One Thing We Can All Agree On. The Wall Street Journal, January 20, 2005.

[112] R. H. Strotz. Myopia and Inconsistency in Dynamic Utility Maximization. *The Review of Economic Studies*, 23(3):165–180, 1955-1956.

[113] O. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, November 1977.

[114] O. Volij. Dynamic consistency, consequentialism and reduction of compound lotteries. *Econometric Letters*, 46(2):121–129, October 1994.

[115] N. Wang, R. Gerrard, and S. Haberman. The premium and the risk of a life policy in the presence of interest rate fluctuations. *Insurance: Mathematics and Economics*, 35(3):537–551, December 2004.

[116] T. Wang. A Class of Dynamic Risk Measures. Working Paper, September 1999.

[117] S. Weber. Distribution-invariant dynamic risk measures. *Mathematical Finance*, 16(2):419–441, April 2006.

[118] P. Weller. Consistent Intertemporal Decision-Making under Uncertainty. *The Review of Economic Studies*, 45(2):263–266, June 1978.

[119] J. E. Zhang and J. Shu. Pricing S&P 500 Index Options with Hestons Model. In *Proceedings. 2003 IEEE International Conference on Computational Intelligence for Financial Engineering*, pages 85–92, March 2003.