

Cost Efficiency in Portfolio Management, Behavioral Finance and Insurance

by

Carole Bernard

Habilitation à Diriger
Des Recherches
Section 06

Université de Rennes 1

Membres du Jury:

Professeur Olivier L'Haridon	Examineur	Université de Rennes 1
Professeur Elisa Luciano	Rapporteur	University of Torino, Collegio Alberto
Professeur Franck Moraux	Directeur d'habilitation	Université de Rennes 1
Professeur Jean-Luc Prigent	Rapporteur	Université de Cergy Pontoise
Professeur Patrick Roger	Rapporteur & Président	Université de Strasbourg

Abstract

This document builds on the concept of cost-efficiency. The idea is to find the cheapest way to achieve a given distribution of wealth at some investment horizon (Part I). In Part II, we develop two direct applications. The first one belongs to behavioural finance. The expected utility framework is shown to be equivalent to any law-invariant setting encompassing recent behavioural theories such as the cumulative prospect theory. A second application is in insurance and more generally in indifference utility pricing. Limitations of the law-invariant settings are then discussed and more realistic preferences are developed in Part III. We consider investors with state-dependent preferences or subject to state-dependent constraints (e.g., coming from benchmarks used in performance and risk assessments). We start with applications to mean variance portfolio selection and to optimal portfolio choice under constraints. A general framework of representation of state-dependent preferences is then presented. Conclusions and recent research progress are discussed in the last chapter.

Acknowledgments

Les travaux de recherche présentés dans ce mémoire ont été majoritairement réalisés à l'université de Waterloo au Canada et ont bénéficié des conditions excellentes pour la recherche qui y sont offertes. En particulier, je souhaite remercier mes étudiants et mes collaborateurs de recherche qui ont participé indirectement à la réalisation de ce document.

Je tiens à remercier Professeur Franck Moraux, qui a accepté de diriger cette habilitation à diriger des recherches et qui m'a soutenu dans cette démarche. Il a su, malgré son emploi du temps très chargé, dégagé du temps afin d'améliorer ma présentation en apportant le recul et les ouvertures nécessaires. Je remercie tout particulièrement les professeurs Elisa Luciano, Jean-Luc Prigent et Patrick Roger pour avoir accepté le rôle de rapporteurs. Je remercie également Olivier L'Haridon d'avoir accepté de faire partie de ce jury. Enfin, je remercie chacun pour le temps qu'il a consacré à la lecture et l'analyse du manuscrit et à l'éventuel déplacement pour participer à la soutenance.

Pendant les six dernières années, Professeur Steven Vanduffel a été mon collaborateur principal. Je lui suis tout particulièrement reconnaissante pour son dynamisme, son enthousiasme, sa passion, son amitié et soutien dans les moments difficiles mais aussi les joies des succès partagés autour de nos recherches communes.

Ma gratitude va aussi à Grenoble Ecole de Management qui m'a intégrée en professeur depuis Janvier 2015 et me soutient dans cette démarche d'habilitation à diriger des recherches. Merci à mes nouveaux collègues pour leur accueil et leurs encouragements.

Je termine en remerciant ma famille, mes amis et mes proches qui me supportent au quotidien. Ils savent à la fois m'encourager et me ramener à la réalité quand mon travail de recherche commence à prendre trop d'importance.

Table of Contents

I	Cost-efficiency	5
1	Cost-efficiency	6
1.1	Assumptions and Definitions	9
1.2	Cost-Efficient Payoffs	10
1.2.1	Characterization	10
1.2.2	Black Scholes Market	13
1.3	Agents with State-Independent Preferences	14
1.4	Examples	15
1.4.1	Forward Contract	15
1.4.2	Put Options	17
1.4.3	Path-Dependent Example: Lookback Option	18
1.5	Cost-efficiency with State-Dependent Constraints	20
1.5.1	Constrained Cost-Efficient Payoffs	21
1.6	Summary and Conclusion	23
II	Law-invariance	25
2	Rationalizing Investors' Choices	26
2.1	Introductory Example	28
2.2	Setting	30
2.3	Explaining Distributions through Expected Utility Theory	32
2.3.1	Standard Expected Utility Maximization	32
2.3.2	Generalized Expected Utility Maximization	35
2.4	From Distributions to Utility Functions	36
2.4.1	Explaining the Demand for Capital Guarantee Products	37
2.4.2	Yaari's Dual Theory of Choice Model	38

2.5	From Distributions to Risk Aversion	39
2.5.1	Risk Aversion Coefficient	39
2.5.2	Decreasing Absolute Risk Aversion	40
2.5.3	Case of a Black-Scholes Market	41
2.6	Distributions and Corresponding Utility Functions	42
2.6.1	Normal Distribution and Exponential Utility	42
2.6.2	Lognormal Distribution and CRRA and HARA Utilities	43
2.6.3	Exponential Distribution	43
2.6.4	Pareto Distribution	44
2.7	Conclusions	44
3	Financial Bounds for Insurance Claims	45
3.1	Setting	47
3.1.1	Agents' Preferences	47
3.1.2	Insurance Contract	49
3.1.3	Financial Market	49
3.2	Traditional Indifference Pricing Approach	50
3.2.1	Indifference Pricing	50
3.2.2	Classical lower bound	50
3.2.3	Comparison between bid and ask prices	51
3.3	Market-based Indifference Pricing Approach	53
3.3.1	General indifference principle	53
3.3.2	Lower bounds on insurance claims	54
3.3.3	Upper bound	57
3.4	Example in the Black and Scholes model	57
3.5	Conclusions	58
III	State-Dependent Preferences	59
4	Mean-Variance Optimal Portfolios with Applications to Fraud Detection	60
4.1	Market Setting	62
4.2	Unconstrained Mean-Variance Optimal Portfolios	63
4.2.1	Mean-Variance Efficiency	63
4.2.2	Example in the Black-Scholes Setting	65

4.2.3	Maximum Sharpe Ratio and Application to Fraud Detection	65
4.3	Mean-Variance Efficiency with a Correlation Constraint	67
4.3.1	Mean-Variance Efficiency	68
4.3.2	Maximum Sharpe Ratio and Application to Fraud Detection	69
4.4	Mean-Variance Optimal Portfolios with a Dependence Constraint	71
4.5	Final Remarks	75
5	Optimal Portfolios under Worst-Case Scenarios	76
5.1	Setting	79
5.1.1	Financial Market	79
5.1.2	Growth Optimal Portfolio (GOP) and Pricing	80
5.1.3	Market Crises	81
5.2	Traditional Diversification Strategies	81
5.2.1	Buy-and-Hold Strategies	81
5.2.2	Constant-Mix Strategies	82
5.2.3	Performance of Investment Strategies during a Crisis	82
5.3	Optimal Tail Diversification	83
5.3.1	Optimal Strategies with Constraints in the Tail	84
5.3.2	Independence in the Tail	85
5.4	Applications	86
5.5	Conclusions	89
6	Optimal Payoffs under State-dependent Preferences	91
6.1	Framework and Notation	93
6.2	Law-invariant Preferences and Optimality of Path-independent Payoffs	94
6.3	Optimal payoffs under state-dependent preferences.	94
6.3.1	Sufficiency of Twins	95
6.3.2	Optimality of Twins	96
6.4	Improving Security Design	97
6.5	Conclusions	99
7	Conclusions and Research Directions	101
	References	106

Introduction

After my PhD studies at ISFA, University of Lyon 1, I have spent nine years at the University of Waterloo in Canada (2006-2014). During these years I have been working in four areas of research: in insurance (optimal insurance design and equity-linked insurance), on financial derivatives pricing and hedging (quantitative finance), on risk management and dependence modelling, and on the theory of cost-efficiency and its applications (optimal design of retail investment products, optimal portfolio selection, mathematical economics and behavioural finance). These four research areas were explored in close collaboration with students at the PhD level (M. Ghossoub, Z. Cui, A. MacKay, J.S. Chen, E. Brechmann, M. Maj, J. Tang), at the master level (W. Gornall, Z. Cui, M. Mülbeyer, X. Jiang) and also at the undergraduate level (Y. Liu, N. McGillivray, J. Zhang, X. Jiang), with colleagues at Waterloo (W. Tian, P. Boyle, M. Hardy, C. Lemieux, A. Kolkiewicz, D. McLeish, R. Wang) as well as with other internationally recognized researchers (M. Augustyniak, O. Bondarenko, A. Chen, C. Czado, M. Denuit, X. He, M. Kwak, O. Le Courtois, M. Ludkovski, F. Moraux, F. Quittard-Pinon, W. Li, L. Rüschendorf, S. Vanduffel, J.A. Yan, X. Y. Zhou. I provide more details on each of these four topics in the introduction hereafter, but the remainder of the document presents the fourth area of research in full detail.

The fourth research direction does not built on work done during my PhD. I consider this part as my most significant research contribution during my time at Waterloo in terms of developing a new theory rather than extending existing directions that were started by myself or by other researchers. This is also the part for which I faced the most challenging criticisms but also for which I received research awards (2011 Young Economist best paper award at the 2011 EGRIE conference, the 2012 Johan de Witt award, 2013 Brains Back to Brussel research funding, and the 2015 Redington prize). My most cited paper with my first PhD student Mario Ghossoub deals with cumulative prospect theory, which is directly related to this fourth research axis in behavioural finance. In this introduction, I now start with an overview of my research by briefly describing my four areas of research.

The first area of research is insurance. I started to work on optimal insurance design during my postdoctoral fellowship in collaboration with W. Tian. The original motivation was to model the insurance and reinsurance demand and to understand the impact of changes in the insurance regulation. For example, we studied the implications of Solvency II requirements on the optimal risk sharing between insurers and reinsurers. The role of regulators is to protect policyholders and market stability. However, we show that regulation can have adverse effects in the reinsurance market, by inducing companies to leave uninsured the highest risks (Bernard and Tian (2009, 2010)). I also worked on the impact of counterparty risk on the optimal risk sharing in the reinsurance market (Bernard and Ludkovski (2012)). This work was pursued with M. Ludkovski (University of California in Santa Barbara) with financial support of the Society of Actuaries (SOA). Optimal insurance design was also extensively studied by my PhD student Mario Ghossoub in his doctoral

dissertation (Ghossoub (2011)). I also published an extension with X. He (Columbia university, US), J.-A. Yan (Chinese Academy, Beijing, China) and X. Zhou (Oxford, UK) on optimal insurance design in a non-expected utility framework, the rank dependent utility theory (Bernard, He, Yan, and Zhou (2015)) and an extension to the presence of ambiguity in the reinsurance market (Bernard, Ji, and Tian (2013)). My work in this area was noticed by G. Dionne who invited me to contribute to a new edition of the “Handbook on Insurance” with a review chapter on reinsurance (Bernard (2013)). I have also extended some work on insurance of bank deposits (Bernard, Le Courtois, and Quittard-Pinon (2005c)) or on equity-linked insurance products that I started during my doctoral studies (Bernard, Le Courtois, and Quittard-Pinon (2005a, 2006, 2010)) and that I have pursued since then (Bernard and Le Courtois (2012a), Bernard and Lemieux (2008), Bernard and Chen (2009)). Specifically, I have worked on the policyholder behavioural risk and surrender risk in equity-linked insurance products with my third Ph.D. student A. MacKay (Bernard, MacKay, and Muehlbeyer (2014); Bernard, Hardy, and MacKay (2014); Bernard and MacKay (2015); MacKay, Augustyniak, Bernard, and Hardy (2016), MacKay (2014)). I still pursue research in this direction with my former postdoctoral fellow, Minsuk Kwak, (Bernard and Kwak (2016b)) as well as with my current PhD student, Junsen Tang, (Bernard, Kolkiewicz, and Tang (2016)).

My second area of research is in mathematical finance, precisely on the pricing and hedging of financial derivatives. I already contributed in this area in my PhD thesis (Bernard, Le Courtois and Quittard-Pinon (2005b; 2008)). In Waterloo, I extended my earlier work on Parisian options with P. Boyle (Bernard and Boyle (2011a)), with O. Le Courtois (Bernard and Le Courtois (2012b)) and also with D. McLeish and our joint PhD student Zhenyu Cui (Bernard, Cui, and McLeish (2012), Cui (2013)). I worked on options’ strategies and particularly on the Madoff’s scandal with P. Boyle (Bernard and Boyle (2009)). I also developed in collaboration with my master and then PhD student, Z. Cui, new methodologies for pricing and hedging volatility derivatives (Bernard and Cui (2011, 2014), Bernard, Cui, and McLeish (2014, 2016)). Among others I have been involved in numerous projects related to the pricing and hedging of financial products, with C. Czado (Germany) on multivariate option pricing (Bernard and Czado (2013)), with W. Li (Delaware, US) on cliquet-options (Bernard and Li (2013)), with P. Boyle on equity-indexed annuities (Bernard and Boyle (2011b)).

As a third research area, I have recently worked on dependence modelling and some applications in risk management. Specifically, I have worked on how to quantify model risk on Value-at-Risk (VaR) assessment of a portfolio of risks in the presence of uncertainty on the dependence among the risks (Bernard, Jiang, and Wang (2014)). The *Rearrangement Algorithm* of Embrechts, Puccetti, and Rüschendorf (2013) makes it possible to obtain approximations for sharp VaR upper and lower bounds of portfolios in which the marginal distributions are known but not their interdependence. It builds on the concept of mixability developed by Wang (Wang and Wang (2011, 2016)). The gap between the upper and lower VaR bound is typically very high and can only be reduced by using dependence information. We extend the Rearrangement Algorithm to find bounds in the case when besides knowledge of the marginal distribution also the joint distribution is known on a subset of the support (Bernard and Vanduffel (2015)). In collaboration with Rüschendorf and Vanduffel, we solve the problem when besides knowledge of the marginal distributions, also an upper bound on the variance of the portfolio sum is known (Bernard, Rüschendorf, and Vanduffel (2016)). The problem in which information on higher order moments is available is studied in Bernard, Denuit, and Vanduffel (2015). With D. McLeish, we propose an improved version of the rearrangement algorithm that is instrumental in solving the above problems (Bernard and McLeish (2016)). Applications to the industry and to the computations of capital requirements are numerous. The first paper (Bernard and Vanduffel (2015)) was awarded the 2014 PRMIA award

for frontiers in risk management and was applied also in Bernard, Rüschendorf, Vanduffel, and Yao (2016) in assessing model risk in credit risk models used in the industry.

The fourth research area is related to the theory of cost-efficiency and its applications to behavioral finance and optimal decision making. My first published contribution in behavioural finance is on optimal portfolio choice in Cumulative Prospect Theory (CPT) (Bernard and Ghossoub (2010)). CPT is a relatively new setting to model preferences (Tversky and Kahneman (1992)). One of its properties is that its objective function to maximize depends only on the distribution of final wealth, similarly as many standard settings (expected utility, Yaari's theory, VaR minimization, mean-variance optimization...). More recently, I have been working on general properties of optimal portfolio selection in the more general setting of "law-invariant" preferences (Bernard, Boyle, and Vanduffel (2014)). This project was initiated in Bernard, Boyle, and Gornall (2009) using empirical data on structured products traded in the US market. The first goal was to attempt to explain the demand for these products but our study proved that retail investment products are too complex, overpriced and not designed to best fit customers' needs. The SEC (US regulator) contacted us about this study to obtain additional information, as they were investigating abusive design of retail investment products. I extended this specific study with S. Vanduffel (Bernard, Boyle, and Vanduffel (2014), Bernard, Maj, and Vanduffel (2011)) to work on optimal portfolio selection and optimal product design for law-invariant preferences. The paper (Bernard, Boyle, and Vanduffel (2014)) now published in *Finance* won the 2012 Johann de Witt prize from the Dutch actuarial society that recognizes the best working paper of the year in actuarial science. A lot of the work with S. Vanduffel was started during an extended visit in Brussels after a successful application in a competitive award offered by the Belgian government (called "*Brains Back to Brussels*") on "optimal design of retail investment products". During this visit, we developed a new tool combining "copulas" and "cost-efficiency". It raised a lot of interest and I have been invited to several seminars and conferences to present this work. For example I was plenary speaker in Ann Arbor in May 2011 (SAFI) and invited speaker in Brussels (AfMath, Feb. 2011), as well as at the CMS meeting in June 2011. I also presented this work and its applications in insurance (Bernard and Vanduffel (2014a)) in many seminars. In particular, it won the 2011 EGRIE young economist award. In November 2011, I was invited to give a mini-course on this topic at RIO 2011. Since then, I have been developing many implications of the theory of cost-efficiency and each subsequent chapter deals with this theory and some of its applications.

The document is organized in three parts. In *Part I* (Chapter 1), we start by explaining what cost-efficiency is and its strong link with law invariant preferences. Specifically, we will give some intuition about this notion and we illustrate it with examples. We then show how to extend the notion of cost-efficiency to *constrained* cost-efficiency. Doing so makes it possible to study the optimal portfolio choice of investors with state-dependent preferences. This is based on a publication in *Finance* in 2014. But a preliminary working paper version of the paper was already presented at the 2010 December meeting of the French Finance Association (AFFI). In *Part II*, we develop direct applications of the theory of cost-efficiency in law-invariant settings. First, it is used to show the limitations of the expected utility setting and more generally any law-invariant decision setting (Chapter 2 published in *Journal of Mathematical Economics*). We also show how to infer the utility function of investors in a fully non-parametric way from their desired distribution of terminal wealth. We develop an application to estimate insurance prices in Chapter 3 (chapter published in *Journal of Risk and Insurance* in 2014). In *Part III*, we show how constrained cost-efficiency can be useful in portfolio management. First, we discuss mean-variance optimization with a correlation constraint (Chapter 4 published in *European Journal of Operational Research* in 2014). Next, we study best strategies that offer diversification in the worst states of the market (Chapter 5 published

in *Quantitative Finance* in 2014 and that received the 2015 Redington prize from the Society of Actuaries (SOA) in North America). In Chapter 6, we formalize state-dependent preferences and the properties of optimal portfolios for investors with state-dependent preferences. This chapter was published in *Quantitative Finance* in 2015. Finally, we conclude and offer a list of potential research directions in Chapter 7.

Part I

Cost-efficiency

Chapter 1

Cost-efficiency

In this chapter, we introduce the theory of cost-efficiency. One of the main contributions is to provide an explicit representation of the lowest cost strategy to achieve a given payoff distribution (that we call “cost-efficient” strategy). For any inefficient strategy, we are able to construct financial derivatives which dominate in the sense of first-order or second-order stochastic dominance. We highlight the connections between cost-efficiency and dependence. This allows us to extend the theory to deal with state-dependent constraints to better reflect real-world preferences. We show in particular that path-dependent strategies (although inefficient in the Black Scholes setting) may become optimal in the presence of state-dependent constraints. Most proofs are omitted and can be found in the published version of this chapter entitled “Explicit Representation of Cost-Efficient Strategies” and published in *Finance* (Bernard, Boyle, and Vanduffel (2014)).

It is common in the academic literature to study optimal investment strategies using an objective function that balances “expected return” and “risk”. This idea was introduced by Markowitz (1952), who together with Roy (1952) first proposed a quantitative approach to determine the optimal trade-off between mean (expected return) and variance (risk). The mean-variance framework has become very influential due to the fact that it combines simplicity with practical applicability. Markowitz’s (1952) paper initiated a tremendous amount of further academic research on optimal portfolio selection problems. Most of the existing approaches have in common that the objective function being optimized is “law-invariant¹” or “state-independent”. In other words the objective is solely driven by the distribution of returns and does not depend on the states of the economy in which they are generated. Preferences that are not law invariant are then called state-dependent.

We illustrate here law-invariant preferences by a non-exhaustive list of most common examples of models used in modelling behaviours in Table 1.0.1. State-dependent preferences appear then naturally when the agent cares about the states of the world in which she receives cash-flows. For example, state-dependent preferences can be modeled as the optimization of a law-invariant objective with a constraint that depends on the state in which the terminal wealth is. In optimal portfolio choice, it is relatively common to evaluate the performance of a portfolio manager relative to a benchmark. For instance, the portfolio manager can then maximize the expected utility of final wealth under the constraint that the probability of her portfolio to be above a given benchmark is larger than a minimum probability. This means that the portfolio manager will not

¹This is the case for a wide range of behavioral theories, among them expected utility theory (von Neuman and Morgenstern, (1947)), cumulative prospect theory (Tversky and Kahneman, 1992), and rank dependent utility theory (Quiggin, 1993). Law-invariant utilities are also recently discussed by Carlier and Dana (2011).

only be interested in the distribution of his portfolio and thus that her preferences become “state-dependent”. Another example is the case of the demand for say fire insurance by an expected utility maximizer. Her objective function is $U(W - X + I - P(I))$ where W is her final wealth, and the loss X is subtracted (in case of a fire) and she received an indemnity payment I from the insurer. To receive this insurance coverage, a premium $P(I)$ is paid upfront at time 0 (we neglect the interest rate effect). The objective function of the policyholder is to look for the best indemnity I and thus to optimize

$$\max_I E[U(W - X + I - P(I))]. \quad (1.1)$$

If there are two states ω_1 (there is a fire and there is a loss $X(\omega_1) = x > 0$) and ω_2 (there is no fire and the loss $X(\omega_2) = 0$). Then given some assumptions on the premium, one can then solve $y_1 := I(\omega_1)$ and $y_2 := I(\omega_2)$ explicitly. For example, when the utility function of the policyholder is concave and the premium $P(I) = (1 + \alpha)E[I]$ with $\alpha > 0$, then the optimal decision will be to choose to receive 0 in state ω_2 and $y_1 > 0$ in the state when there is a fire. The objective optimized in (1.1) is state-dependent, define $V_1(I) := E[U(W - X + I - P(I))]$. To evaluate it, one needs the joint distribution between X and I . Note that it would make no sense to find the optimal indemnity I by optimizing directly the expected utility of the indemnity I such as $V_2(I) = E[U(I - P(I))]$. Here $V_2(I)$ only depends on the probability distribution of I . Assume that the probability of having a fire is $1/2$ then the objective function $V_2(I)$ will be the same whether the indemnity is paid in state ω_1 when there is a fire or in state ω_2 when there is no a fire. However, it is intuitive that a policyholder will only value the insurance payment if it compensates for an actual loss. Therefore, the policyholder cares about the state in which cashflows are received and the objective function is called “state-dependent.”

Law-invariant preferences	Reference
Expected Utility Theory	von Neumann and Morgenstern (1947)
Mean-variance optimization	Markowitz (1952)
Rank Dependent Utility	Quiggin (1993)
Yaari Theory	Yaari (1987)
Cumulative Prospect Theory	Tversky and Kahneman (1992)
SP/A Theory	Lopes (1987); Lopes and Oden (1999)
Loss Aversion	Berkelaar, Kouwenberg, and Post (2004)
Minimization of a law-invariant risk measure Ex: Value-at-Risk, Expected Shortfall...	Artzner, Delbaen, Eber, and Heath (1999) Jorion (2007)...

Table 1.0.1: Examples of Law-invariant Settings

In this chapter, we first focus on law-invariant preferences. Investors are thus only interested in the distribution of final wealth and not about the states in which outcomes are received. In the first part of the chapter, we assume that an agent specifies a desired distribution² of terminal wealth up-front. Under fairly general assumptions we derive a formula for the unique investment portfolio that minimizes the cost of achieving this distribution (“cost-efficient” portfolio). It is then clear that investors with increasing state-independent preferences (for example, expected utility maximizers) will only consider cost-efficient strategies. We exploit the connection between efficiency and stochastic dominance to develop new insights. Two approaches to improve an inefficient financial strategy are proposed. Any specified inefficient strategy can be improved by either condi-

²Goldstein, Johnson and Sharpe (2008) propose an approach to construct the desired probability distribution of wealth of a given individual at maturity.

tioning (in the sense of second-order stochastic dominance), or by giving the explicit payoff of the unique efficient financial derivative that strictly dominates this strategy (in the sense of first-order stochastic dominance). In a Black Scholes market the strictly dominating strategies are shown to be path-independent and increasing in the risky asset. This feature of perfect correlation implies that optimal strategies incur their worst outcomes in bear markets (for example during a financial crisis). This arguably does not suit most investors who desire protection against downward shocks or, more generally, who also care about the way their investments interact with economic conditions. This idea is reinforced in the context of an increasing concern to evaluate systemic risk (Acharya, (2009)). The basic principle here is that, while the risk of a failure of one business can never be totally eliminated, simultaneous failures should be avoided. The performance of a portfolio should then not be assessed solely by its final distribution but also by its interaction with the financial market which then gives rise to state-dependent preferences.

Hence, the second part of the chapter extends the cost-efficiency framework to include state-dependent preferences. We develop a constructive method to obtain the cheapest investment strategy under state-dependent constraints, where each constraint is interpreted as information on the relationship between the outcomes of the strategy and the states of the financial market. We discuss several cases and in particular observe that in this setting path-dependent payoffs can be optimal even in a Black Scholes framework.

Our work builds on the work of Cox and Leland (2000)³ and Dybvig (1988a; 1988b). We briefly summarize their main results. In the Black Scholes framework, Cox and Leland (2000) provide necessary and sufficient conditions for a dynamic investment strategy to satisfy each of the following properties: (i) the strategy is self financing, (ii) the strategy yields path-independent returns, and (iii) the strategy is consistent with the optimal behavior of some expected utility maximizer. Dybvig (1988a; 1988b) studies optimal investment strategies mainly in the context of discrete markets. In the case of an arbitrage-free and complete market, Dybvig characterizes the efficient payoffs for rational agents with increasing preferences and who care only about the distribution of terminal wealth. In particular, Dybvig's study leads to suboptimality of popular dynamic investment strategies such as the stop loss strategy, CPPI strategies (Bertrand and Prigent (2001))...

Apart from the papers of Cox and Leland (2000) and Dybvig (1988a; 1988b), the literature on this topic is limited. Amin and Kat (2003) in an empirical paper conclude that investing in hedge funds is inefficient. Their measure of efficiency is based on the distributional price defined by Dybvig. These results have been obtained in a frictionless market. Frictions such as transaction costs have been studied by Pelsser and Vorst (1996) and Jouini and Khallal (2001) among others. Pelsser and Vorst (1996) show that certain inefficient strategies may become efficient in the presence of transaction costs in a complete financial market. See also Dumas and Luciano (1991). Jouini and Khallal (2001) quantify the inefficiency of strategies in an incomplete financial market and in a multiperiod framework, and show that inefficient strategies may be optimal in the presence of frictions.

Dybvig (1988a; 1988b) shows that a cost-efficient strategy with distribution F puts more weight on the least expensive states while still maintaining the same distribution F . In other words, the cheaper the state the higher the outcome of the strategy. The same intuition underlies Basak and Shapiro's (2001) paper where an agent maximizes expected utility subject to a Value-at-Risk constraint. When the agent has to satisfy the VaR constraint it may be optimal for her to incur

³The original Cox and Leland paper was written in Cox and Leland (1982) but not published until 2000 in Cox and Leland (2000).

losses in those states which are most expensive to insure. The state-price density then plays a critical role in ranking payoffs in terms of their cost and characterizes optimal strategies when agents have state-independent preferences.

In practice, agents face state-dependent constraints when making investment decisions. Purchasing a put option may mean that the agent would like to protect an investment in the underlying asset. Investors may want to outperform a benchmark or to maintain some correlation with some specific market indices. In this chapter we characterize cheapest possible strategies under state-dependent constraints. The results in the present chapter require that the market is frictionless and arbitrage-free. We make the following main contributions.

First, we connect the notion of “cost-efficiency” with the notion of extreme dependence and anti-monotonicity. This allows us to study cost-efficiency in a unified way. In particular we provide a sufficient condition for a payoff to be cost-efficient (Prop.1.2.1) which allows to include all the different cases discussed in Dybvig (1988a; 1988b).

Second, we discuss that cost-efficient payoffs do not offer protection against a decline in the economy and we introduce state-dependent constraints allowing the investors to specify the interaction between their strategy and the economic conditions. We are able to characterize cost-efficient strategies in the presence of state-dependent preferences and we provide explicit constructions.

Third, we highlight the suboptimality of inefficient strategies for agents with a fixed investment horizon and state-independent increasing preferences⁴ through the explicit construction of *two* preferred strategies and relate these to stochastic dominance rankings. The inefficient strategy is dominated in the sense of first-order stochastic dominance by the optimal cost-efficient strategy which is cheaper and has the same distribution. In this chapter we show it can also be dominated in the sense of second-order stochastic dominance with a strategy that has the same cost but a different distribution. Both approaches are of interest and lead to different strategies. We compare them using a lookback option as an example.

The rest of the chapter is organized as follows. Section 1.1 describes the framework and discusses the main assumptions. Section 1.2 unifies existing results and presents new work on cost-efficiency and optimal cost-efficient payoffs in a preference-free framework. We derive the implications for agents’ optimal investment strategies in Section 1.3. Section 1.4 illustrates our propositions using numerical examples based on some well-known strategies and payoffs. Section 1.5 introduces the concept of constrained cost-efficient payoffs, provides a general characterization result and discusses some applications. The final section summarizes the chapter.

1.1 Assumptions and Definitions

This section is devoted to our main assumptions and definitions. We assume that the financial market is free of arbitrage, perfectly liquid and frictionless (no transaction costs, no trading constraints). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the corresponding probability space. Under these assumptions, there exists a state-price process $(\xi_t)_t$ such that $(\xi_t S_t)_t$ is a martingale for all traded assets S in this market. In an incomplete market, the state-price process is not unique. We assume that the choice of the state-price process is fixed for the decision maker. We first define the cost of a payoff.

⁴A functional f defined on random variables is said to be state-independent or law-invariant if for two random variables X and Y with the same distribution, then $f(X) = f(Y)$.

Definition 1.1.1. The cost of a strategy (or of a financial investment contract) with terminal payoff X_T is given by

$$c(X_T) = E[\xi_T X_T],$$

where the expectation is taken under the physical measure \mathbb{P} .

In this chapter it is always tacitly assumed that all payoffs and state-prices are square integrable. This is sufficient to ensure that all the expectations we write down exist. In particular $c(X_T) < +\infty$. Let F denote the cumulative distribution function (cdf) of a financial payoff. Its inverse F^{-1} is given by

$$F^{-1}(y) = \inf \{x \mid F(x) \geq y\}.$$

Next we define cost-efficiency and the distributional price (first defined by Dybvig (1988a)). These concepts will be extensively used in the rest of the document.

Definition 1.1.2. A strategy (or a payoff) is cost-efficient if any other strategy that generates the same distribution costs at least as much.

Definition 1.1.3 (Dybvig (1988a)). The “*distributional price*” of a cdf F is defined as

$$P_D(F) = \min_{\{Y_T \mid Y_T \sim F\}} c(Y_T) \quad (1.2)$$

where $\{Y_T \mid Y_T \sim F\}$ denotes the set of payoffs that have the same distribution as F . The “*efficiency loss*” of a strategy with payoff X_T at maturity T with cdf F is equal to $c(X_T) - P_D(F)$.

Note that the cost c and the distributional price P_D both depend on the state-price process $(\xi_t)_t$. However for ease of exposition, we omit this dependence in the notation.

1.2 Cost-Efficient Payoffs

Dybvig (1988a; 1988b) discussed how to construct cost-efficient payoffs in a complete market over either finitely many equally probable states and over a continuum of states (using heuristic arguments). We provide a unifying approach (Prop.1.2.1) that allows to cover all the cases in Dybvig (1988a; 1988b). This approach is extended in Section 1.5 in a natural way to include state-dependent constraints. We also give new expressions of cost-efficient strategies in the Black Scholes Market (Prop.1.2.5) and show these exhibit path-independence.

1.2.1 Characterization

It turns out that the construction of cost-efficient payoffs is intimately connected to the concept of anti-monotonicity⁵. Loosely speaking anti-monotonicity for ξ_T and X_T means that the higher the outcome for ξ_T the lower the outcome for X_T . An important property of anti-monotonic pairs is that they exhibit the lowest possible correlation coefficient. More precisely when the distributions for both X_T and ξ_T are fixed, we have

$$(X_T, \xi_T) \text{ is anti-monotonic} \Rightarrow \text{corr}[X_T, \xi_T] \text{ is minimal.} \quad (1.3)$$

⁵A more formal treatment of anti-monotonicity is provided in the Appendix of the published version in *Finance*.

This property (1.3) is key in constructing cost-efficient payoffs. Recall that the cost of a strategy is given by the expected value of the product, i.e. $E[\xi_T X_T]$. Therefore minimizing the cost of a strategy X_T with a given distribution F amounts to minimizing the correlation between ξ_T and X_T ⁶. The next proposition is now intuitive.

Proposition 1.2.1 (Sufficient condition for cost-efficiency). *Any random payoff X_T with the property that (X_T, ξ_T) is anti-monotonic is cost-efficient.*

Proof. Let (X_T, ξ_T) be an anti-monotonic pair with fixed marginal distributions. Then property (1.3) states that $\text{corr}[\xi_T, X_T]$ and thus $E[\xi_T X_T]$ is minimal. Hence X_T is a cost-efficient payoff. \square

Note that Prop.1.2.1 is valid in discrete and continuous-time markets without additional assumptions and is not restricted to non-negative payoffs⁷. Moreover this proposition also suggests a way to construct cost-efficient payoffs distributed with a given cdf F . Intuitively all one needs to do is to reshuffle the outcomes inherent in F such that these are reversely ordered with the realizations of ξ_T , where some care is necessary to ensure that the distribution F is preserved effectively. The following two corollaries put this idea into practice in two important market settings. Their proofs follow directly from Prop.1.2.1.

Corollary 1.2.2. *Let Ω be finite. Assume all states $\omega \in \Omega$ are equally probable and that all state-prices $\xi_T(\omega)$ are strictly different. Define*

$$Y_T^* = F^{-1} \left(1 - F_{\xi_T}^-(\xi_T) \right), \quad (1.4)$$

where $F_{\xi_T}^-(\cdot)$ is the left limit of $F_{\xi_T}(\cdot)$. Then, Y_T^ is the cheapest way to achieve the distribution F . It is also unique.*

Note that $F_{\xi_T}^-(\xi_T)$ assumes the values k/n ($k = 0, 2, \dots, n-1$), all with equal probability. This property ensures that Y_T^* has the right distribution F . This corollary first appeared in Theorem 3 of Dybvig (1988a) in a slightly different form. We now provide two observations to show that the assumptions in Corollary 1.2.2 cannot be readily relaxed and that the condition in Prop.1.2.1 is indeed sufficient but not necessary in general. (i) If in Corollary 1.2.2 some state-prices are equal then Y_T^* is not in general the unique cost-efficient payoff with distribution F . Indeed, let $\xi_T(\omega_1) = \xi_T(\omega_2)$ and $Y_T^*(\omega_1) \neq Y_T^*(\omega_2)$. Define another strategy Y_T such that $Y_T(\omega_1) = Y_T^*(\omega_2)$, $Y_T(\omega_2) = Y_T^*(\omega_1)$ and finally $Y_T(\omega) = Y_T^*(\omega)$ for all other states $\omega \in \Omega$. Then, Y_T is also a cost-efficient strategy with distribution F . (ii) If in Corollary 1.2.2 not all state-prices are equally probable then cost-efficient strategies are no longer anti-monotonic with the state-price ξ_T in general. For example assume that there exists a state ω_0 with probability $\mathbb{P}(\omega_0)$ strictly different from all other state probabilities, and let $\mathbb{P}(\omega_0)$ be assigned to the highest desired outcome. It is then clear that whenever $\xi_T(\omega_0)$ is not the smallest state-price available, the cost-efficient strategy with the desired distribution cannot be anti-monotonic with the state-price ξ_T . The same example also shows why the condition in Prop.1.2.1 is sufficient but not necessary. In a continuous-time setting we obtain the following result.

⁶ $\text{corr}(X_T, \xi_T) = \frac{E[\xi_T X_T] - E[\xi_T]E[X_T]}{\text{std}(\xi_T)\text{std}(X_T)}$. Since the marginal distributions for X_T and ξ_T are fixed, $E[X_T]$, $E[\xi_T]$, $\text{std}(X_T)$ and $\text{std}(\xi_T)$ are constant, and minimizing $E[\xi_T X_T]$ is the same as minimizing correlation.

⁷Dybvig (1988a) assumes non-negative payoffs in the proofs for continuous markets. See Assumption A1 page 389 in the Appendix A of Dybvig (1988a).

Corollary 1.2.3. *Let ξ_T be continuously distributed. Define*

$$Y_T^* = F^{-1}(1 - F_{\xi_T}(\xi_T)). \quad (1.5)$$

Then, Y_T^ is the cheapest way to achieve the distribution F . It is also almost surely unique.*

Note that $F_{\xi_T}(\xi_T)$ has a uniform $U(0, 1)$ distribution (as well as $1 - F_{\xi_T}(\xi_T)$) and thus Y_T^* is distributed with F . The proof of Corollary 1.2.3 is sketched in the Appendix of Dybvig (1988a). It has also been proven using involved arguments in Lemma B.1, in the Appendix of Jin and Zhou (2008) but our proof is simpler and holds more generally. Note that it is possible to relax the assumption on continuity of the distribution of ξ_T but optimal payoffs contain a random component (Bernard, Rüschendorf, and Vanduffel (2014)).

In the remainder of the chapter, mainly for ease of exposition, we will restrict the analysis to the case when ξ_T is continuously distributed. Under this assumption, Prop.1.2.1 becomes a sufficient and necessary condition for cost-efficiency.

Proposition 1.2.4 (Characterization of cost-efficiency). *A payoff is cost-efficient if and only if it is non-increasing in the state-price ξ_T almost surely.*

This characterization can be used to show that path-dependent strategies are not cost-efficient under some conditions. We say that a strategy with payoff X_T at time T is path-dependent if X_T depends not only on the final value (at time T) of the traded assets but on previous values as well. Otherwise the strategy is called path-independent. Bondarenko (2003) shows that the existence of a path-independent pricing kernel ξ_T is equivalent to the absence of statistical arbitrage opportunities, where a statistical arbitrage opportunity is defined as a zero-cost trading strategy with a positive expected payoff and nonnegative conditional expected payoffs for each ξ_T . He also shows that the assumption of a path-independent state-price process is satisfied by many popular asset pricing models, including the CAPM, the consumption-based models, the multifactor pricing models and the Black Scholes model, amongst others. Prop.1.2.4 states that cost-efficient payoffs have to be non-increasing in the state-price ξ_T . Hence, when ξ_T is path-independent (as in the Black Scholes model in Section 1.2.2 hereafter) then path-dependent payoffs are not cost-efficient. However ξ_T may be path-dependent in some more complicated markets. For example, in the model of Miltersen and Schwartz (1998) for the valuation of options on futures (revisited by Manoliu and Tompaidis (2002)) the state-price process becomes clearly path-dependent and involves $\int_0^T \gamma_t^j dW_t^j$ for some general adapted process γ_t^j and the Brownian motion W_t^j (see for instance formula (6) of Manoliu and Tompaidis (2002) for an explicit form of the state-price process in that model). Cost-efficient strategies write as (1.5) and are thus path-dependent as they inherit the path-dependence of the state-price process. Note that in the case in which we assume deterministic volatilities, the γ_t^j are deterministic and the state-price process ξ_T is lognormal.

Remark 1.2.1. Our study allows to derive bounds on the cost of a strategy with a given distribution. The strategy Z_T^* that generates the distribution F at highest cost is unique almost surely and is equal to $Z_T^* = F^{-1}(F_{\xi_T}(\xi_T))$. For any payoff X_T that is distributed with the cdf F , one has the following bounds

$$P_D(F) = \int_0^1 F_{\xi_T}^{-1}(\gamma) F^{-1}(1 - \gamma) d\gamma \leq c(X_T) \leq E[\xi_T Z_T^*] = \int_0^1 F_{\xi_T}^{-1}(v) F^{-1}(v) dv, \quad (1.6)$$

where the distributional price $P_D(F)$ is calculated as $P_D(F) = E[\xi_T Y_T^*]$ and Y_T^* is the cost-efficient strategy given in Corollary 1.2.3.

1.2.2 Black Scholes Market

We now give some specific results for the one-dimensional Black Scholes market. This is a convenient model since there is a simple analytical relation between the state-price and the price of the underlying asset. In this setting, we are able to derive simple expressions for cost-efficient strategies in Prop.1.2.5 which will be useful to derive closed-form expressions in Section 1.4. Here the dynamics of the underlying stock price are given by

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t,$$

where W_t is a standard Brownian motion under the physical measure \mathbb{P} , μ is the drift and σ is the volatility. The solution for the stock price is $S_T = S_0 \exp\left(\left(\mu - \frac{\sigma^2}{2}\right)T + \sigma W_T\right)$, with cdf F_{S_T} at time T given by

$$F_{S_T}(x) = \mathbb{P}(S_T \leq x) = \Phi\left(\frac{\ln\left(\frac{x}{S_0}\right) - \left(\mu - \frac{\sigma^2}{2}\right)T}{\sigma\sqrt{T}}\right), \quad (1.7)$$

where Φ is the cdf of a standard normal random variable. Let r denote the continuously compounded risk-free interest rate. The (unique) state-price process can be computed explicitly as $\xi_t = e^{-rt} e^{-\frac{1}{2}\left(\frac{\mu-r}{\sigma}\right)^2 t} e^{-\left(\frac{\mu-r}{\sigma}\right)W_t}$. Consequently, ξ_T can be written as an explicit function of the stock price S_T as follows

$$\xi_T = \alpha \left(\frac{S_T}{S_0}\right)^{-\beta} \quad (1.8)$$

where $\alpha = \exp\left(\frac{\theta}{\sigma}\left(\mu - \frac{\sigma^2}{2}\right)T - \left(r + \frac{\theta^2}{2}\right)T\right)$, $\beta = \frac{\theta}{\sigma}$, $\theta = \frac{\mu-r}{\sigma}$. When $\mu = r$ the state-price ξ_T is constant and Prop.1.2.1 implies that every payoff is cost-efficient. Otherwise ξ_T is stochastic with a continuous distribution. The “cost-efficient payoff” distributed with cdf F has now the closed-form expression as a function of ξ_T given by (1.5). The functional relation between ξ_T and S_T , established in (1.8), allows us to express cost-efficient payoffs in terms of S_T . This is the subject of the next proposition.

Proposition 1.2.5 (Cost-efficiency in the Black Scholes model). *Under the assumptions of the Black Scholes model, when $\mu > r$, the (almost surely) cost-efficient payoff with cdf F is given by $F^{-1}(F_{S_T}(S_T))$, where F_{S_T} is the cdf of the stock price at time T . When $\mu < r$ the (almost surely) cost-efficient payoff is given by $F^{-1}(1 - F_{S_T}(S_T))$.*

When a closed-form expression for F^{-1} is available the cost-efficient payoff Y_T^* can be computed explicitly using the expression (1.7) for the cdf of the stock-price S_T . Examples are given in Section 1.4. Under the assumptions of the Black Scholes model with $\mu > r$ (respectively $\mu < r$) a payoff is cost-efficient if and only if it is a non-decreasing (respectively non-increasing) payoff in the underlying stock price S_T . This leads to the next (obvious) corollary.

Corollary 1.2.6 (Path-dependent payoffs in the Black Scholes model). *In the Black Scholes model, path-dependent payoffs are not cost-efficient unless $\mu = r$.*

Using this corollary, it follows that stop loss strategies, CPPI strategies and popular dynamic strategies that rebalance at discrete dates may never be cost-efficient. Therefore, it is possible to construct dominating strategies. Examples will follow later in this chapter. See Bernard, Maj and Vanduffel (2011) for examples in the multidimensional Black Scholes model. For convenience in the remainder of the chapter all results presented in the context of a Black Scholes market assume that $\mu > r$.

1.3 Agents with State-Independent Preferences

In this section we establish the close links between cost-efficiency of payoffs and their optimality properties for agents with state-independent preferences, meaning that they only care about the distribution of their wealth. We also assume that agents have a fixed investment horizon and no intermediate consumption.

The next proposition points out the strong connection between cost-efficiency and first-order stochastic dominance. It makes clear that cost-efficient payoffs are “better” for agents with increasing preferences. Next in Prop.1.3.2 we establish new connections with second-order stochastic dominance. In particular these results give additional ways to construct “better” payoffs. In what follows consider a payoff X_T with cdf F .

Proposition 1.3.1 (First-order stochastic dominance).

1. Taking into account the initial costs, the cost-efficient payoff Y_T^* for the payoff X_T is a.s. equal to X_T or dominates X_T in the first-order stochastic dominance sense (we write \prec_{fsd}),

$$(X_T - c(X_T)e^{rT}) \prec_{fsd} (Y_T^* - P_D(F)e^{rT}) \quad (1.9)$$

2. The dominance ordering (1.9) is strict unless X_T is non-increasing in ξ_T .

The proposition implies that for any given payoff X_T the cost-efficient payoff Y_T^* will be at least as good for all agents who prefer “more to less”.⁸ These agents only need to consider payoffs which are non-increasing in the terminal value of the state-price process, ξ_T .

In particular these statements hold true for expected utility maximizers with a non-decreasing utility function, including risk-loving investors (with convex utility), risk-averse investors (with concave utility) and loss averse⁹ investors (with S-shaped utility)), and CPT-investors (defined as investors acting consistently with the cumulative prospect theory developed by Tversky and Kahneman (1992)). Note that also target probability maximizers (see Browne (1999)) prefer cost-efficient strategies. Browne makes use of stochastic control theory in the Black Scholes setting to show that the optimal strategy amounts to the purchase of a digital option. From our results we know the optimal strategy must be non-decreasing in S_T implying Browne’s result for a target probability maximizer in a straightforward way.

In Prop.1.3.1, the inefficient strategy X_T is dominated in the sense of first-order stochastic dominance by the optimal cost-efficient strategy Y_T^* , which is cheaper and has the same distribution. Next, we show that X_T can also be dominated in the sense of second-order stochastic dominance with another strategy that has the same cost but a different distribution. The idea is to make the payoff X_T less dispersed. By virtue of Jensen’s inequality all one needs to do is to substitute X_T by its conditional expectation with respect to any other variable. Intuitively, one may then expect that decreasing the randomness, and thus risk, of the payoff X_T goes hand in hand with increasing its cost. But this property does not always hold. Actually one can easily show that conditioning on the state-price ξ_T always strictly preserves the cost, and thus generates a payoff which is preferred to the original payoff X_T . The following proposition states this formally.

⁸This observation is the key element needed to show that the optimal wealth for an investors with non-decreasing state-independent preferences is also the optimal wealth for an expected utility maximizer (Bernard, Chen and Vanduffel, (2015)).

⁹The terminology “loss-averse” refers here to the paper by Berkelaar, Kouwenberg and Post 2004.

Proposition 1.3.2 (Second-order stochastic dominance).

1. Taking into account the initial costs, the payoff $H_T^* = E[X_T | \sigma(\xi_T)]$ (where $\sigma(\xi_T)$ is the σ -algebra generated by ξ_T) is a.s. equal to X_T or dominates X_T in the second-order stochastic dominance sense (we write \prec_{ssd}),

$$(X_T - c(X_T)e^{rT}) \prec_{ssd} (H_T^* - c(H_T^*)e^{rT}). \quad (1.10)$$

2. The dominance ordering given by (1.10) is strict unless X_T is a function of ξ_T .

In the specific context of a one-dimensional Lévy market with so-called Esscher pricing (which amounts to making a particular choice for the state price process) this result can also be found in Vanduffel et al. (2009). Here we show that the result holds for fairly general markets. To the best of our knowledge we are first in doing so. In particular, no restriction on the dimensionality of the market is needed and no particular form for the state-price process is assumed.

Prop.1.3.1 and 1.3.2 show that a given reference payoff X_T can be improved either by making it cost-efficient (Y_T^*) or by conditioning (H_T^*). It is clear that H_T^* may not be cost-efficient. It will be cost-efficient only if it is a non-increasing function of ξ_T . Since H_T^* and X_T have the same mean and the same cost, H_T^* can never dominate X_T in the sense of first-order stochastic dominance. In a Black Scholes market, H_T^* is of particular interest when the expected return is difficult to estimate; see Merton 1972 for evidence that this is often the case in practice. Indeed, the payoff H_T^* does not depend on the drift μ^{10} and thus H_T^* improves upon X_T also in case the investor has difficulty in estimating μ . For example, the strategy that Vanduffel et al. (2012) propose to outperform the popular dollar cost averaging strategy does not depend on μ . In contrast, the cost-efficient payoff Y_T^* usually depends on μ explicitly.

1.4 Examples

In this section we illustrate the results obtained so far using some common payoff distributions. We first investigate a short forward contract and a long put option both of which are path-independent but inefficient. Then we examine a path-dependent payoff: the lookback option. All results in this section are new at the time they were first derived,¹¹ and detailed derivations can be found in the published version of the chapter (Bernard, Boyle, and Vanduffel (2014)). For simplicity, we assume a Black Scholes market and use the same notation as in Section 1.2.2.

1.4.1 Forward Contract

One of the simplest examples of a financial product with a decreasing payoff in relation to the stock price is a short forward contract. A short position in a forward contract has a payoff $X_T = K - S_T$. Its cdf F is given by $F(x) = \mathbb{P}(S_T > K - x) = 1 - F_{S_T}(K - x)$ where F_{S_T} is given in (1.7). Then, $F^{-1}(y) = K - F_{S_T}^{-1}(1 - y)$. Since the payoff is decreasing in S_T , it is inefficient. Applying Prop.1.2.5,

¹⁰This property follows from the bijection between the state price ξ_T and the risky asset price S_T and the observation that the drift parameter does not affect the distribution of a Brownian bridge (see for instance Beghin and Orsingher (1999)).

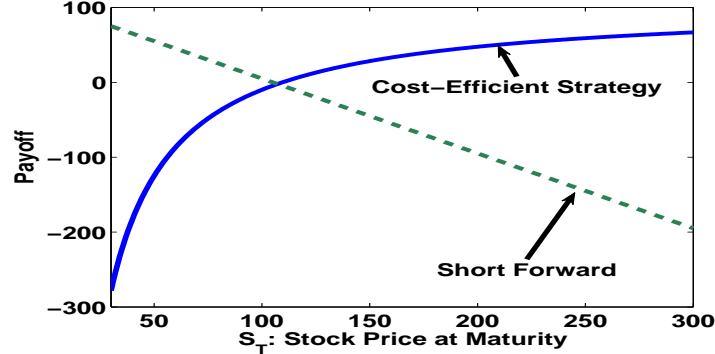
¹¹Bernard and Boyle (2010)

the corresponding cost-efficient strategy is given by $Y_T^* = F^{-1}(F_{S_T}(S_T)) = K - F_{S_T}^{-1}(1 - F_{S_T}(S_T))$. After some straightforward calculations, we obtain

$$Y_T^* = K - \frac{c}{S_T} \quad (1.11)$$

where $c = S_0^2 e^{2(\mu - \frac{\sigma^2}{2})T}$. The payoff profile is given in Figure 1.4.1. The no-arbitrage price of the short forward contract is equal to $Ke^{-rT} - S_0$, the cost of the cost-efficient strategy is equal to $c(Y_T^*) = Ke^{-rT} - S_0 e^{2(\mu - r)T}$. Since $\mu > r$, we can readily verify that $c(Y_T^*) < c(X_T)$.

Figure 1.4.1: Payoff of the short forward contract and its cost-efficient counterpart.



Parameters: $\sigma = 20\%$, $\mu = 9\%$, $r = 5\%$, $S_0 = 100$, $T = 1$ year, strike $K = 100$. The graph shows the payoffs (as functions of the stock price S_T) of the short forward position and its cost-efficient counterpart that has the same distribution.

We can give an intuitive interpretation of the payoff in formula (1.11). Write S_T as $S_T(z)$ where $z \sim \mathcal{N}(0, 1)$ and

$$S_T(z) = S_0 \exp \left(\left(\mu - \frac{\sigma^2}{2} \right) T + \sigma \sqrt{T} z \right). \quad (1.12)$$

Then the payoff of the short forward is $K - S_T(z)$ and the payoff of the cost-efficient strategy (1.11) can be written as $Y_T^* = K - S_T(-z)$. Notice that $S_T(z)$ and $S_T(-z)$ are equal in distribution. Then, $K - S_T(z)$ has the same distribution as $K - S_T(-z)$. This relationship illustrates how the efficient payoff is again a short forward contract, but now on an asset that is perfectly negatively correlated (in logs) with the original asset.

Finally note that both contracts have potentially unlimited liability and that this occurs at opposite extremes of the stock price distribution. The maturity value of the standard short forward tends to minus infinity as the stock price tends to infinity, while for the cost-efficient payoff the infinite loss occurs when the stock price is zero.

From this example, it is already clear that an individual who enters a long position in a forward contract is usually not only interested in the (stand-alone) distribution of his contract. Rather he is hedging an existing position and is mainly concerned with the total position $K - S_T + L_T$ where L_T is some other asset that is correlated with S_T . This feature will become even clearer in the example with the put option. The cost-efficient alternative will then be $F^{-1}(F_{S_T}(S_T))$ where F is the cdf of $K - S_T + L_T$.

1.4.2 Put Options

Another simple example of a payoff distribution with a decreasing payoff in the stock price is the standard put option. From Prop.1.2.5 it is not optimal for investors who have state-independent and increasing preferences (i.e. who only care about the distribution of final wealth and prefer more to less) to hold only puts in their portfolio. A put option with strike K and maturity T has the following payoff $X_T = (K - S_T)^+$, and hence X_T is increasing in the state-price process ξ_T . Let F denote the cumulative distribution function of the payoff of the put option. The put option is therefore the (a.s.) unique payoff that has the highest possible cost and is distributed with cdf F given by

$$F(x) = \mathbb{P}(X_T \leq x) = \begin{cases} 1 & \text{if } x \geq K \\ \mathbb{P}(S_T > K - x) = \Phi\left(\frac{(\mu - \frac{\sigma^2}{2})T - \ln(\frac{K-x}{S_0})}{\sigma\sqrt{T}}\right) & \text{if } 0 \leq x < K \\ 0 & \text{if } x < 0 \end{cases}$$

It is straightforward to invert F and $F^{-1}(y) = \left(K - S_0 e^{(\mu - \frac{\sigma^2}{2})T - \sigma\sqrt{T}\Phi^{-1}(y)}\right)^+$ for $y \in (0, 1)$. Applying Prop.1.2.5, the cost-efficient payoff that gives the same distribution as a put option is $Y_T^* = F^{-1}(F_{S_T}(S_T))$. Since F_{S_T} is given by (1.7), we find

$$Y_T^* = \left(K - \frac{c}{S_T}\right)^+, \quad (1.13)$$

where $c = S_0^2 e^{2(\mu - \frac{\sigma^2}{2})T}$. This cost-efficient payoff is represented graphically in Figure 1.4.2. The no-arbitrage price at time zero of this payoff is given by

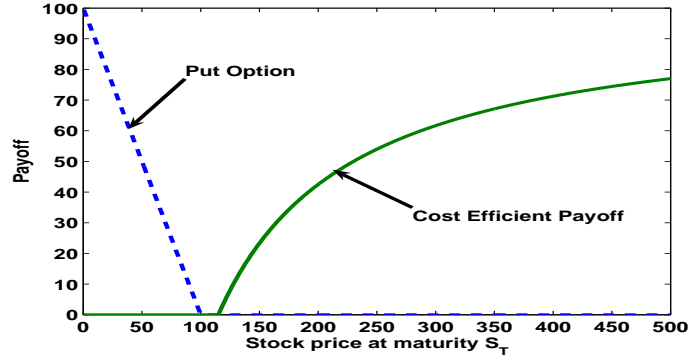
$$K e^{-rT} \Phi(-d_4) - S_0 e^{2(\mu-r)T} \Phi(-d_3), \quad (1.14)$$

where $d_3 = \frac{\ln(\frac{S_0}{K}) + (2\mu-r)T + \frac{\sigma^2 T}{2}}{\sigma\sqrt{T}}$ and $d_4 = d_3 - \sigma\sqrt{T}$. A portfolio consisting of a put option (or entirely of put options) is clearly not optimal for expected utility maximizers with increasing utility. It is dominated in the sense of first-order stochastic dominance by the European payoff given by Y_T^* in (1.13). Figure 1.4.2 displays both payoffs.

Based on the parameters given in Figure 1.4.2, the cost of the put option is 5.574 while the cost of the corresponding cost-efficient payoff is 3.146. This cost-efficient payoff has the same distribution under the real probability measure as the payoff of the put option. Hence the efficiency loss for this put option is 2.428 (that is 44% of its market price). Buying a put option in this context rather than the payoff Y_T^* represents a significant efficiency loss. This implies that a rational agent with state-independent preferences (and a state price density that is path independent) would not invest her entire wealth in a put option.

Similarly as for the short forward contract, there is an intuitive interpretation of formula (1.13). Using the expression (1.12), the payoff of the put option is $(K - S_T(z))^+$ and the payoff of the cost-efficient strategy is $(K - S_T(-z))^+$. This relationship illustrates how the efficient payoff is also a put option but now on the asset that is perfectly negatively correlated with the initial payoff. It also explains why its price as displayed in (1.14) is the Black Scholes formula for the price of another put option with adjusted parameters.

Figure 1.4.2: Payoff of a put option and its cost-efficient counterpart



Parameters: $\sigma = 20\%$, $\mu = 9\%$, $r = 5\%$, $S_0 = 100$, $T = 1$ year, strike $K = 100$. The graph shows the payoffs (as functions of the stock price S_T) of the put option and its cost-efficient counterpart that has the same payoff distribution.

1.4.3 Path-Dependent Example: Lookback Option

We now consider a path-dependent payoff: a lookback option. The paper published in *Finance* also solves the case of the Asian option. We already know that they cannot be cost-efficient from Corollary 1.2.6. We are able to obtain expressions for payoffs that are dominating.

Consider a lookback call option with strike K . The payoff of this option is given by

$$L_T = \left(\max_{0 \leq t \leq T} \{S_t\} - K \right)^+.$$

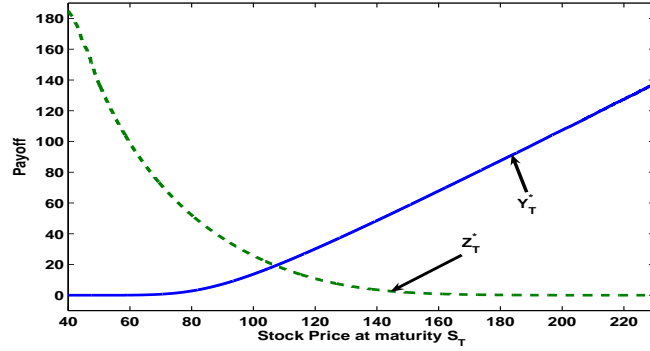
We now derive the cost-efficient strategy for this payoff distribution. Let F_L be the cdf of L_T . The payoff that has the lowest cost and has the same distribution as L_T is given by $Y_T^* = F_L^{-1}(F_{S_T}(S_T))$. The payoff that has the highest cost and has the same distribution as L_T is given by $Z_T^* = F_L^{-1}(1 - F_{S_T}(S_T))$. Hence the payoff Y_T^* can be written as

$$Y_T^* = F_L^{-1} \left(\Phi \left(\frac{\ln \left(\frac{S_T}{S_0} \right) - \left(\mu - \frac{\sigma^2}{2} \right) T}{\theta \sqrt{T}} \right) \right). \quad (1.15)$$

Figure 1.4.3 plots the payoff distribution of the cost-efficient distribution associated with the lookback option as the solid line for the specified parameter values. The price of the lookback option in this case is 19.17. The price of the associated cost-efficient payoff is 18.85. Hence the efficiency loss in this case is 0.32 (which is about 1.7% of the market price). Notice that the payoff profile of the cost-efficient payoff is similar to that of a standard European call. For comparison purposes, we display the profile of the payoff Z_T^* which has the same distribution as the lookback call and has the highest initial cost (25.26). This inefficient payoff is similar to the payoff of a put option.

Observe that an investor who maximizes the expected utility would get the same expected utility for the three payoffs (the Lookback contract and the two payoffs displayed in Figure 1.4.3). However, one of the contract is strictly cheaper, which means that it strictly dominates the two

Figure 1.4.3: Payoffs related to the cdf of a lookback call option



Parameters $\sigma = 20\%$, $\mu = 9\%$, $r = 5\%$, $S_0 = 100$, $T = 1$ year, strike $K = 100$. The graph shows the payoff of the cost-efficient derivative Y_T^* (as a function of the stock price S_T) that has the same distribution at time T as the lookback option. Z_T^* shows the most expensive payoff with the same distribution as the lookback option.

other payoffs as the extra budget can be invested in a bank account to ensure a strictly higher expected utility for any increasing utility function (first-order stochastic dominance).

We have discussed two methods of constructing a payoff that dominates a path-dependent payoff. The first method constructs a payoff with the same distribution and a strictly lower cost: the cost-efficient payoff dominates the lookback option in the sense of first-order stochastic dominance. The second method is discussed in Prop.1.3.2 in Section 1.3. There we noted that the conditional expectation with respect to the σ -algebra generated by ξ_T provides an alternative way of obtaining a strictly “better” payoff than the lookback option in the sense of second-order stochastic dominance. This constructed payoff has the same no-arbitrage price as the original lookback call option.

We denote the conditional expectation by H_T^* , where $H_T^* = E[L_T | \sigma(\xi_T)]$. In the Black Scholes setting, this payoff is equal to $H_T^* = E[L_T | \sigma(S_T)]$ because there is a measurable bijection between S_T and ξ_T (this relationship was given in (1.8)).

$$H_T^* = E \left[\left(\max_{0 \leq t \leq T} \{S_t\} - K \right)^+ \middle| \sigma(S_T) \right]$$

Using the properties of the maximum of a Brownian bridge, one can derive the following formula for H_T^*


$$H_T^* = \begin{cases} S_T - K + \frac{\sqrt{2\pi T}}{2} S_0 \sigma e^{\frac{1}{8T} \left(\frac{2}{\sigma} \ln \left(\frac{S_T}{S_0} \right) + \sigma T \right)^2} \Phi \left(\frac{-\frac{2}{\sigma} \ln \left(\frac{S_T}{S_0} \right) + \sigma T}{2\sqrt{T}} \right) & \text{if } S_T \geq K, \\ \frac{\sqrt{2T\pi}}{2} S_0 \sigma e^{\frac{1}{8T} \left(\frac{2}{\sigma} \ln \left(\frac{S_T}{S_0} \right) + \sigma T \right)^2} \Phi \left(\frac{\frac{2}{\sigma} \ln \left(\frac{S_0 S_T}{K^2} \right) + T\sigma}{2\sqrt{T}} \right) & \text{if } S_T < K. \end{cases} \quad (1.16)$$

The payoff H_T^* is a function of S_T . Since it is not path-dependent it can be represented on the graph. Panel A of Figure 1.4.4 displays the payoffs of H_T^* and the cost-efficient payoff Y_T^* in the case of the lookback call option as a function of S_T . They are quite close to each other.

We proved that H_T^* dominates the payoff L_T in the sense of the second-order stochastic dominance (see Prop.1.3.2). This can be easily seen from Panel B in Figure 1.4.4 where the two cdfs cross

just once. The distribution of H_T^* and of the lookback payoff have the same mean but obviously the distribution of H_T^* has less spread.

Figure 1.4.4: Comparison of payoffs of H_T^* and Y_T^*



Parameters: $\sigma = 20\%$, $\mu = 9\%$, $r = 5\%$, $s_0 = 100$, $T = 1$ year, strike $k = 100$. In Panel A, the solid line represents the payoff H_T^* and the dashed line represents the payoff of the cost-efficient payoff Y_T^* (as functions of the stock price S_T). In Panel B, the solid line represents the cdf of H_T^* and the dashed line the cdf of the lookback call option.

1.5 Cost-efficiency with State-Dependent Constraints

Recall that the state-price density ξ_T is the ratio of the price for an Arrow-Debreu security per unit of probability. Standard economic theory predicts a high price of the state when the market is down or the economy is in recession, since buying consumption in states of scarcity is a form of insurance (see for example Dybvig and Ross (1997)). This is also confirmed in the Black Scholes market where there usually is an explicit negative relationship between the state-price and the price of the risky asset (see also (1.8)). A given reference payoff can then be made cheaper, without impact on its distribution, by rearranging consumption so that it becomes reversely ordered with the state price density. Unfortunately the optimized payoff does no longer offer protection against poor economic situations, showing that cost-efficiency comes at some price eventually. To illustrate this insight, consider a Black Scholes market and let X_T be a strategy and X_T^* its cost-efficient counterpart, so that both are equally distributed at maturity. Then it holds for all (finite) $a > 0$ that

$$E[X_T^* \mid S_T \leq a] < E[X_T \mid S_T \leq a], \quad (1.17)$$

reflecting that the cost-efficient payoff X_T^* provides strictly lower average consumption in downside scenarios. Recall also that in Section 1.4.2 we compare the payoff $X_T := (K - S_T)^+$ of a standard put with a suitable power put $X_T^* := (K - \frac{c}{S_T})^+$ with $c = S_0^2 e^{2(\mu - \frac{\sigma^2}{2})T}$. Whilst they share the same distribution, these strategies interact with S_T in a fundamentally different way. A standard put provides income during a recession when consumption is more expensive, whereas the cost-efficient counterpart delivers no wealth in this instance. Standard put options are not used in isolation but for their feature of offering protection when the economy is in a downturn, and investors are prepared to pay an extra premium for obtaining this protection. The same feature also holds

in an insurance context where investors are prepared to pay a higher price than the distributional (financial) price for a given insurance contract (Bernard and Vanduffel (2014a)). These observations strengthen the case for extending cost-efficiency to deal with state-dependent preferences.

We still assume that the agent wants a given distribution F at maturity but we also assume his investment strategy is subject to additional state-dependent constraints. In our setting each constraint gives information on the interaction between the strategy Y_T and the financial market at time T (through the state-price process at time T). This formulation permits us to accommodate the constraints in a very natural way using tools from the theory of copulas.¹²

First we explain how additional information about the joint behavior between the final payoff of the strategy and the market can be formulated and modeled. We then introduce the concept of “constrained cost-efficient payoffs” (cheapest strategies that satisfy a given set of constraints) and provide a characterization result. We illustrate the theoretical results with several examples for which “constrained cost-efficient” strategies can be derived explicitly. In particular we discuss strategies that are optimal for agents who want some degree of protection for their portfolio when the market is in a crisis. Throughout this section it is assumed that for all $0 < t < T$ the state-prices ξ_t and ξ_T are jointly continuously distributed.

1.5.1 Constrained Cost-Efficient Payoffs

We have shown in Section 1.2 that cost-efficient payoffs are non-increasing with the state-price process at time T (Prop.1.2.4). We recall that this result is obtained by identifying the dependence structure between the final payoff Y_T with cdf F and the state-price process ξ_T which minimizes the cost $E[\xi_T Y_T]$ of the strategy. As a result cost-efficiency is intimately related to the study of dependence or “copulas”. It is known that the joint distribution of (Y_T, ξ_T) is characterized by the respective marginal distributions of ξ_T and Y_T and the manner in which these variables interact. More precisely we have

$$\mathbb{P}(\xi_T \leq s, Y_T \leq y) = C(F_{\xi_T}(s), F(y)), \quad (1.18)$$

where C is the joint distribution for a pair of uniform random variables over $(0, 1)$. The distribution function C is called a copula¹³.

From (1.18) it is clear that the cost $E[\xi_T Y_T]$ is a function of F_{ξ_T} , F and the copula C . Furthermore it can be shown that minimizing this cost amounts to finding the minimal copula C (see Prop.1.5.1 below). This forms the basis for the proof of Prop.1.2.1 where we show that the minimal cost is obtained for the anti-monotonic dependence (in other words for the *anti-monotonic copula*) $C(u, v) = \max(0, u + v - 1)$ which is the joint distribution of $(U, 1 - U)$ where U is a standard uniform random variable. This section extends Corollary 1.2.3 and Prop.1.2.4 to include constraints on the dependence between the payoff Y_T and the market.

To illustrate the nature of the constraints considered here, we provide a simple example. We consider a Black Scholes market with a single asset S . Assume an agent wants to achieve the same distribution F as S_T but is subject to additional constraints. For example he requires that

$$\mathbb{P}(Y_T > 100 \mid S_T < 95) = 0.8. \quad (1.19)$$

¹²Cherubini, Luciano, and Vecchiato (2004), Nelsen (2006),

¹³This result is also known as Sklar’s theorem and shows that a copula C combines the marginal distributions F_{ξ_T} and F to obtain the joint distribution for (ξ_T, Y_T) . When Y_T is continuously distributed, the copula C is unique almost surely. It is also invariant under strictly increasing transformations of ξ_T and Y_T .

This constraint reflects the desire of the agent to receive some protection when the market declines. It is just one particular example of a possible state-dependent preference. Of course one can also express this constraint in terms of the joint distribution between Y_T and S_T , and more generally between Y_T and ξ_T , which is the setting we use in this chapter.

Formally, for a given distribution F we are interested in finding the cheapest possible payoff Y_T with the same distribution as F and which satisfies additional constraints of the form

$$\mathbb{P}(\xi_T \leq x, Y_T \leq y) = Q(F_{\xi_T}(x), F(y)), \quad (1.20)$$

with $x > 0, y \in \mathbb{R}$ and Q a given feasible function. It is sufficient to assume that $Q(u, v)$ is a copula (that is a joint distribution of a pair of standard uniform random variables). Each constraint gives information on the dependence between the state-price ξ_T and Y_T and is, for a given function Q , determined by the pair $(F_{\xi_T}(x), F(y))$. Denote the set of all such constraints by \mathbb{S} . This set can be finite or infinite. For convenience we denote its elements by (a, b) .

Hence, we seek payoffs that solve the following problem:

$$\min_{\{Y_T \mid Y_T \sim F, \mathbb{S}\}} \mathbb{E}[\xi_T Y_T]. \quad (1.21)$$

A payoff that solves problem (1.21) is called a \mathbb{S} -constrained cost-efficient payoff. Note that when \mathbb{S} is empty the formulation (1.21) for the constrained cost-efficiency problem coincides exactly with the original problem (1.2). We say that a copula C satisfies \mathbb{S} if $C(a, b) = Q(a, b)$ for all $(a, b) \in \mathbb{S}$. Note that when \mathbb{S} is the full unit square, then the copula¹⁴ is fully specified and thus also the joint distribution. This case is extensively studied by Bernard, Moraux, Rüschendorf and Vanduffel (2015). The following proposition clarifies the crucial role of finding bounds for copulas when determining “constrained cost-efficient” payoffs.

Proposition 1.5.1 (Constrained cost-efficient payoffs). *Let \mathbb{S} be a set of constraints and $t \in (0, T)$. If there exists a copula L satisfying \mathbb{S} such that $L \leq C$ (pointwise¹⁵) for all other copulas C satisfying \mathbb{S} then the payoff Y_T^* given by*

$$Y_T^* = F^{-1}(f(\xi_T, \xi_t)) \quad (1.22)$$

solves (1.21) with $f(\xi_T, \xi_t) = h\left(F_{\xi_T}(\xi_T), j(F_{\xi_T}(\xi_T), F_{\xi_t}(\xi_t))\right)$, where the functions $j(u, v)$ is defined as the first partial derivative for $(u, v) \rightarrow J(u, v)$ where J denotes the copula for the random pair (ξ_T, ξ_t) . Formally $j(u, v) = \frac{\partial}{\partial u} J(u, v)$. For each given u , we define $h(u, v)$ as the inverse of the following function of v

$$v \mapsto \frac{\partial}{\partial u} L(u, v).$$

The partial derivatives can be interpreted in terms of conditional distributions: If (U, V) has a copula J then $j(u, v) = \mathbb{P}(V \leq v \mid U = u)$. Since $j(F_{\xi_T}(\xi_T), F_{\xi_t}(\xi_t))$ is independent of ξ_T , the dependence between Y_T^* and ξ_T is given by the minimum copula L . Prop.1.5.1 gives a sufficient condition (in terms of a minimal copula L) for the existence of an optimal solution to Problem (1.21), and provides an explicit construction when this condition is met. Interestingly, the \mathbb{S} -constrained cost-efficient payoff does not necessarily depend only on the value ξ_T of the state-price process at

¹⁴Similarly one may study state-dependence by imposing constraints on the correlation between the strategy and one or more market indices (see Bernard and Vanduffel Bernard and Vanduffel (2014b) for an application of such constraints).

¹⁵ $L \leq C$ “pointwise” means that $\forall (u, v) \in [0, 1]^2, L(u, v) \leq C(u, v)$.

time T but also on the value ξ_t at time t which can be chosen freely for any $0 < t < T$. Thus it follows that the \mathbb{S} -constrained cost-efficient payoff is not necessarily unique.

Here is some intuition why constrained cost-efficient payoffs generally do not depend on ξ_T only. When there are no constraints one can make a given payoff cost-efficient by redistributing wealth levels across the different states of the economy such that they become reversely ordered with the state-price (cost-efficient strategies are non-increasing in the state-price ξ_T). In other words the payoff and the state-price become functions of each other. Adding constraints will however impose structure on the dependence between both variables ex-ante, and thus enforces departures from this perfect dependence pattern between the payoff and the state-price. It may thus no longer be possible to write one as function of the other.

In a Black Scholes framework, we know that ξ_t is a function of S_t and ξ_T a function of S_T only. Thus optimal payoffs (as in Prop.1.5.1) are weakly path-dependent in general, in the sense that they are functions of S_t and S_T . In the absence of state-dependent constraints however it should follow that $f(\xi_T, \xi_t) = 1 - F_{\xi_T}(\xi_T)$. The next example shows that the more complex expression (1.22) reduces to the original expression (1.5) for cost-efficient payoffs in the case of no constraints.

Example 1 (No constraints). When there are no constraints, \mathbb{S} is the empty set. Furthermore it is known from the Fréchet-Hoeffding bounds¹⁶ on copulas that $\forall (u, v) \in [0, 1]^2$, $C(u, v) \geq \max(0, u + v - 1)$. Note then that $L(u, v) := \max(0, u + v - 1)$ is indeed a copula as the joint distribution of the random pair $(U, 1 - U)$ where U is a standard uniform random variable. It is then straightforward to show that $\frac{\partial L(u, v)}{\partial u} = 1$ if $v > 1 - u$ and that $\frac{\partial L(u, v)}{\partial u} = 0$ if $v < 1 - u$. Hence, we find that $h(u, p) = 1 - u$ for all $0 < p \leq 1$ which implies that $f(\xi_T, \xi_t) = 1 - F_{\xi_T}(\xi_T)$. It follows that Y_T^* is given by

$$Y_T^* = F^{-1}(1 - (F_{\xi_T}(\xi_T))), \quad (1.23)$$

which corresponds to expression (1.5) obtained for the unconstrained cost-efficient payoff.

Prop.1.5.1 provides a construction for a constrained cost-efficient strategy when a minimal copula L exists. A construction of L is proposed by Tankov (2011) and sufficient conditions to ensure that L is a copula have been derived in Bernard, Jiang, and Vanduffel (2012), Bernard, Liu, MacGillivray, and Zhang (2013). In the published version of this chapter, we first examine simple examples with a finite number of constraints. We then look at examples with an infinite number of constraints. Such examples will be further developed in Chapter 5.

1.6 Summary and Conclusion

This chapter uses a preference-free framework for ranking different investment strategies. For a given investment strategy, we derive an explicit analytical expression for the cheapest strategy that has the same payoff distribution as the given strategy. We also provide an expression for an equal cost strategy that has a different, but less dispersed, payoff distribution. These payoffs are all expressed in terms of the state-price density. Our framework is used to analyze the strategies that produce some common option type payoffs and to discuss the inefficiency of path-dependent strategies. This in turn leads to some new types of relationships and insights.

We motivate and develop a new approach to handle constraints in this framework. The approach we propose to accommodate state-dependent preferences is novel and makes use of the theory of

¹⁶Fréchet (1960), Hoeffding (1963).

copulas. We show that the optimal investment for an investor with state-independent preferences is closely related to finding a lower bound for a certain copula. This important observation allows us to solve for new optimal strategies when there are constraints on the dependence between the strategy and the financial market. Although these strategies are not optimal for state-independent preferences, they become optimal for managers who have state-dependent preferences and seek some particular properties of their final wealth.

There are many interesting applications of the concept of cost-efficiency in the context of law-invariant preferences. We will only present two applications in full detail in Chapters 2 and 3. Specifically, in Chapter 2, we will show that there is an equivalence between cost-efficiency and law-invariant preferences. It can be used to improve the design of retail investment products (Bernard, Maj, and Vanduffel (2011)). In Bernard, Chen, and Boyle (2016), we use cost-efficiency to design executive compensation packages that improve upon the Asian executive option proposed by Tian (2013). See also the PhD thesis of Jit Seng Chen cosupervised with S. Vanduffel (Chen (2015)). In Chapter 3, we show the implications of cost-efficiency in pricing insurance (e.g. equity-linked insurance where mortality risk and financial risks are covered simultaneously). The second part of this chapter has introduced the notion of “constrained cost-efficiency” to allow to account for constraints that are state-dependent. Part III gives several applications of constrained cost-efficiency. Specifically, in Chapter 4, we revisit mean variance optimal portfolio and give simple proofs of well known results. We then show how to account for some realistic constraints, such as correlation constraints. A potential application is to detect fraud such as the one of Madoff. In Chapter 5, we solve investment problems with state-dependent constraints. For example, we construct strategies that are independent of the market when markets fall. Finally, Chapter 6 can be seen as an extension of this chapter as we further characterize the optimal strategies for investors who seek to achieve some terminal distribution of wealth but also with some dependence to a benchmark.

Part II

Law-invariance

Chapter 2

Rationalizing Investors' Choices

This chapter is our first application of the concept of cost-efficiency. We are able to show that the Expected Utility paradigm can rationalize optimal investment choices of investors with preferences that satisfy first-order stochastic dominance. Specifically, the optimal investment strategy in any behavioral law-invariant (state-independent) setting corresponds to the optimum for an expected utility maximizer with an explicitly derived concave non-decreasing utility function. This result enables us to infer the utility and risk aversion of agents from their investment choice in a non-parametric way. We relate the property of decreasing absolute risk aversion (DARA) to distributional properties of the terminal wealth and of the financial market. Specifically, we show that DARA is equivalent to a demand for a terminal wealth that has more spread than the opposite of the log pricing kernel at the investment horizon. All detailed proofs can be found in the published version in *Journal of Mathematical Economics* (Bernard, Chen, and Vanduffel (2015)).

The von Neumann and Morgenstern Expected Utility Theory (EUT) has for decades been the dominant theory for making decisions under risk. Nonetheless, this framework has been criticized for not always being consistent with agents' observed behavior (e.g., the paradox of Allais (1953), Starmer (2000)). In response to this criticism, numerous alternatives have been proposed, most notably dual theory (Yaari (1987)), rank-dependent utility theory (Quiggin (1993)) and cumulative prospect theory (Tversky and Kahneman (1992)). These competing theories differ significantly, but all three satisfy first-order stochastic dominance (FSD). Indeed, many economists consider violation of this property as grounds for refuting a particular model; see for example Birnbaum and Navarrette (1998), and Levy (2008) for empirical evidence of FSD violations. Recall also that although the original prospect theory of Kahneman and Tversky (1979) provides explanations for previously unexplained phenomena, it violates FSD. To overcome this potential drawback, Tversky and Kahneman (1992) developed cumulative prospect theory.

In the presence of a continuum of states, we show that the optimal portfolio in any behavioral theory that respects FSD can be rationalized by the expected utility setting, i.e., it is the optimal portfolio for an expected utility maximizer with an explicitly known concave utility function. This implied utility function is unique up to a linear transformation among concave functions and can thus be used for further analyses of preferences, such as to infer the risk aversion of investors.

A surprising feature is that we only assume that the preferences respect FSD, which contrasts with earlier results on the rationalization of investment choice under expected utility theory. Dybvig ((1988a), Appendix A) and Peleg and Yaari (1975) among others, have worked on this problem assuming that preferences preserve second-order stochastic dominance (SSD). However, being SSD-preserving is quite a strong assumption, and while consistency with FSD is inherent, and even

enforced, in most decision theories, this is not readily the case for SSD. For instance, rank dependent utility theory satisfies FSD but not SSD (Ryan (2006)), and the same holds true for cumulative prospect theory (see e.g., Baucells and Heukamp (2006)). Thus, our results show that for any agent behaving according to the cumulative prospect theory (i.e., for a “CPT investor”) there is a corresponding expected utility maximizer with concave utility purchasing the same optimal portfolio, even if the CPT investor can exhibit risk seeking behavior with respect to losses. This approach, however, is not intended to dispense with alternative models to expected utility theory, as they have been developed mainly to compare gambles and not to deal with optimal portfolio selection per se. For instance, prospect theory can be particularly useful to explain puzzles in finance (Broihanne, Merli, and Roger (2008)).

Our results are rooted in the basic insight that under some assumptions, the marginal utility at a given consumption level is proportional to the ratio of risk-neutral probabilities and physical probabilities (Duffie (2010)). At first, it then seems obvious to infer a (concave) utility function and the risk aversion from the optimal consumption of the investor. However, the characterization that the marginal utility is proportional to the pricing kernel at a given consumption level is valid only if the utility is differentiable at this consumption level. This observation renders the rationalization of investment choices by the expected utility theory non-trivial, as there are many portfolios for which the implied utility is not differentiable at all consumption levels, such as the purchase of options or capital guarantee products. Furthermore, in a discrete setting (with a finite number of equiprobable states), the utility function that is consistent with optimal consumption is not unique. In this context, Peleg and Yaari (1975) give one potential implied utility, but there are many others. In the presence of a continuum of states, when the pricing kernel is continuously distributed, we are able to derive the unique (up to a linear transformation) concave utility function that is implied by the optimal consumption of any investor who respects FSD.

The proof of our main results builds on Dybvig’s (1988a; 1988b) seminal work on portfolio selection. Instead of optimizing a value function, Dybvig (1988a) specifies a target distribution and solves for the strategy that generates the distribution at the lowest possible cost.¹ Here, we seek to infer preferences of consumers who are investing in the financial market. We show that if their portfolio satisfies some conditions, then it can be rationalized by expected utility theory with a non-decreasing and concave function. We assume that there is an infinite number of states in which it is possible to invest, whereas previous work considered a finite number of states. The assumption that we make is natural in the context of optimal portfolio selection problems. It allows us to obtain the uniqueness of the implied concave utility and to be able to use this inferred utility to estimate risk aversion. Inference of risk preferences from observed investment behavior has also been studied by Sharpe (2007) and Dybvig and Rogers (1997) for instance.² Sharpe (2007) assumes a static setting and relies on Dybvig’s (1988b) results to estimate the coefficient of constant relative risk aversion for a CRRA utility based on target distributions of final wealth.

In this chapter, we establish a link between expected utility theory (EUT) and all other theories that respects FSD. This connection can be used to estimate the agents’ utility functions and risk

¹It may indeed be more natural for an investor to describe her target distribution of terminal wealth instead of her utility function. For example, Goldstein, Johnson and Sharpe (2008) discuss how to estimate the distribution at retirement using a questionnaire. The pioneering work in portfolio selection by Markowitz (1952) is based solely on the mean and variance of returns and does not invoke utility functions. Black (1988) calls a utility function “a foreign concept for most individuals” and states that “instead of specifying her preferences among various gambles the individual can specify her consumption function”.

²Under some conditions, Dybvig and Rogers (1997) infer utility from dynamic investment decisions. Our setting is static and well adapted to the investment practice by which consumers purchase a financial contract and do not trade thereafter.

aversion coefficients in a non-expected utility setting. Our approach in doing so is non-parametric and is based solely on knowledge of the distribution of optimal wealth and of the financial market. This is in contrast with traditional approaches to inferring utility and risk aversion, which specify an exogenous parametric utility function in isolation of the market in which the agent invests and then calibrate this utility function using laboratory experiments and econometric analysis of panel data.

It is widely accepted that the Arrow-Pratt measure of absolute risk aversion is decreasing with wealth. This feature - i.e., decreasing absolute risk aversion (DARA) - is often the motivation for using the CRRA utility instead of the exponential utility to model investors' preferences. In this chapter, we show that the DARA property is completely characterized by a demand for final wealth W that exhibits more spread than a certain market variable (the opposite of the log pricing kernel). Our characterization of DARA can be used to empirically test DARA preferences based on observed investment decisions.

The chapter is organized as follows. The introductory example in Section 2.1 explains in a simplified setting why a distribution of terminal wealth can always be obtained as the optimum of the maximization of expected utility for a risk-averse agent. The general setting is presented in Section 2.2 with the strong connection between law-invariance and first-order stochastic dominance. Section 2.3 shows that *any* distribution of terminal wealth can be obtained as the optimum of an expected utility maximizer with non-decreasing and (possibly non strictly) concave utility. Section 2.4 provides some applications of the results derived in Section 2.3. In particular, we illustrate how a non-decreasing concave utility function can be constructed to explain the demand for optimal investments in Yaari's (1987) setting. In Section 2.5, we show how to derive the coefficients of risk aversion directly from the choice of the distribution of final wealth and the financial market. In this section, we also explore the precise connections between decreasing absolute risk aversion and the variability of terminal wealth. In Section 2.6, we derive new utilities corresponding to well-known distributions and discuss their properties.

2.1 Introductory Example

Throughout this chapter, we consider agents with law-invariant and non-decreasing preferences,³ $V(\cdot)$. We say that $V(\cdot)$ is *non-decreasing* if, for consumptions X and Y satisfying $X \leq Y$, one has that $V(X) \leq V(Y)$. We say that $V(\cdot)$ is *law-invariant* if $X \sim Y$ implies that $V(X) = V(Y)$, where " \sim " reflects equality in distribution. This is often referred to as a "state-independent" set of preferences. We also assume that the agent's initial budget is finite.

In this section, we present an example in order to introduce the notation and to explain in a simplified setting (a space with a finite number of equiprobable states) why a distribution of terminal wealth can always be obtained as the optimum of the maximization of expected utility for a risk averse agent. We will also show the limitations of this discrete setting and how it fails to identify *the* implied concave utility function and implied risk aversion of the investor.

The introductory example takes place in a finite state space $\Omega = \{\omega_1, \omega_2, \dots, \omega_N\}$ consisting of N equiprobable states (with probability $\frac{1}{N}$) at some terminal time T . Denote by $\frac{\xi(\omega_i)}{N}$ the initial (positive) cost at time 0 of the Arrow-Debreu security that pays one unit in the i^{th} state, ω_i , at

³This assumption is present in most traditional decision theories including the von Neumann and Morgenstern expected utility theory, Yaari's dual theory (Yaari (1987)), the cumulative prospect theory (Tversky and Kahneman (1992)) and rank dependent utility theory (Quiggin (1993)).

time T and zero otherwise. Let us call $\xi := (\xi_1, \xi_2, \dots, \xi_N)$ the pricing kernel where $\xi_i := \xi(\omega_i)$. It is clear that any state-contingent consumption $X := (x_1, x_2, \dots, x_N)$ (with $x_i := X(\omega_i)$) at time T writes as a linear combination of the N Arrow–Debreu securities.

The optimal investment problem of the agent with preferences $V(\cdot)$ is to find the optimal consumption X^* by solving the optimization problem,

$$\max_{X \mid E[\xi X] = X_0} V(X), \quad (2.1)$$

where the budget constraint $E[\xi X] = \frac{1}{N} \sum_{i=1}^N \xi_i x_i = X_0$ reflects that the agent's initial wealth level is X_0 . We assume that an optimum X^* to Problem (2.1) exists (it is always the case when restricting to non-negative consumptions). Denote by $x_i^* := X^*(\omega_i)$. Observe then that X^* and ξ must be anti-monotonic⁴; in other words, the outcomes for X^* are ordered in reverse of the ones for ξ , or

$$\forall \omega, \omega' \in \Omega, \quad (\xi(\omega) - \xi(\omega')) (X^*(\omega) - X^*(\omega')) \leq 0. \quad (2.2)$$

Let us prove (2.2) by contradiction. To this end, assume that there exist two states, ω_i and ω_j , such that $\xi_i > \xi_j$ and $x_i^* > x_j^*$. Let Y be another consumption such that $Y(\omega) = X^*(\omega)$ for all $\omega \in \Omega \setminus \{\omega_i, \omega_j\}$, and $Y(\omega_i) = x_j^*$, $Y(\omega_j) = x_i^*$. Since all states are equiprobable, $X \sim Y$ implies that $V(X) = V(Y)$ by the law-invariance of $V(\cdot)$. However, Y has a strictly lower cost, i.e.,

$$E[\xi X^*] - E[\xi Y] = \frac{1}{N} (\xi_i - \xi_j) (x_i^* - x_j^*) > 0.$$

Hence, X^* cannot be optimal as $V(X^*) < V(Z)$ in which $Z = Y + \frac{E[\xi X^*] - E[\xi Y]}{E[\xi]}$ (note that $E[\xi Z] = X_0$). Without loss of generality, using (2.2), we can thus assume that $x_1^* \leq x_2^* \leq \dots \leq x_N^*$ and $\xi_1 \geq \xi_2 \geq \dots \geq \xi_N$. For ease of exposition, we suppose in addition that the inequalities are strict (see Peleg and Yaari (1975) for the most general case):

$$x_1^* < x_2^* < \dots < x_N^*, \quad \xi_1 > \xi_2 > \dots > \xi_N.$$

Proposition 2.1.1 (Rationalizing Investment in a Discrete Setting). *The optimal solution of (2.1), denoted by X^* also solves the maximum expected utility problem*

$$\max_{X \mid E[\xi X] = X_0} E[U(X)] \quad (2.3)$$

for any concave utility $U(\cdot)$ such that the left derivative⁵ denoted by U' exists in x_i^* for all i and satisfies

$$\forall i \in \{1, 2, \dots, N\}, \quad U'(x_i^*) = \xi_i. \quad (2.4)$$

The proof of Proposition 2.1.1 can be found in Bernard, Chen, and Vanduffel (2015). It uses the concavity of U and is a basic application of pathwise optimization. It is clear from Proposition 2.1.1 that the utility function $U(\cdot)$ that rationalizes the optimal investment choice X^* in this setting is not unique. The candidate utility $U_P(\cdot)$ proposed by Peleg and Yaari (1975) is given by

$$U_P(x) = \int_0^x v(y) dy$$

⁴This observation appeared in Peleg and Yaari (1975) (as the principle of decreasing willingness in Theorem 1) and in Dybvig (1988a; 1988b).

⁵The results can also be written in terms of right derivative.

where $v(y) = \xi_1 - y + x_1^*$ for $y < x_1^*$ and $v(y) = \xi_j + (y - x_j^*) \frac{(\xi_{j+1} - \xi_j)}{(x_{j+1}^* - x_j^*)}$ for $y \in [x_j^*, x_{j+1}^*)$ ($j = 1, 2, \dots, N-1$) and $v(y) = \xi_N$ for $y \geq x_N^*$. $U_P(\cdot)$ is differentiable at all x_i^* and satisfies the condition $U'_P(x_i^*) = \xi_i$ of Proposition 2.1.1. It will be clear later that if we use this utility and generalize the financial market to have a continuum of states, then X^* no longer solves (2.3). In this chapter, we introduce⁶ a utility function $U(\cdot)$

$$U(x) = \int_{x_1^*}^x F_\xi^{-1}(1 - F(y))dy, \quad (2.5)$$

where $F(\cdot)$ is the distribution function (cdf) of X^* , $F_\xi(\cdot)$ is the cdf of ξ and the quantile function F_ξ^{-1} is defined as $F_\xi^{-1}(p) = \inf\{t \mid F_\xi(t) \geq p\}$ with the convention that $F_\xi^{-1}(0) = 0$. One observes that $y \mapsto F_\xi^{-1}(1 - F(y))$ is a decreasing right-continuous step function (taking the value ξ_1 for $y < x_1^*$, ξ_{i+1} when $y \in [x_i^*, x_{i+1}^*)$ for $i = 1, \dots, N-1$, and the value 0 for $y \geq x_N^*$). Thus, $U(\cdot)$ is continuous, piecewise linear and concave satisfying $U(x_1^*) = 0$. It is differentiable at all points except at each x_i^* , $i = 1, 2, \dots, N$, where it has a distinct right and left derivative. Denote by $U'(\cdot)$ the left derivative of $U(\cdot)$. We have that condition (2.4) of Proposition 2.1.1 is satisfied. Moreover, Theorem 2.3.2 will show that it is the unique (generalized) utility that explains the demand for X^* in the market when ξ is continuously distributed.

Henceforth, we assume that the state price ξ is continuously distributed, and we extend the above construction to the more general market (infinite state space) in Section 2.3. Doing so is not only a technical extension, but rather natural in the context of making optimal investment choices and at least consistent with the literature on it. As the example illustrates, it is also necessary in order to define *the* utility function (up to a linear transformation) that rationalizes the demand for a given optimal consumption X^* , and that can thus be used to compute the implied risk aversion of the investor.

When requiring that the utility function is differentiable at all points, we will show that we can only explain continuous distributions. However, there are many situations in which the investor wants a discrete distribution of wealth or a mixed distribution. He and Zhou (2011) show that, under some assumptions, optimal payoffs in Yaari's dual theory have a discrete distribution, whereas in the case of cumulative prospect theory, the optimal final wealth has a mixed distribution. While these observations point to differences between the decision theories, we show that these optimal payoffs can be rationalized by (generalized) expected utility theory. Section 2.4 provides some applications of the results derived in Section 2.3. In particular, we illustrate how a non-decreasing concave utility function can be constructed to explain the demand for optimal investment in Yaari's (1987) setting. One of the key findings presented toward the end of the chapter (Section 2.5) is to infer risk aversion and to show that DARA is equivalent to a demand for terminal wealth that exhibits more spread than the market variable $H_T := -\log(\xi_T)$.

2.2 Setting

We assume an arbitrage-free and frictionless financial market $(\Omega, \mathcal{F}, \mathbb{P})$ with a fixed investment horizon of $T > 0$. Let ξ_T be the pricing kernel that is agreed upon by all agents. We assume that

⁶The formula for $U(\cdot)$ in (2.5) will appear more intuitive after reading the proof of Theorem 2.3.1 in the following section. It is built so that the optimum for the expected utility problem (2.3) corresponds to the cheapest strategy with distribution F .

it has a positive density on $\mathbb{R}^+ \setminus \{0\}$. The value X_0 at time 0 of a consumption X_T at T is then computed as

$$X_0 = E[\xi_T X_T].$$

We consider only terminal consumptions X_T such that X_0 is finite. Throughout the chapter, agents have law-invariant and non-decreasing preferences $V(\cdot)$. Theorem 2.2.1 shows that these properties are equivalent to preferences $V(\cdot)$ that respect FSD. $X \prec_{fsd} Y$ means that Y is (first-order) stochastically larger than X , i.e., for all $x \in \mathbb{R}$, $F_X(x) \geq F_Y(x)$, where F_X and F_Y denote the cumulative distribution functions (cdfs) of X and Y respectively. Equivalently for all non-decreasing functions v , $E[v(X)] \leq E[v(Y)]$.

Theorem 2.2.1. *Preferences $V(\cdot)$ are non-decreasing and law-invariant if and only if $V(\cdot)$ satisfies FSD.*

An agent with preferences $V(\cdot)$ finds her optimal terminal consumption X_T^* by solving the following optimization problem:

$$\max_{X_T \mid E[\xi_T X_T] = X_0} V(X_T). \quad (2.6)$$

When an optimal strategy X_T^* exists, denote its cdf by F . Intuitively, since $V(\cdot)$ is non-decreasing and law-invariant, then among all strategies with cdf F , the optimum X_T^* must be the cheapest possible one. This observation is made precise in the following lemma, which is instrumental to the rest of the chapter.

Lemma 2.2.1 (Cost efficiency). *Assume that an optimum X_T^* of (2.6) exists and denote its cdf by F . Then, X_T^* is the cheapest (cost-efficient) way to achieve the distribution F at the investment horizon T , i.e., X_T^* also solves the following problem:*

$$\min_{X_T \mid X_T \sim F} E[\xi_T X_T]. \quad (2.7)$$

Furthermore, for any given cdf F , the solution X_T^* to Problem (2.7) is almost surely (a.s.) unique and writes as $X_T^* = F^{-1}(1 - F_{\xi_T}(\xi_T))$. Payoffs are cost-efficient if and only if they are non-increasing in the pricing kernel ξ_T .

This lemma is proved in Bernard, Boyle and Vanduffel (2014) (Chapter 1).⁷ It is also closely related to results that first appeared in Dybvig (1988a; 1988b).

Remark 2.2.1. The assumption of a continuous state space is needed to prove the equivalence between the cost efficiency of a payoff and its monotonicity in the pricing kernel ξ_T . In a continuous setting, $F_{\xi_T}(\xi_T)$ is uniformly distributed over $(0, 1)$, so that $F^{-1}(1 - F_{\xi_T}(\xi_T))$ has distribution F . In a discrete setting and assuming equiprobable states, Lemma 2.2.1 is proved by Dybvig (1988a; 1988b). It relies on the fact that $1 - F_{\xi_T}(\xi_T)$ has a uniform distribution over n states and therefore all (discrete) distributions F are attainable by $F^{-1}(1 - F_{\xi_T}(\xi_T))$. But when the states are not equiprobable, this is no longer valid. Specifically, any payoff that is non-increasing in the pricing kernel is cost-efficient but the reverse is incorrect. Counterexamples can easily be constructed. For instance, take $\Omega = \{\omega_0, \omega_1, \omega_3\}$ with $P(\omega_0) = \frac{1}{2}$ and $P(\omega_1) = P(\omega_2) = \frac{1}{4}$. The pricing kernel takes the values $\xi_T(\omega_0) = 2$, $\xi_T(\omega_1) = 3$ and $\xi_T(\omega_2) = 1$. Suppose that the desired distribution F is

⁷The first part of this lemma is related to Proposition 1.3.1 on First-order stochastic dominance. The second part corresponds to Proposition 1.2.4 on the characterization of cost-efficient payoffs and their Corollary 1.2.3 for the explicit expression of the unique cost-efficient payoff with distribution F .

given by $F(x) = 0$ for $x < 1$, $F(x) = \frac{1}{4}$ for $1 \leq x < 2$, $F(x) = \frac{1}{2}$ for $2 \leq x < 3$, and $F(x) = 1$ for $x \geq 3$. Then, the cost-efficient way of achieving F is given by assigning $X_T^*(\omega_0) = 3$, $X_T^*(\omega_1) = 1$ and $X_T^*(\omega_2) = 2$. However, X_T^* is not anti-monotonic with ξ_T . In general, assume there exists a state ω_0 with probability $\mathbb{P}(\omega_0)$ strictly different from all other state probabilities, and let $\mathbb{P}(\omega_0)$ be assigned to the highest desired payoff outcome. Whenever $\xi_T(\omega_0)$ is not the smallest state-price available, the cost-efficient strategy with the desired distribution cannot be anti-monotonic with the state-price ξ_T .

Lemma 2.2.1 provides us with an alternative approach to portfolio optimization. Usually, one resorts to a value function $V(\cdot)$ in order to model preferences and then finds the optimal consumption by solving Problem (2.6). Using Lemma 2.2.1, one specifies a desired distribution F of terminal wealth up-front⁸ and determines the cheapest strategy that is distributed with F .

In general, there may be more than one solution to (2.6). However, two different solutions must have different cdfs because the cost-efficient payoff generating a given distribution is unique. In the context of EUT, $V(X_T) = E(U(X_T))$ for some utility function $U(\cdot)$. When $U(\cdot)$ is not concave, a standard approach to solving Problem (2.6) is to introduce the concave envelope of $U(\cdot)$, denoted by $U_C(\cdot)$, which is the smallest concave function larger than or equal to $U(\cdot)$. Reichlin (2013) shows that under some technical assumptions, the maximizer for $U_C(\cdot)$ is also the maximizer for $U(\cdot)$. However, this maximizer is only unique under certain cases (see Lemma 5.9 of Reichlin (2013)).

In the following section, we reconcile different decision theories by showing that an optimal portfolio in any behavioral theory that respects FSD can be obtained as an optimal portfolio for a risk averse investor maximizing a (generalized) expected utility.

2.3 Explaining Distributions through Expected Utility Theory

In the first part of this section, we show how the traditional expected utility setting with a strictly increasing and strictly concave utility function on an interval can be used to explain F when it is strictly increasing and continuous on this interval. The second part shows that, more generally, any distribution of optimal final wealth can be explained using a “generalized” utility function defined over \mathbb{R} . Since a solution X_T^* to the optimization problem (2.6) is completely characterized by its distribution F , it follows that X_T^* is also optimal for some expected utility maximizer. The last part discusses tests for verifying whether investor behavior can be rationalized by the (generalized) Expected Utility Theory.

2.3.1 Standard Expected Utility Maximization

Definition 2.3.1 ($\mathcal{U}_{(a,b)}$: set of utility functions). *Let $(a,b) \subset \mathbb{R}$ where $a, b \in \overline{\mathbb{R}}$ (with $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, +\infty\}$). We define a set $\mathcal{U}_{(a,b)}$ of utility functions U on (a,b) such that $U : (a,b) \rightarrow \mathbb{R}$ is continuously differentiable, strictly increasing on (a,b) , U' is strictly decreasing on (a,b) (so that the investor is risk averse), $U(c) = 0$ for some $c \in (a,b)$, $U'(a) := \lim_{x \searrow a} U'(x) = +\infty$, and $U'(b) := \lim_{x \nearrow b} U'(x) = 0$.*

Note that Inada’s conditions correspond to $a = c = 0$ and $b = +\infty$.

⁸As aforementioned, it is presumably easier for many investors to describe a target terminal wealth distribution F than to articulate the value function $V(\cdot)$ governing their investment decision (see e.g., Goldstein, Johnson and Sharpe (2008)).

Definition 2.3.2 (Rationalization by Standard Expected Utility Theory). *An optimal portfolio choice X_T^* with a finite budget X_0 is rationalizable by the standard expected utility theory if there exists a utility function $U \in \mathcal{U}_{(a,b)}$ such that X_T^* is also the optimal solution to*

$$\max_{X_T \mid E[\xi X_T] = X_0} E[U(X_T)]. \quad (2.8)$$

The following lemma finds the optimal payoff for an expected utility maximizer with a utility function in $\mathcal{U}_{(a,b)}$.

Lemma 2.3.1. *Consider a utility function U in $\mathcal{U}_{(a,b)}$. Assume that $X_0 \in (E[\xi_T a], E[\xi_T b])$. The (a.s.) unique optimal solution X_T^* to the expected utility maximization (2.8) is given by*

$$X_T^* := [U']^{-1}(\lambda^* \xi_T)$$

where $\lambda^* > 0$ is such that $E[\xi_T X_T^*] = X_0$. Furthermore, X_T^* has a continuous distribution F , which is strictly increasing on (a, b) with $\lim_{x \searrow a} F(x) = 0$, $\lim_{x \nearrow b} F(x) = 1$.

This lemma is proved by Merton (1971) when Inada's conditions are satisfied. Note in particular that in the statement of the lemma, the condition on the budget automatically disappears when Inada's conditions are satisfied. The result is presented here in a slightly more general setting, as it will be needed in what follows. The following theorem gives, for any strictly increasing continuous distribution of final wealth, an explicit construction of the utility function that explains the investor's demand in the expected utility maximization framework.

Theorem 2.3.1 (Strictly increasing continuous distribution). *Consider a strictly increasing and continuous cdf F on $(a, b) \subset \mathbb{R}$ with $a, b \in \bar{\mathbb{R}}$. Assume that the cost of the unique cost-efficient payoff X_T^* solving (2.7) is finite and denote it by X_0 . Then X_T^* is also the optimal solution of the expected utility maximization problem (2.8) with the following explicit utility function $U \in \mathcal{U}_{(a,b)}$*

$$U(x) = \int_c^x F_{\xi_T}^{-1}(1 - F(y)) dy \quad (2.9)$$

for some c such that $F(c) > 0$. This utility U is unique in $\mathcal{U}_{(a,b)}$ up to a linear transformation.

If an expected utility maximizer chooses a particular investment with distribution of terminal wealth F , then the only utility function in $\mathcal{U}_{(a,b)}$ that rationalizes her choice is given by (2.9) (up to a linear transformation). Note also that this utility function (2.9) involves properties of the financial market at the horizon time T (through the cdf F_{ξ_T} of the pricing kernel ξ_T).

In the second part of this section, we generalize Theorem 2.3.1 to include more general distributions (discrete and mixed distributions). Obviously, any distribution F can always be approximated by a sequence of continuous strictly increasing distributions, F_n . Then, for each F_n , Theorem 2.3.1 allows us to obtain the corresponding strictly concave and strictly increasing utility function $U_n \in \mathcal{U}_{(a,b)}$ so that the optimal investment for an expected utility maximizer with utility function U_n is distributed with the cdf F_n . Thus, Theorem 2.3.1 already explains *approximately* the demand for all distributions.

Rationalization by Standard Expected Utility Theory

To summarize our findings in this section, we formulate the following characterization of consumptions that can be rationalizable by standard expected utility theory.

Corollary 2.3.1 (Rationalizable consumption by Standard EUT). *Consider a terminal consumption X_T at time T purchased with an initial budget X_0 and distributed with a continuous and strictly increasing cdf F . The following five conditions are equivalent:*

- (i) X_T is the solution to a maximum portfolio problem for some objective $V(\cdot)$ that satisfies FSD.
- (ii) X_T is non-increasing in ξ_T a.s.
- (iii) X_T is cost-efficient with cdf F .
- (iv) X_T is rationalizable by standard Expected Utility Theory.
- (v) X_T is the solution to a maximum portfolio problem for some objective $V(\cdot)$ that satisfies SSD.

Note that Corollary 2.3.1 highlights the strong link between the concept of “cost efficiency” and “rationalization” of consumption by Expected Utility Theory. Recall that $X \prec_{f_{sd}} Y$ means that for all non-decreasing utility functions $u(\cdot)$, $E[u(X)] \leq E[u(Y)]$, and $X \prec_{ssd} Y$ means that the expected utilities are ordered for concave and non-decreasing utility functions. It is thus clear that if V satisfies SSD, then it satisfies FSD so that (v) implies (i) trivially. Our results show that the converse is true and thus for the rationalization of consumption by Expected Utility Theory, it is only required that preferences satisfy FSD, which, unlike the case of SSD, is an assumption that is postulated by most common decision theories.

Remark 2.3.1. Corollary 2.3.1 cannot be generalized in a discrete setting when the states are not equiprobable. In particular, (i) is not equivalent to (iv).⁹ For example, take $\Omega = \{\omega_1, \omega_2\}$ with $P(\omega_1) = \frac{1}{3}$ and $P(\omega_2) = \frac{2}{3}$, and $\xi(\omega_1) = \frac{3}{4}$ and $\xi(\omega_2) = \frac{9}{8}$. Take a budget $X_0 = 1$ and consider X_T with $X_T(\omega_1) = a_1$ and $X_T(\omega_2) = a_2$ satisfying the budget condition $\frac{a_1}{4} + \frac{3a_2}{4} = 1$. Let the objective be defined as $V(X_T) := VaR_{1/3}^+(X_T) \mathbb{1}_{P(X_T < 0)=0}$ (where $VaR_{\alpha}^+(X_T)$ is defined as $VaR_{\alpha}^+(X_T) := \sup\{x \in \mathbb{R}, F_{X_T}(x) \leq \alpha\}$). Note that $V(\cdot)$ is law-invariant and non-decreasing. Thus, $V(\cdot)$ satisfies FSD. It is clear that $V(\cdot)$ is maximized for X_T^* defined through $X_T^*(\omega_1) = 0$ and $X_T^*(\omega_2) = \frac{4}{3}$. By contrast, we can show that X_T^* is never optimal for an expected utility maximizer with non-decreasing concave utility U on $[0, \frac{4}{3}]$ (range of consumption). To this end, assume without loss of generality that $U(0) = 0$ and $U(\frac{4}{3}) = 1$. Consider Y_T such that $Y_T(\omega_1) = \frac{4}{3}$ and $Y_T(\omega_2) = \frac{8}{9}$. Observe that $E[\xi Y_T] = E[\xi X_T^*] = 1$. By concavity of U , $\frac{U(\frac{8}{9}) - U(0)}{8/9} \geq \frac{U(\frac{4}{3}) - U(0)}{4/3}$, and thus $U(\frac{8}{9}) \geq \frac{2}{3}$. Hence, $E[U(Y_T)] = \frac{1}{3}U(\frac{4}{3}) + \frac{2}{3}U(\frac{8}{9}) \geq \frac{7}{9} > E[U(X_T^*)] = \frac{2}{3}$.

A related connection between cost efficiency and expected utility theory rationalization can be found in another context in Polissou and Quah (2013) and Green and Srivastava (1986). In these papers, the authors deal with sets of contingent consumption quantities and prices in a finite state-space setting and seek to rationalize this set of observations using expected utility theory. Here, we only observe one price and the contingent demand, X_T . We may use this observation to test whether or not an observed demand for a terminal consumption can be rationalized by expected utility theory.

⁹Using Proposition 2.1.1 and Lemma 2.2.1 formulated in the discrete setting, it is possible to prove a discrete version of Corollary 2.3.1 assuming that states are equiprobable. See also Remark 2.2.1.

Specifically, we use the equivalence between (iii) and (iv) of Corollary 2.3.1. Assume that an investor chooses to invest her initial budget X_0 from 0 to time T in such a way that he is generating a payoff X_T at time T . A simple test to investigate whether this investment choice is rationalizable by Expected Utility Theory consists in checking condition (iii) of Corollary 2.3.1, i.e., that $X_0 = E[\xi_T F^{-1}(1 - F_{\xi_T}(\xi_T))]$ as $F^{-1}(1 - F_{\xi_T}(\xi_T))$ is the a.s. unique cost-efficient strategy that provide F at maturity. Note that this test is based only on the cost of the payoff. If the distribution of returns is not obtained in the cheapest way, it may be caused by an optimal investment criterion that does not satisfy FSD. This is also a potential explanation for the pricing kernel puzzle (Hens and Reichlin (2013)).

2.3.2 Generalized Expected Utility Maximization

In the previous section, we used cost efficiency to construct a utility function that is continuously differentiable, strictly concave and strictly increasing on an interval (a, b) with $a, b \in \mathbb{R}$ to explain the demand for continuous and strictly increasing distributions on (a, b) . The same approach can be used to construct a “generalized” utility function defined on the entire real line \mathbb{R} , which explains the demand for *any* distribution. A generalized utility function does not need to be either differentiable on (a, b) or strictly concave. It is formally defined as follows.

Definition 2.3.3 ($\tilde{\mathcal{U}}_{(a,b)}$: set of generalized utility functions). *Let $(a, b) \subset \mathbb{R}$. We say that $\tilde{U} : \mathbb{R} \rightarrow \mathbb{R}$ belongs to the set of generalized utility functions $\tilde{\mathcal{U}}_{(a,b)}$ if $\tilde{U}(x)$ writes as*

$$\tilde{U}(x) := \begin{cases} U(x) & \text{for } x \in (a, b), \\ -\infty & \text{for } x < a, \\ \lim_{x \searrow a} U(x) & \text{for } x = a, \\ \lim_{x \nearrow b} U(x) & \text{for } x \geq b, \end{cases}$$

where $U : (a, b) \rightarrow \mathbb{R}$ is strictly increasing and concave on (a, b) . We then define \tilde{U}' (with abuse of notation) on \mathbb{R} as follows. On (a, b) , \tilde{U}' denotes the left derivative of \tilde{U} . For $x < a$, $\tilde{U}'(x) := +\infty$. $\tilde{U}'(a) := \lim_{x \searrow a} U'(x)$ and $\tilde{U}'(b) := \lim_{x \nearrow b} U'(x)$, and for $x > b$, $\tilde{U}'(x) = 0$. Conventions: if $a = -\infty$ then $\tilde{U}(a) := \lim_{x \searrow a} U(x)$ and $\tilde{U}'(a) := +\infty$, if $b = \infty$ then $\tilde{U}(b) := \lim_{x \nearrow b} U(x)$ and $\tilde{U}'(b) := 0$.

Definition 2.3.4 (Rationalization by Generalized Expected Utility Theory). *An optimal portfolio choice X_T^* with a finite budget X_0 is rationalizable by the generalized expected utility theory if there exists a utility function $\tilde{U} \in \tilde{\mathcal{U}}_{(a,b)}$ such that X_T^* is also the optimal solution to*

$$X_T \mid \max_{E[\xi_T X_T] = X_0} E[\tilde{U}(X_T)]. \quad (2.10)$$

The following lemma finds the optimal payoff for a generalized expected utility maximizer with a utility function in $\tilde{\mathcal{U}}_{(a,b)}$.

Lemma 2.3.2. *Consider a generalized utility function \tilde{U} of $\tilde{\mathcal{U}}_{(a,b)}$. Assume that $X_0 \in (E[\xi_T a], E[\xi_T b])$. The optimal solution X_T^* to the generalized expected utility maximization (2.10) exists, is a.s. unique and is given by*

$$X_T^* := [\tilde{U}']^{-1}(\lambda^* \xi_T),$$

where \tilde{U}' is as defined in Definition 2.3.3 and where $\lambda^* > 0$ is such that $E[\xi_T X_T^*] = X_0$ and where

$$\left[\tilde{U}'\right]^{-1}(y) := \inf \left\{ x \in (a, b) \mid \tilde{U}'(x) \leq y \right\}, \quad (2.11)$$

with the convention that $\inf\{\emptyset\} = b$. Furthermore, X_T^* may have mass points.

Lemma 2.3.2 allows us to derive a unique implied generalized expected utility to rationalize the demand for any distribution. This is done in Theorem 2.3.2.

Theorem 2.3.2. *Let F be a distribution. Let X_T^* be the optimal solution to (2.7) for the cdf F . Denote its cost by X_0 and assume that it is finite. Then, X_T^* is also the optimal solution to Problem (2.10) where $\tilde{U} : \mathbb{R} \rightarrow \bar{\mathbb{R}}$ is a generalized utility function defined as*

$$\tilde{U}(x) := \int_c^x F_{\xi_T}^{-1}(1 - F(y)) dy \quad (2.12)$$

with some $c \geq a$ such that $F(c) > 0$. Conventions: $F_{\xi_T}^{-1}(1) = +\infty$, $F_{\xi_T}^{-1}(0) = 0$, if $x_1 < x_2$ then $\int_{x_1}^{x_2} (+\infty) dy = +\infty$, and $\int_{x_2}^{x_1} (+\infty) dy = -\infty$ and $\int_{x_1}^{x_1} g(y) dy = 0$ for all g valued in $\bar{\mathbb{R}}$. $\tilde{U}(\cdot)$ is unique in the class of generalized utilities up to a linear transformation.

Rationalization by Generalized Expected Utility Theory

We can summarize our findings in this section by the following corollary, which is similar to Corollary 2.3.1 but now includes all possible distributions of final wealth.

Corollary 2.3.2. *Consider a terminal consumption X_T at time T purchased with an initial budget X_0 and distributed with F . The following five conditions are equivalent:*

- (i) X_T is the solution to a maximum portfolio problem for some objective $V(\cdot)$ that satisfies FSD.
- (ii) X_T is non-increasing in ξ_T a.s.
- (iii) X_T is cost-efficient with cdf F .
- (iv) X_T is rationalizable by Generalized Expected Utility Theory.
- (v) X_T is the solution to a maximum portfolio problem for some objective $V(\cdot)$ that satisfies SSD.

The proof of this corollary is identical to that of Corollary 2.3.1 by replacing expected utility with generalized expected utility and is thus omitted.

2.4 From Distributions to Utility Functions

In this section, we use the results of the previous sections to derive utility functions that explain the demand for some financial products and for some distributions of final wealth for agents with preferences satisfying FSD. Let us start with a simple example showing how one can recover the popular CRRA utility function from a lognormally distributed final wealth. This example is particularly useful when explaining the optimal demand for a retail investor who chooses an equity-linked structured product with capital guarantee. The last example deals with an example of non-expected

utility: Yaari's dual theory of choice. In this case, we are able to exhibit the non-decreasing concave utility function such that the optimal strategy in Yaari's (1987) theory is also obtained in an expected utility maximization framework.

For ease of exposition, we restrict ourselves to the one-dimensional Black-Scholes model with one risky asset, S_T .¹⁰ In this case, the pricing kernel ξ_T is unique and can be expressed explicitly in terms of the stock price S_T as in (1.8) from Chapter 1 and has a lognormal distribution.

Assume first that consumption is restricted on $(0, \infty)$ and that the investor wants to achieve a lognormal distribution $\mathcal{LN}(M, \Sigma^2)$ at maturity T of her investment. The desired cdf is $F(x) = \Phi\left(\frac{\ln x - M}{\Sigma}\right)$, and from (1.8) it follows that $F_{\xi_T}^{-1}(y) = \exp\left\{\Phi^{-1}(y)\theta\sqrt{T} - rT - \frac{\theta^2 T}{2}\right\}$. Applying Theorem 2.3.1, the utility function explaining this distribution writes as

$$U(x) = \begin{cases} a \frac{x^{1-\frac{\theta\sqrt{T}}{\Sigma}}}{1-\frac{\theta\sqrt{T}}{\Sigma}} & \frac{\theta\sqrt{T}}{\Sigma} \neq 1 \\ a \log(x) & \frac{\theta\sqrt{T}}{\Sigma} = 1 \end{cases}, \quad (2.13)$$

where $a = \exp\left(\frac{M\theta\sqrt{T}}{\Sigma} - rT - \frac{\theta^2 T}{2}\right)$. This is a CRRA utility function with relative risk aversion $\frac{\theta\sqrt{T}}{\Sigma}$. A more thorough treatment of risk aversion is provided in Section 2.5.

2.4.1 Explaining the Demand for Capital Guarantee Products

Many structured products include a capital guarantee and have a payoff of the form $Y_T = \max(G, S_T)$, where S_T is the stock price and G is the (deterministic) guaranteed level. S_T has a lognormal distribution, $S_T \sim \mathcal{LN}(M, \sigma^2 T)$, where $M := \ln S_0 + \left(\mu - \frac{\sigma^2}{2}\right)T$ so that the cdf for Y_T is equal to $F_{Y_T}(y) = \mathbb{1}_{y \geq G} \Phi\left(\frac{\ln y - M}{\sigma\sqrt{T}}\right)$. Since Y_T has a mixed distribution (with mass point at G), we can apply Theorem 2.3.2 to derive the corresponding utility function. Let $p := \Phi\left(\frac{\ln G - M}{\sigma\sqrt{T}}\right)$ and define the following discrete and continuous cdfs:

$$F_{Y_T}^D(y) = \begin{cases} 0 & y < G \\ 1 & y \geq G \end{cases}, \quad F_{Y_T}^C(y) = \begin{cases} 0 & y < G \\ \frac{\Phi\left(\frac{\ln y - M}{\sigma\sqrt{T}}\right) - p}{1-p} & y \geq G \end{cases}. \quad (2.14)$$

Then, we can see that $F_{Y_T}(y) = pF_{Y_T}^D(y) + (1-p)F_{Y_T}^C(y)$. When $x < G$, $\tilde{U}(x) = -\infty$. The utility function \tilde{U} belongs to $\tilde{\mathcal{U}}_{(G, \infty)}$ and is given by

$$\tilde{U}(x) = \begin{cases} -\infty & x < G, \\ a \frac{x^{1-\frac{\theta}{\sigma}} - G^{1-\frac{\theta}{\sigma}}}{1-\frac{\theta}{\sigma}} & x \geq G, \frac{\theta}{\sigma} \neq 1, \\ a \log\left(\frac{x}{G}\right) & x \geq G, \frac{\theta}{\sigma} = 1, \end{cases} \quad (2.15)$$

with $a = \exp\left(\frac{M\theta}{\sigma} - rT - \frac{\theta^2 T}{2}\right)$. The mass point is explained by a utility that is infinitely negative for any level of wealth below the guaranteed level. The CRRA utility above this guaranteed level ensures the optimality of a lognormal distribution above the guarantee, as aforementioned.

¹⁰All developments can be executed in the general market setting given in Section 2.2. However, closed-form solutions are more complicated or unavailable.

2.4.2 Yaari's Dual Theory of Choice Model

Optimal portfolio selection under Yaari's (1987) dual theory involves maximizing the expected value of the terminal payoff under a distorted probability function. Specifically, under Yaari's dual theory of choice, decision makers evaluate the "utility" of their non-negative final wealth X_T (with cdf F) by calculating its distorted expectation $\mathbb{H}_w[X_T]$:

$$\mathbb{H}_w[X_T] = \int_0^\infty w(1 - F(x)) dx, \quad (2.16)$$

where the (distortion) function $w : [0, 1] \rightarrow [0, 1]$ is non-decreasing with $w(0) = 0$ and $w(1) = 1$.

The investor's initial endowment is $X_0 \geq 0$. He and Zhou (2011) find the optimal payoff when the distortion function is given by $w(z) = z^\gamma$ where $\gamma > 1$. They show that there exists¹¹ c such that the optimal payoff is equal to

$$X_T^* = B \mathbb{1}_{\xi_T \leq c}, \quad (2.17)$$

where $B = X_0 e^{rT} \Phi\left(\frac{\ln c + rT - \frac{\theta^2 T}{2}}{\theta\sqrt{T}}\right) > 0$ is chosen to satisfy the budget constraint $E[\xi_T X_T^*] = X_0$. The corresponding cdf is

$$F(x) = \begin{cases} \Phi\left(\frac{-rT - \frac{\theta^2 T}{2} - \ln c}{\theta\sqrt{T}}\right) & 0 \leq x < B \\ 1 & x \geq B \end{cases}. \quad (2.18)$$

We find that the utility function $\tilde{U} \in \tilde{\mathcal{U}}_{(0,B)}$ is given by

$$\tilde{U}(x) = \begin{cases} -\infty & x < 0, \\ c(x - c) & 0 \leq x \leq B, \\ c(B - c) & x > B. \end{cases} \quad (2.19)$$

The utility function such that the optimal investment in the expected utility setting is similar to the optimum in Yaari's framework (1987) is simply linear up to a maximum $c(B - c)$ and then constant thereafter.

Remark 2.4.1. He and Zhou (2011) derive the form of the optimal solution when the distortion is convex (characterizing risk aversion in Yaari's dual theory). It is a digital option, and thus always has mass points. As a consequence, buying a stock is never optimal in Yaari's dual theory, as it yields a terminal wealth that is lognormally distributed. We have shown that the (generalized) concave expected utility maximization can rationalize any investment in a law invariant increasing setting. By contrast, Yaari's dual theory with a convex distortion cannot explain all decisions that are optimal in some law invariant and increasing setting.

In the setting of the rank-dependent utility theory (Quiggin (1993)) and the cumulative prospect theory (Tversky and Kahneman (1992)), the expressions for the optimal portfolios are more complicated and beyond the scope of this chapter. Both settings include EUT as a limiting case (when there is no probability distortion and the reference level of wealth is zero). Therefore, in both settings, we would be able to rationalize all distributions of wealth, but possibly in a non unique way.

¹¹It is proved that c is the unique root on $(1, \gamma e^{-rT})$ of the function $\left[\Phi\left(\frac{\ln x + rT + \frac{\theta^2 T}{2}}{\theta\sqrt{T}}\right)\right]^{\gamma-1} \times \left[x\Phi\left(\frac{\ln x + rT + \frac{\theta^2 T}{2}}{\theta\sqrt{T}}\right) - \gamma e^{-rT} \Phi\left(\frac{\ln x + rT + \frac{\theta^2 T}{2}}{\theta\sqrt{T}}\right)\right]$.

2.5 From Distributions to Risk Aversion

Let us first recall the definition of the classical Arrow-Pratt coefficients of risk aversion.

Definition 2.5.1 (Arrow-Pratt Risk Aversion Coefficients). *consider an expected utility maximizer with a twice differentiable utility function U . His absolute risk aversion coefficient $A(x)$ at the wealth level x and corresponding relative risk aversion coefficient $R(x)$ are respectively defined as*

$$A(x) = -\frac{U''(x)}{U'(x)}, \quad R(x) = -\frac{xU''(x)}{U'(x)}.$$

In this section, we propose to define risk aversion coefficients directly from the choice of the distribution desired by the investor and from the distribution of the log pricing kernel (without specifying a utility function). We then show that this definition coincides with the Arrow-Pratt measures for risk aversion when the distribution is continuous (Theorem 2.5.1). Next, we show that decreasing absolute risk aversion (DARA) is equivalent to terminal wealth exhibiting a spread greater than the market variable

$$H_T := -\log(\xi_T).$$

In a Black-Scholes setting, agents thus have DARA preferences if and only if they show a demand for distributions with tails that are “fatter than normal.” Our characterization for DARA also allows us to construct an empirical test for DARA preferences that is based on observed investor behavior.

We denote by G the cdf of H_T and by g its density, and we assume that $g(x) > 0$ for all $x \in \mathbb{R}$.

2.5.1 Risk Aversion Coefficient

We propose the following definitions of the distributional absolute and relative risk aversion coefficients using solely the distribution of H_T and the target distribution F .

Definition 2.5.2 (Distributional Risk Aversion Coefficient). *Let F be the distribution desired by the investor. Let G and g be the cdf and density of $-\log(\xi_T)$, respectively. Consider a level of wealth x such that $x = F^{-1}(p)$ for some $0 < p < 1$. We define the distributional absolute and relative risk aversion at x as*

$$\mathcal{A}(x) = \frac{f(F^{-1}(p))}{g(G^{-1}(p))}, \quad \mathcal{R}(x) = F^{-1}(p) \frac{f(F^{-1}(p))}{g(G^{-1}(p))} \quad (2.20)$$

where

$$f(y) := \lim_{\varepsilon \rightarrow 0} \frac{F(y + \varepsilon) - F(y)}{\varepsilon} \quad (2.21)$$

when this limit exists. If F is differentiable at y , then $f(y) = F'(y)$. If F has a density, then $F'(y) = f(y)$ at all y .

The next theorem shows that the above definitions of the distributional absolute and relative risk aversion are consistent with the Arrow-Pratt measures for absolute and relative risk aversion recalled in Definition 2.5.1.

Theorem 2.5.1 (Arrow-Pratt measures for risk aversion). *Consider an investor who targets a cdf F for her terminal wealth (with corresponding f defined as (2.21)). Let U obtained in (2.9) be the utility that rationalizes F . Assume that F and G are such that U is twice-differentiable. Consider a level of wealth x in the interior of the support of the distribution F . Then $x = F^{-1}(p)$ for some $0 < p < 1$. The measures for the distributional absolute and relative risk aversion at x defined in (2.20) coincide with the Arrow-Pratt measures:*

$$\mathcal{A}(x) = A(x), \quad \mathcal{R}(x) = R(x) \quad (2.22)$$

Expression (2.22) shows that the coefficient for absolute risk aversion can be interpreted as a likelihood ratio and is linked directly to the financial market (through the cdf G of the negative of the log pricing kernel, H_T).

2.5.2 Decreasing Absolute Risk Aversion

We provide precise characterizations of DARA in terms of distributional properties of the final wealth and of the financial market. In what follows, we only consider distributions F and G that are twice differentiable.

Theorem 2.5.2 (Distributional characterization of DARA). *Consider an investor who targets some distribution F for her terminal wealth. The investor has (strictly) decreasing absolute risk aversion (DARA) if and only if*

$$y \mapsto F^{-1}(G(y)) \text{ is strictly convex on } \mathbb{R}. \quad (2.23)$$

The investor has asymptotic DARA (DARA for a sufficiently high level of wealth) if and only if there exists $y^ \in \mathbb{R}$ such that $y \mapsto F^{-1}(G(y))$ is strictly convex on (y^*, ∞) .*

The convexity of the function $F^{-1}(G(x))$ reflects the fact that the target distribution F is “fatter-tailed” than the distribution G . In other words, F is larger than G in the sense of transform convex order (Shaked and Shantikumar (2007), p. 214).

Theorem 2.5.2 extends recent results by Dybvig and Wang (2012) in another direction. These authors show that if agent A has lower risk aversion than agent B , then agent A purchases a distribution that is larger than the other in the sense of SSD. Here, we show that the risk aversion of an agent is decreasing in available wealth if and only if the agent purchases a payoff that is heavier-tailed than the market variable H_T .

Theorem 2.5.3. *Consider an investor with optimal terminal wealth $W_T \sim F$. The investor has decreasing absolute risk aversion if and only if W_T is increasing and strictly convex in H_T .*

Note that an investor with optimal terminal wealth $W_T \sim F$ such that F has right-bounded support does not exhibit DARA.¹²

¹²Recall that $k(x) := F^{-1}(G(x))$ is non-decreasing. Since F is right-bounded, there exists $b \in \mathbb{R}$, $k(x) \leq b$ ($x \in \mathbb{R}$). We need to show that $k(x)$ cannot be strictly convex ($x \in \mathbb{R}$). We proceed by contradiction, so let $k(x)$ be strictly convex. This implies that $k(x)$ remains above its tangent. As it is non-constant and non-decreasing, there exists a point x for which the tangent has a positive slope and thus goes to infinity at infinity. It is thus impossible for k to be bounded from above.

2.5.3 Case of a Black-Scholes Market

In a Black-Scholes market, H_T is normally distributed, with mean $rT + \frac{\theta^2 T}{2}$ and variance $\theta^2 T$. It is then straightforward to compute

$$\mathcal{A}(x) = \theta f(x) \sqrt{2\pi T} \exp\left(\frac{1}{2} [\Phi^{-1}(1 - F(x))]^2\right), \quad \mathcal{R}(x) = \theta x f(x) \sqrt{2\pi T} \exp\left(\frac{1}{2} [\Phi^{-1}(1 - F(x))]^2\right).$$

The financial market thus influences the risk aversion coefficient $\mathcal{A}(x)$ through the instantaneous Sharpe ratio $\theta = \frac{\mu - r}{\sigma}$. Interestingly, the effect of the financial market on the risk aversion coefficient $\mathcal{A}(x)$ is proportional and does not depend on available wealth x . This also implies that in a Black-Scholes market, the properties of the function $x \rightarrow \mathcal{A}(x)$ are solely related to distributional properties of final wealth and do not depend on particular market conditions. This observation is implicit in the following theorem.

Corollary 2.5.1 (DARA in a Black-Scholes market). *Consider an investor who targets some cdf F for her terminal wealth. In a Black-Scholes market the investor has decreasing absolute risk aversion if and only if*

$$\frac{f(F^{-1}(p))}{\phi(\Phi^{-1}(p))} \text{ is strictly decreasing on } (0, 1) \quad (2.24)$$

or, equivalently,

$$F^{-1}(\Phi(x)) \text{ is strictly convex on } \mathbb{R}, \quad (2.25)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the density and the cdf of a standard normal distribution.

When F is the distribution of a normal random variable, the ratio (2.24) becomes constant. This confirms the well-known fact that the demand for a normally distributed final wealth in a Black-Scholes market is tied to a constant absolute risk aversion (see Section 2.6.1 for a formal proof). In a Black-Scholes setting, the property of DARA remains invariant to changes in the financial market. This is, however, not true in a general market, where risk aversion, demand for a particular distribution and properties of the financial market are intertwined.

If the final wealth W has a lognormal distribution F , then one has $F^{-1}(\Phi(x)) = \exp(x)$. Since the exponential function is clearly convex, Theorem 2.5.1 implies that, in a Black-Scholes market, the demand for a lognormal distribution corresponds to DARA preferences. One can also readily show that the exponential distribution also corresponds to DARA preferences in this setting. Effectively, from (2.25), we only need to show that the survival function $1 - \Phi(x)$ of a standard normal random variable is log-concave.¹³

Consider for all $x \in \mathbb{R}$ such that $F(x) < 1$, the hazard function $h(x) := \frac{f(x)}{1 - F(x)}$, which is a useful device for studying heavy-tailed properties of distributions. The hazard function for an exponentially distributed random variable rate function is clearly constant, so that a non-increasing hazard function reflects a distribution that is heavier-tailed than an exponentially distributed random variable. It is therefore intuitive that investors exhibiting a demand for distributions with a non-increasing hazard function exhibit DARA preferences. The following theorem makes this precise.

¹³This is well-known in the literature and can be seen as a direct consequence of a more general result attributed to Prékopa (1973), who shows that the differentiability and log-concavity of the density implies log-concavity of the corresponding distribution and survival function. It is clear that a normal density is log-concave and thus also its survival function.

Theorem 2.5.4 (Sufficient conditions for DARA). *Consider an investor who targets some cdf F for her terminal wealth. Denote its density by f . If the hazard function $h(x)$ is non-increasing (resp. non-increasing for $x > k$) or, equivalently, if $1 - F(x)$ is log-convex (resp. log-convex for $x > k$), then the investor has decreasing absolute risk aversion (resp. asymptotic DARA).*

We remark from the proof of Theorem 2.5.4 (see details in Bernard, Chen, and Vanduffel (2015)) that a random variable with non-increasing hazard function $h(x)$ must assume values that are almost surely in an interval $[a, \infty)$ where $a \in \mathbb{R}$. A lognormally distributed random variable thus has no non-increasing hazard function (but still satisfies the DARA property).

2.6 Distributions and Corresponding Utility Functions

We end this chapter with a few examples illustrating the correspondence between the distribution of final wealth and the utility function. We use the same setting as the examples in Section 2.4. In particular, the pricing kernel is given by (1.8) in Chapter 1.

2.6.1 Normal Distribution and Exponential Utility

Let F be the distribution of a normal random variable with mean M and variance Σ^2 . Then, from Theorem 2.3.1, the utility function explaining this distribution writes as

$$u(x) = -\frac{\Sigma a}{\theta \sqrt{T}} \exp\left(-\frac{x\theta}{\Sigma} \sqrt{T}\right) \quad (2.26)$$

where $a = \exp\left(\frac{M\theta\sqrt{T}}{\sigma} - rT - \frac{\theta^2 T}{2}\right)$ and $\theta = \frac{\mu-r}{\sigma}$. This is essentially the form of an exponential utility function with constant absolute risk aversion

$$\mathcal{A}(x) = \frac{\theta \sqrt{T}}{\Sigma}.$$

Note that the absolute risk aversion is constant and inversely proportional to the volatility Σ of the distribution. A higher volatility of the optimal distribution of final wealth corresponds to a lower absolute risk aversion. This is consistent with Dybvig and Wang (2012), who show that lower risk aversion leads to a larger payoff in the sense of SSD.¹⁴ Reciprocally, consider the following exponential utility function defined over \mathbb{R} ,

$$U(x) = -\exp(-\gamma x), \quad (2.27)$$

where γ is the risk aversion parameter and $x \in \mathbb{R}$. This utility has constant absolute risk aversion $\mathcal{A}(x) = \gamma$. The optimal wealth obtained with an initial budget X_0 is given by

$$X_T^* = X_0 e^{rT} - \frac{\theta}{\gamma \sigma} \left(r - \frac{\sigma^2}{2}\right) T + \frac{\theta}{\gamma \sigma} \ln\left(\frac{S_T}{S_0}\right), \quad (2.28)$$

where $\theta = \frac{\mu-r}{\sigma}$ is the instantaneous Sharpe ratio for the risky asset S , and thus X_T^* follows a normal distribution $\mathcal{N}\left(X_0 e^{rT} + \frac{\theta}{\gamma \sigma} (\mu - r)T, \left(\frac{\theta}{\gamma}\right)^2 T\right)$.

¹⁴In case of two normal distributions with equal mean, increasing SSD is equivalent to increasing variance.

2.6.2 Lognormal Distribution and CRRA and HARA Utilities

The HARA utility is a generalization of the CRRA utility and is given by

$$U(x) = \frac{1-\gamma}{\gamma} \left(\frac{ax}{1-\gamma} + b \right)^\gamma, \quad (2.29)$$

where $a > 0$, $b + \frac{ax}{1-\gamma} > 0$. For a given parametrization, this restriction puts a lower bound on x if $\gamma < 1$ and an upper bound on x when $\gamma > 1$. If $-\infty < \gamma < 1$, then this utility displays DARA. In the case $\gamma \rightarrow 1$, the limiting case corresponds to a linear utility. In the case $\gamma \rightarrow 0$, the utility function becomes logarithmic: $U(x) = \log(x + b)$. Its absolute risk aversion is

$$\mathcal{A}(x) = a \left(\frac{ax}{1-\gamma} + b \right)^{-1}. \quad (2.30)$$

The optimal wealth obtained with an initial budget X_0 is given by

$$X_T^* = C \left(\frac{S_T}{S_0} \right)^{\frac{\theta}{\sigma(1-\gamma)}} - \frac{b(1-\gamma)}{a}$$

where $C = \frac{X_0 e^{rT} + \frac{b(1-\gamma)}{a}}{\exp\left(\frac{\theta}{\sigma(1-\gamma)}\left(r - \frac{\sigma^2}{2}\right)T + \left(\frac{\theta}{1-\gamma}\right)^2 \frac{T}{2}\right)}$. Its cdf is

$$F_{HARA}(y) = \Phi \left(\frac{\ln \left(\frac{y + \frac{b(1-\gamma)}{a}}{C} \right) + \frac{\theta}{\sigma(\gamma-1)} \left(\mu - \frac{\sigma^2}{2} \right) T}{\frac{\theta}{\gamma-1} \sqrt{T}} \right).$$

Observe that the optimal wealth for a HARA utility is a lognormal distribution translated by a constant term.

When $b = 0$, the HARA utility reduces to CRRA and the optimal wealth is LogNormal $\mathcal{LN}(M, \Sigma^2)$ over \mathbb{R}^+ . We showed in Section 2.4 that the utility function explaining this distribution is a CRRA utility function. It has decreasing absolute risk aversion

$$\mathcal{A}(x) = \frac{\theta\sqrt{T}}{x\Sigma} \text{ where } \theta = \frac{\mu - r}{\sigma}.$$

2.6.3 Exponential Distribution

Consider an exponential distribution with cdf $F(x) = 1 - e^{-\lambda x}$ where $\lambda > 0$. The utility function in (2.9) in Theorem 2.3.1 cannot be obtained in closed-form and does not correspond to a well-known utility. The coefficient of absolute risk aversion is given by

$$\mathcal{A}(x) = \theta\lambda\sqrt{2\pi T} \exp \left(-\lambda x + \frac{1}{2} \left[\Phi^{-1} \left(e^{-\lambda x} \right) \right]^2 \right).$$

We have already shown that $\mathcal{A}(x)$ is decreasing and thus that the exponential distribution corresponds to a utility function exhibiting DARA.

2.6.4 Pareto Distribution

Consider a Pareto distribution with scale $m > 0$ and shape $\alpha > 0$, defined on $[m, +\infty)$. Its pdf is $f(x) = \alpha \frac{m^\alpha}{x^{\alpha+1}} \mathbb{1}_{x \geq m}$. The coefficient of absolute risk aversion for $x \geq m$ is given by

$$\mathcal{A}(x) = \frac{\theta \alpha m^\alpha \sqrt{2\pi T}}{x^{\alpha+1}} \exp \left(\frac{1}{2} \left[\Phi^{-1} \left(\frac{m^\alpha}{x^\alpha} \right) \right]^2 \right).$$

Bergstrom and Bagnoli (2005) show that $1 - F(x)$ is log-convex. Theorem 2.5.4 thus implies that the Pareto distribution corresponds to DARA preferences.

2.7 Conclusions

When investors' preferences satisfy FSD, we are able to construct a concave utility function such that the optimal portfolio for the corresponding expected utility maximizer coincides with the investor's optimal portfolio when optimizing his preferences. This result can be used for non-parametric estimation of the utility function and of the risk aversion of an investor. This construction is indeed based solely on knowledge of the distribution of her final wealth (when optimally investing) and of the distribution of the pricing kernel. Another application consists in inferring the utility function that rationalizes optimal investment choice in non-expected utility frameworks (e.g., cumulative prospect theory or rank dependent utility theory). An application in a dynamic setting is given in Bernard and Kwak (2016a).

A further research direction could be to use this theory to better understand the demand for skewed distributions. Specifically, we show in this chapter that the exponential utility can rationalize the demand for a normal distribution in a Black-Scholes market. It would be of interest to understand what deformation of the exponential utility can explain the demand for a skewed normal distribution. Such an analysis could help to understand the properties of the utility function that explain the demand for positive or negative skewness. The link between the properties of the distribution and the utility have been recently studied in the economics literature by looking at lotteries and comparing lotteries (see for instance Denuit and Eeckhoudt (2010)). We would then be able to discuss how the presence of a financial market affects the existing results on the connection between the properties of the utility functions and the skewness preferences (Kraus and Litzenberger (1976), Machina (1982), Brockett and Kahane (1992), Mitton and Vorkink (2007), Chiu (2010)).

Our results suggest inherent deficiencies in portfolio selection within a decision framework that satisfies FSD (equivalently with a law-invariant non-decreasing objective function). When the market exhibits a positive risk premium, Bernard, Boyle and Vanduffel (2014) show that for every put contract there exists a derivative contract that yields the same distribution at a strictly lower cost but without protection against a market crash. Then, the demand for standard put contracts can only be rationalized in a decision framework that violates FSD and in which background risk as a source of state-dependent preferences comes into play. Similarly, insurance contracts provide protection against certain losses and provide a payout when it is needed. For this reason, they often appear more valuable to customers than (cheaper) financial payoffs with the same distribution (Bernard and Vanduffel (2014a)). Thus, it would be useful to seek to develop decision frameworks in which FSD can be violated, for example considering ambiguity on the pricing kernel, Almost Stochastic Dominance (ASD) (Levy (2006), chapter 13) or state-dependent preferences (Bernard, Boyle and Vanduffel (2014), Chabi-Yo, Garcia and Renault (2008)).

Chapter 3

Financial Bounds for Insurance Claims

In this chapter, we develop an application of cost-efficiency in insurance that was published in the *Journal of Risk and Insurance* (Bernard and Vanduffel (2014a)). In this chapter, insurance claims are priced using an indifference pricing principle. We first revisit the traditional economic framework and then extend it to incorporate a financial (sub)market as a tool to invest and to (partially) hedge. In this context, we derive lower bounds for claims' prices, and these bounds correspond to the market prices of some explicitly known financial payoffs. In particular, we show that the discounted expected value is no longer valid as a classical lower bound for insurance prices in general: it has to be corrected by a covariance term which reflects the interaction between the insurance claim and the financial market. An example of equity-linked insurance contract illustrates the chapter. All proofs can be found in the appendix of the published version of this chapter in *Journal of Risk and Insurance* (Bernard and Vanduffel (2014a)).

The valuation of insurance claims is at the core of actuarial science. The traditional actuarial premium principle is based on a quantity such as the expectation, the standard deviation, the variance, the quantile or any other quantity derived from the claim distribution under the physical probability. A second approach consists of specifying a set of reasonable properties that the premium principle should satisfy. Such approach is intimately connected with the axiomatic approach to risk measures, see Artzner et al. (1999). A third approach incorporates the preferences of the decision makers involved (i.e. the insurance buyer and insurance seller) in the determination of insurance prices. Such premia are then typically derived from economic indifference principles (using for example the expected utility theory from von Neumann and Morgenstern (1953), see also the zero-utility premium principle proposed by Bühlmann (1980)). We refer to Young (2004) for a review of these three approaches.

As Brockett et al. (2005) note, a “striking feature of the actuarial valuation principles is that they are formulated within a framework that generally ignores the financial market.” Indeed, the different approaches proposed in the literature for pricing insurance claims usually assume that, apart from the availability of a risk-free bond, there is no financial market and even if there is one it cannot be used to hedge insurance claims and to determine insurance premia. However, it is now clear that insurance claims should be priced by taking into account the financial market. First, life insurance contracts often include financial guarantees and index-linked features so that at least for these components the pricing of the contract should make reference to the financial market.

Moreover, the decision makers involved in the pricing process do not only invest in risk-free bonds but use more diversified portfolios. In addition, when the insurance claim can be replicated using financial instruments, the price (premium) for it should be *market consistent*, effectively meaning that any good pricing rule in insurance should be such that it preserves market prices when applied to financial payoffs. Finally, Bernard, Boyle and Vanduffel (2014) (Chapter 1) have shown that, given the distribution of an insurance payoff C_T , it is possible to construct a financial payoff that generates the same distribution as C_T at minimal (market) cost, which further suggests that there should be a link between an insurance pricing principle and pricing in financial markets.

Traditionally, the discounted expectation of the future insurance claim is a lower bound for the insurance premium (calculated through an actuarial valuation principle ignoring the financial market). In other words, premium principles have a “non-negative loading”. It is argued that a premium principle that does not satisfy this requirement can lead to the insurer’s ruin (assuming the insurer faces a series of independent claims so that the law of large numbers holds). Our research shows that in presence of a financial market such *no-undercut principle* does not necessarily hold.

This chapter is related to the literature on pricing of claims in incomplete markets. Specifically, we assume throughout the chapter that there is a financial market that can be used to perfectly hedge financial risk but that the insurance contract also depends on additional sources of uncertainty that could not be hedged using financial instruments. There is already an important literature related to the pricing of insurance contracts that have hedgeable parts and non-hedgeable parts in the presence of a financial market. The idea that the premium can be invested in the financial market first appeared in Kahane and Nye (1975). Brockett et al. (2005) use indifference pricing in the presence of a financial market for weather derivatives and expected utility maximizers. They show how the hedging part is important and how indifference prices could significantly differ to actuarial prices obtained using the discounted expectation. Pelsser (2014) and Wüthrich et al. (2010) extensively discuss market consistency, which is an important issue given the current regulation in the insurance industry about mark-to-market valuation. Recent papers on participating policies, equity-indexed annuities and variable annuities propose to combine the financial and actuarial approach by assessing the risk under the real-world measure and pricing in the risk-neutral world. See for example Chapter 12 and 14 of Dickson et al. (2013) for pricing equity-linked insurance with deterministic or stochastic cash-flow analysis and by using specific risk measures. Recently, Graf et al. (2011) select a risk minimizing asset allocation (under the real world measure) and distribute terminal surplus such that the contract is fair (evaluated using the risk neutral pricing measure). In doing so, they extend the earlier two-step procedure of Barbarin and Devolder (2005). This method consists of first determining a guaranteed interest rate such that certain solvency requirements are satisfied, then obtaining fair contracts (using risk-neutral valuation) by adjusting the participation in terminal surplus accordingly. These two papers ignore mortality risk and actuarial pricing is done through given premium principles. A similar approach is used by Gaillardetz and Lakhmiri (2011).

In this chapter, we propose an alternative pricing approach that takes into account financial and non-financial risk. The prices obtained are shown to be consistent with arbitrage-free pricing valid for financial claims, and incorporate preferences of agents otherwise. Our main contributions can then be summarized as follows:

First, under very general assumptions on preferences, we show that the discounted expected value is a lower bound for bid and ask prices derived from economic indifference pricing principles in the absence of a financial market. *Second*, we propose to price insurance contracts using an indifference principle which incorporates the financial market as a tool to (partially) hedge and as

a way to invest the initial premia received by the insurer. This indifference principle can already be found in Hodges and Neuberger (1989) but is presented here more generally. We also show that it is consistent with arbitrage-free pricing. *Third*, assuming decision makers involved are risk-averse we use results on cost-efficient financial payoffs (see Bernard, Boyle and Vanduffel (2014)) to derive a lower bound on insurance prices (which generalizes and extends Hobson (2005)). *Fourth*, we show that this bound effectively corresponds to the market price of an explicitly known financial payoff. *Fifth*, we prove that the discounted expected value is no longer valid as a classical lower bound for insurance prices in general, and has to be corrected by a covariance term which reflects the interaction between the insurance claim and the financial market. *Finally*, the different results are illustrated by an equity-linked life insurance contract. More generally, such method is applicable to determine explicitly indifference prices of contracts that depend on hedgeable risk and non-hedgeable risk (in an expected utility or non-expected utility framework).

The chapter is further structured as follows. Section 3.1 presents the context, notations and assumptions. In Section 3.2 we recall how bid and ask prices are calculated under a traditional indifference pricing approach. We also derive and discuss the traditional lower bound under very general assumptions. Section 3.3 extends this framework to include the presence of a financial market and derives the new lower bound. In Section 3.4 we illustrate the theoretical results with several examples in the context of a Black-Scholes market. Section 3.5 concludes.

3.1 Setting

Consider an insurance market, described by a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with two market participants: the insured (insurance buyer) and the insurer (insurance seller). We first model the preferences of each agent, then the insurance contract and finally the financial (sub)market. In what follows, the notations X_T and Y_T denote random wealths (hence positive realizations reflect income in these instances) whereas the notation C_T is used to denote a random loss (hence positive outcomes reflect losses). All mentioned expectations are tacitly assumed to exist.

3.1.1 Agents' Preferences

We denote by $U(\cdot)$ and by $V(\cdot)$ the respective objective functions of the buyer and the seller and we make the following assumptions on their preferences.

- Market participants all have a fixed investment horizon $T > 0$ and there is no intermediate consumption. Our model is therefore a one-period model. Specifically, we assume that a premium is paid at the beginning of the period by the buyer to the insurance company, and the indemnity is paid at the end of the period upon occurrence of a loss.
- Market participants all prefer “more to less”, in other words their respective objective functions preserve first order stochastic dominance relationships (denoted by \prec_{fsd}). Hence, if $Y_T \prec_{fsd} X_T$ then $U(X_T) \geq U(Y_T)$ and $V(X_T) \geq V(Y_T)$. Therefore $U(\cdot)$ and $V(\cdot)$ are both non-decreasing.
- Agents have “state-independent preferences” or “law-invariant preferences”: if Y_T has the same distribution as X_T then $U(Y_T) = U(X_T)$ and $V(Y_T) = V(X_T)$.

Such set of preferences is quite general and consistent with a wide range of decision theories, including the expected utility theory (von Neumann and Morgenstern (1953)), Yaari's dual theory of choice (Yaari (1987)), the cumulative prospect theory (Tversky and Kahneman (1992)) and the rank dependent utility theory (Quiggin (1993)). For example, in the particular case of expected utility the preferences for a final wealth X_T would be calculated as $U(X_T) = \mathbb{E}[u(X_T)]$ and $V(X_T) = \mathbb{E}[v(X_T)]$, where u and v are the respective utility functions for the buyer and the seller. Instead of maximizing an objective function, one may also minimize a risk measure, and any law-invariant risk measure that preserves first stochastic dominance can also be used here as well (for example the quantile or any other distorted expectation). This will be illustrated later with an example under Yaari's dual theory of choice under risk.

In the remainder of the chapter, we assume that the individual is *risk-averse* if the following property is verified:

Assumption (Risk aversion) The preferences $U(\cdot)$ of an agent are considered “risk-averse” if for any two random wealths X_T and Y_T ,

$$\begin{cases} \mathbb{E}[X_T] = \mathbb{E}[Y_T] \\ \forall d \in \mathbb{R}, \quad \mathbb{E}[(X_T - d)^+] \leq \mathbb{E}[(Y_T - d)^+] \end{cases} \Rightarrow U(X_T) \geq U(Y_T), \quad (3.1)$$

where $(\cdot)^+ = \max(\cdot, 0)$ and where all expectations (which are assumed to exist) are calculated under the physical measure \mathbb{P} . When X_T and Y_T satisfy the stated conditions then we also say that they are ordered in convex order (denoted by \leq_{cx}). Loosely speaking, $X_T \leq_{cx} Y_T$ states that X_T is less volatile than Y_T for the same expectation, a situation that is preferred by risk averse agents.

The concept of risk-aversion we use in this chapter is standard in actuarial science. In general, it is stronger than so-called *weak risk aversion* (where certainty $\mathbb{E}(X_T)$ is preferred to the randomness X_T itself). However, in the context of expected utility theory both risk aversion concepts coincide and go along with a concave utility function. This is no longer true in Yaari's theory (and most other decision theories) where each notion gives rise to possibly different decision making under risk; we refer to Cohen (1995) for a detailed discussion of risk aversion concepts and their relation with decision theories.

Specifically, we will need the following well-known properties.

Property 1: The notion of risk aversion (3.1) can also be written as¹

$$\begin{cases} \mathbb{E}[X_T] = \mathbb{E}[Y_T] \\ \forall f \text{ convex}, \quad \mathbb{E}[f(X_T)] \leq \mathbb{E}[f(Y_T)] \end{cases} \Rightarrow U(X_T) \geq U(Y_T), \quad (3.2)$$

Property 2: If the objective function $U(\cdot)$ of the individual verifies (3.1) then for all random variables X_T and Z_T

$$U(\mathbb{E}[X_T|Z_T]) \geq U(X_T), \quad (3.3)$$

in particular,

$$U(\mathbb{E}[X_T]) \geq U(X_T). \quad (3.4)$$

Property 3 [“cut criterion”]: A sufficient condition for convex order is the “single crossing property” (or second order stochastic dominance), namely, if X_T and Y_T , with respective distributions F and G , are such that $\mathbb{E}[X_T] = \mathbb{E}[Y_T]$ and

$$\exists c \geq 0, \quad \begin{cases} \forall x \in (0, c), & F(x) \leq G(x) \\ \forall x \in (c, +\infty), & F(x) \geq G(x). \end{cases} \quad (3.5)$$

¹Rigorously it should be for all convex functions f such that the expectations involved exist.

See for example Müller and Stoyan (2002) for a proof of Property 3.

3.1.2 Insurance Contract

We denote by w_b and w_a respectively the initial wealth of the insurance buyer and the insurance seller. We suppose that the buyer faces a non-negative random loss C_T during the period under study, and that the insurance company sells a contract that reimburses $h(C_T)$.

Insurance contracts are designed to avoid as much as possible moral hazard and manipulation of information by the insured. Hence, we impose that $0 \leq h(x) \leq x$. Moreover, we assume that contracts are offered such that $x \mapsto h(x)$ is non-decreasing. Otherwise, the insured has an incentive to partly hide his losses to receive a higher indemnity and this would significantly increase verification costs by the insurer. Finally, to ensure that the insured has no incentive to inflate losses, the retention of the loss $C_T - h(C_T)$ is non-decreasing. Hence, it is assumed that $x \mapsto x - h(x)$ is non-decreasing as well. Thus C_T and $C_T - h(C_T)$ are both non-decreasing in C_T , in other words they form a comonotonic pair². Note that for traditional insurance contracts such as deductible insurance, capped insurance or proportional insurance the properties for $x \mapsto h(x)$ are obviously satisfied.

Assuming that $h(x)$ and $x - h(x)$ are both non-decreasing³ is standard in the optimal insurance design literature. Assume also that this loss C_T is not traded and that in general it cannot be fully hedged in the financial market.

3.1.3 Financial Market

A financial (sub)market is available to all market participants (insured and insurer). We assume it is free of arbitrage, perfectly liquid, frictionless (no transaction costs, no trading constraints) and complete. Let us denote by $(\Omega_1, \mathcal{F}_1, \mathbb{P}_1)$ the corresponding probability space, where \mathcal{F}_1 is the smallest σ -algebra generated by the risky assets traded in the market. Under these assumptions, there exists a unique state-price process $(\xi_t^1)_t$ such that $(\xi_t^1 S_t^1)_t$ is a martingale for all traded assets S^1 . The state-price process is also known as the “pricing kernel”, “deflator” or “stochastic discount factor” (Cochrane (2005)). All \mathcal{F}_1 -measurable payoffs X_T^1 can be hedged and have a price $c(X_T^1)$ given by

$$c(X_T^1) = \mathbb{E}_{\mathbb{P}_1}[\xi_T^1 X_T^1]. \quad (3.6)$$

Here the expectation is taken with respect to the physical probability measure \mathbb{P}_1 on Ω_1 . We further assume that ξ_T^1 is continuously distributed. In particular, one has that

$$c(1) = \mathbb{E}_{\mathbb{P}_1}[\xi_T^1] := e^{-rT}, \quad (3.7)$$

where r is the risk free rate. It is also well-known that $c(X_T^1)$ can be presented as $c(X_T^1) = e^{-rt} \mathbb{E}_{\mathbb{Q}_1}[X_T^1]$, where the expectation is now taken with respect to the so-called risk-neutral measure \mathbb{Q}_1 defined through $\xi_t^1 = e^{-rt} (\frac{d\mathbb{Q}_1}{d\mathbb{P}_1})_t$. We refer to Harrison and Kreps (1979) or more recently to Björk (2004) for an extensive theory on arbitrage-free pricing.

At this stage, the financial claims are not yet well defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. In order to embed the financial market in the global market, we effectively assume that $(\Omega, \mathcal{F}, \mathbb{P})$

²A standard reference in actuarial science on co-monotonicity and anti-monotonicity is Dhaene et al. (2002).

³For example, Carlier and Dana (2011) also restrict themselves to this set and note on page 878 that the map h is then 1-Lipschitz.

is obtained as a suitable product space of $(\Omega_1, \mathcal{F}_1, \mathbb{P}_1)$ with another probability space $(\Omega_2, \mathcal{F}_2, \mathbb{P}_2)$, which represents “pure” insurance (non-financial) risk.⁴ Financial claims X_T can now be readily defined on $\Omega = \Omega_1 \times \Omega_2$ by $X_T(\omega_1, \omega_2) := X_T^1(\omega_1)$. Then $\xi_t(\omega_1, \omega_2) := \xi_t^1(\omega_1)$ entails a state-price process $(\xi_t)_t$ on $(\Omega, \mathcal{F}, \mathbb{P})$, and also induces a risk neutral measure which we denote by \mathbb{Q} . The state-price process at maturity T , ξ_T , is thus obtained as the product of the pricing kernel for the financial claims, times a random variable identically equal to 1 for the non-financial risks.

Furthermore, the (pure) financial payoffs X_T are hedgeable and have a unique price that writes as $c(X_T) = \mathbb{E}_{\mathbb{P}}[\xi_T X_T]$. In general, a claim C_T contains a (pure) financial part and a (pure) insurance part and is no longer hedgeable. It is then not clear how to price such C_T . Indeed, while $(\xi_t^1)_t$ is unique on $(\Omega_1, \mathcal{F}_1, \mathbb{P}_1)$, the corresponding $(\xi_t)_t$ is just one of the many state-price processes that can be defined on the incomplete market $(\Omega, \mathcal{F}, \mathbb{P})$. However, it turns out that the process $(\xi_t)_t$ as defined above gives rise to the construction of a lower bound on all other pricing rules (therefore all other state-price processes) of the incomplete market $(\Omega, \mathcal{F}, \mathbb{P})$. In the remainder we will always explicitly write “ $\mathbb{E}_{\mathbb{Q}}$ ” when working under the \mathbb{Q} -measure.

3.2 Traditional Indifference Pricing Approach

3.2.1 Indifference Pricing

We now define bid and ask premiums (denoted by p_b and p_a respectively), following the traditional actuarial approach for indifference pricing (also called “certainty equivalents”).

For an insurance indemnity $h(C_T)$, bid and ask prices are respectively obtained as solutions to

$$U((w_b - p_b)e^{rT} - C_T + h(C_T)) = U(w_b e^{rT} - C_T), \quad (3.8)$$

and

$$V(w_a e^{rT}) = V((w_a + p_a)e^{rT} - h(C_T)). \quad (3.9)$$

Here, the left hand side of (3.8) corresponds to the utility obtained by buying insurance at a premium p_b from the insurer and the right hand side corresponds to the situation when the insured did not buy any insurance and suffers from the loss C_T . The value at T of any initial wealth w is assumed to be equal to $w e^{rT}$ effectively meaning that in this setting only the risk-free investment is available. The ask price p_a is derived similarly using the insurer’s preferences $V(\cdot)$ in (3.9). See also Bühlmann (1980). Note that apart from the availability of a risk-free asset the presence of the financial market does not intervene in (3.8) and (3.9).

3.2.2 Classical lower bound

We are already able to give a well-known lower bound valid for both bid and ask prices.

⁴We refer to Møller (2003a; 2003b) for detailed information on this construction; here we merely provide the main ideas. As noted by Møller (2003a; 2003b), “This leads to a new market consisting of the same financial assets, but where the class of claims is extended in that claims can depend on the additional insurance risk as well as on the financial assets” and also, “[t]his construction has the advantage that it allows to merge directly classical insurance risk models with models for financial markets.”

Proposition 3.2.1 (Classical lower bound). *Assume that both the insured and the insurer are risk-averse in the sense of assumption (3.1) and that bid and ask prices are defined through (3.8) and (3.9). Denote by p_\bullet either the bid price p_b or the ask price p_a , then*

$$p_\bullet \geq e^{-rT} \mathbb{E}[h(C_T)]. \quad (3.10)$$

The proof of Proposition 3.2.1 is given in the appendix. Such result is standard and typically appears in textbooks in the context of full insurance (i.e. $h(C_T) = C_T$) and assuming agents have concave increasing preferences. Here we present and prove the result more generally.

This classical bound (3.10) conforms well with common intuition and is consistent with the desire to avoid ruin. Note that such bound does not depend on the respective initial wealths and on the specifications of preferences.

Remark 3.2.1 (Full insurance case). In the case of full insurance, $h(C_T) = C_T$, the lower bound (3.10) already holds when decision makers are *weakly risk-averse* so that assumption (3.1) is not strictly required. More precisely, (3.10) holds whenever the preferences satisfy $U(X_T) \leq U(\mathbb{E}[X_T])$ and $V(X_T) \leq V(\mathbb{E}[X_T])$.

3.2.3 Comparison between bid and ask prices

In the context of classical indifference pricing, insurance agreements are only possible when the bid price p_b is larger than the ask price p_a , implying that it is of interest to analyze to which extent such ordering exists or not. Equations (3.8) and (3.9) show that p_b and p_a generally depend on the preferences $U(\cdot)$ and $V(\cdot)$ as well as w_b and w_a in a non-straightforward way, indicating that strong conclusions regarding the ordering that may exist between both are not obviously in reach. For example, let $U(\cdot) = V(\cdot) = \mathbb{E}[u(\cdot)]$ with utility function $u(x) = 1 - \frac{1}{x}$, and let both agents also have the same initial wealth $w_a = w_b = w$. Take for C_T a uniformly distributed random variable on $(0, 2)$. Then Figure 3.2.1 below represents bid and ask prices with respect to initial wealth. It is clear that no general assertions can be made, In particular, it is possible to construct cases where $p_b < p_a$.

However, under the additional assumption that the insurer is risk neutral, which is not so unreasonable, it follows that $p_b \geq p_a$ always holds so that insurance agreements are possible in these instances. We make this clear in the following proposition.

Proposition 3.2.2 (Risk neutral insurer). *If the insurance seller is risk neutral (i.e. $V(\cdot) = \mathbb{E}[\cdot]$) then it follows that the bid price for partial insurance verifies*

$$p_b \geq p_a = e^{-rT} \mathbb{E}[h(C_T)].$$

The proof of Proposition 3.2.2 follows from Proposition 3.2.1 and from the fact that equation (3.9) immediately implies that under risk neutrality the ask price p_a coincides with the lower bound.

Note that the proposition holds for partial and full insurance schemes. Also, when both insurer and the insured are risk neutral, it follows that $p_b = p_a = e^{-rT} \mathbb{E}[h(C_T)]$ holds.

We now provide some more specialized results when agents have preferences according to exponential expected utility or when their set of preferences can be described using Yaari (1987)'s theory of decision making under risk.

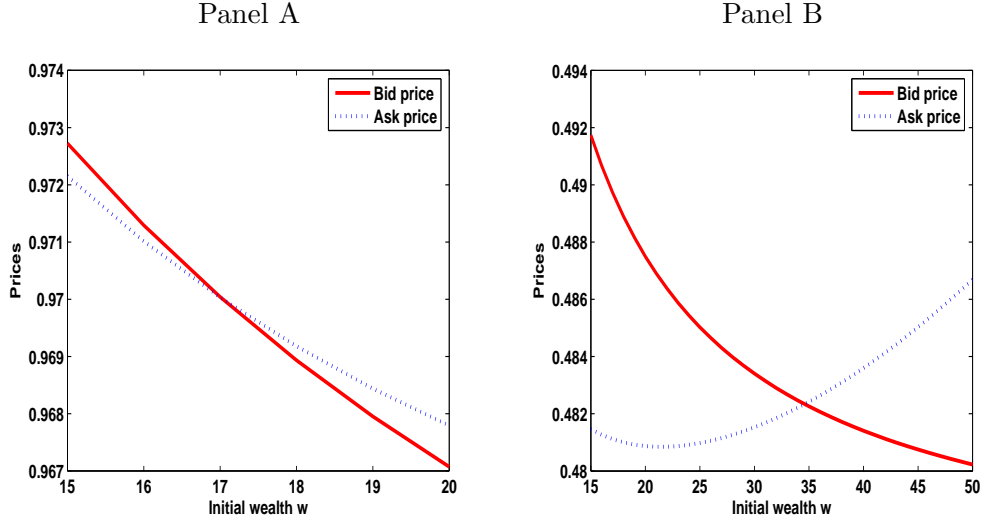


Figure 3.2.1: Bid and ask prices as a function of initial wealth $w_a = w_b = w$, $r = 5\%$ and $T = 1$. Panel A corresponds to full insurance. Panel B corresponds to partial insurance with $h(x) = .5x$ (proportional insurance).

a. The case of Exponential Expected Utility

Let $U(X) = V(X) = \mathbb{E}[u(X)]$, where $u(x) = \frac{1}{\gamma}(1 - e^{-\gamma x})$ with $\gamma > 0$. It is straightforward to derive from equations (3.8) and (3.9) that one has

$$p_a e^{rT} = \frac{1}{\gamma} \ln(\mathbb{E}[\exp(\gamma h(C_T))])$$

and

$$p_b e^{rT} = \frac{1}{\gamma} \ln \left(\frac{\mathbb{E}[\exp(\gamma C_T)]}{\mathbb{E}[\exp(\gamma(C_T - h(C_T)))]} \right). \quad (3.11)$$

Then, since $C_T = h(C_T) + (C_T - h(C_T))$ and observing that

$$\text{Cov}[\exp(\gamma h(C_T)), \exp(\gamma(C_T - h(C_T)))] \geq 0,$$

because of the comonotonicity of $h(C_T)$ and $C_T - h(C_T)$, it is clear that $p_b \geq p_a$. Finally, when $h(C_T) = C_T$ one has that $p_b = p_a$.

b. The case of Yaari (1987)'s theory

Under Yaari's dual theory of choice, decision makers evaluate the "utility" of a random variable X by calculating its distorted expectation $\mathbb{H}_f[X]$ which is formally defined as the Choquet-integral

$$\mathbb{H}_f[X] = - \int_{-\infty}^0 [1 - f(\bar{F}_X(x))] dx + \int_0^{\infty} f(\bar{F}_X(x)) dx, \quad (3.12)$$

where $\bar{F}_X(x)$ is denoting the de-cumulative distribution function and where the (distortion) function $f : [0, 1] \rightarrow [0, 1]$ is non-decreasing with $f(0) = 0$ and $f(1) = 1$; see also Wang and Young (1997) and Dhaene et al. (2006). A subclass of distortion functions that is often considered in the literature is the class of concave distortion functions. It is well-known that when a distortion function f

is concave the corresponding premium principle is coherent in the sense of Artzner et al. (1999) meaning that it satisfies the properties of sub-additivity, monotonicity, positive homogeneity and translation invariance.

Furthermore, when f is differentiable it follows that expression (3.12) can also be rewritten as

$$\mathbb{H}_f[X] = \mathbb{E} \left[F_X^{-1}[1 - U] f'(U) \right]. \quad (3.13)$$

Hence, in the context of Yaari's theory, it follows that the indifference principles (3.8) and (3.9) for deriving the bid and ask prices p_b and p_a give rise to the following expressions for full insurance

$$p_a = p_b = -e^{-rT} \mathbb{H}_f[-C_T] = e^{-rT} \mathbb{H}_{\bar{f}}[C_T], \quad (3.14)$$

where \bar{f} is the so-called 'dual distortion function' of f i.e.

$$\bar{f}(x) = 1 - f(1 - x), \quad x \in [0, 1], \quad (3.15)$$

and note that \bar{f} is indeed a distortion function.

In the case of partial insurance, the ask price writes as

$$p_a = e^{-rT} \mathbb{H}_{\bar{f}}(h(C_T)),$$

whereas the bid price is given by

$$p_b = e^{-rT} \left[\mathbb{H}_{\bar{f}}(C_T) - \mathbb{H}_{\bar{f}}(C_T - h(C_T)) \right].$$

Since Yaari's objective function exhibits additivity for sums of comonotonic risks, and since both $h(C_T)$ and $(C_T - h(C_T))$ in the above sum are comonotonic (thanks to assumptions on the insurance design in Section 3.1.2), we have $p_b = p_a$.

The indifference principles (3.8) and (3.9) have been extensively used in actuarial science. However, they assume that the insurer and the insured do not make use of the financial market and invest only in bonds yielding a risk-free rate r . Section 3.3 generalizes the indifference pricing of insurance claims by considering the presence of a financial market explicitly.

3.3 Market-based Indifference Pricing Approach

3.3.1 General indifference principle

Hodges and Neuberger (1989) generalize the static utility indifference pricing to include the presence of financial markets as a dynamic way to invest and a tool to hedge, and we follow their approach to (re)define the bid price p_b and the ask price p_a for insurance claims.

Hence, for a given initial wealth $w > 0$ we define $A(w)$ as the set of random wealths X_T that can be generated at maturity $T > 0$ by trading in the financial market. Assume now that during the considered horizon the investor is exposed to a risk with random payoff C_T . As before one needs to consider the viewpoint of the buyer (insured) as well as of the seller (insurer). Then, from the viewpoint of the insured the (bid) price p_b is such that

$$\sup_{X_T \in A(w_b - p_b)} U[X_T - C_T + h(C_T)] = \sup_{X_T \in A(w_b)} U[X_T - C_T], \quad (3.16)$$

whereas from the viewpoint of the insurer, the ask price p_a follows from

$$\sup_{X_T \in A(w_a + p_a)} V[X_T - h(C_T)] = \sup_{X_T \in A(w_a)} V[X_T]. \quad (3.17)$$

We now formulate some remarks.

Remark 3.3.1 (Well-defined prices). We assume that preferences $U(\cdot)$ and $V(\cdot)$ are such that all the above suprema exist, are finite and reached. For example the risk neutrality case is excluded.

Remark 3.3.2 (Extension of the classical actuarial framework). When the financial market is reduced to the existence of a bank account, indifference prices derived from (3.16) and (3.17) coincide with the ones derived in the classical setting (3.8) and (3.9). Formulas (3.16) and (3.17) are therefore a natural extension of the classical setting of utility indifference pricing presented in the previous section.

Remark 3.3.3 (Consistency). There is not a general consistency between prices derived in the presence of a financial market and prices derived in the absence of such market. Even when the claim is independent of the financial market, the indifference prices derived from (3.16) and (3.17) are in general not equal to the ones derived in the classical setting (3.8) and (3.9) where the market is ignored. Intuitively, this is because in the former case, even when the claim is independent of the market, the possibility to invest the premium in the financial market is reflected in the indifference price. There is however equality in some very particular cases.

The next proposition shows that the general indifference principle is consistent with arbitrage-free pricing. For a proof in the expected utility framework, see Hendersen and Hobson (2004). A general proof is given in the appendix.

Proposition 3.3.1 (Market consistency). *Assume $h(C_T)$ is hedgeable, then*

$$p_b = p_a = \mathbb{E}[\xi_T h(C_T)]. \quad (3.18)$$

Market consistency is an important issue given the current regulation in the insurance industry about mark-to-market valuation. See for example Pelsser (2014) and Wüthrich et al. (2010). While there are some techniques⁵ available in the literature, it remains difficult to compute explicitly bid and ask prices from (3.16) and (3.17). This further underpins the need for determining bounds that can be computed easily. In the remainder of the chapter, we discuss bounds on bid and ask prices in the presence of a financial market.

3.3.2 Lower bounds on insurance claims

The main result of this section can be found in Theorem 3.3.3 where we give a lower bound for bid and ask prices in the presence of a financial market. To prove and understand this theorem, we first need the following proposition.

Proposition 3.3.2 (Optimal wealth). *Under the assumptions in Section 3.1.1 for an objective function $U(\cdot)$, the assumptions in Section 3.1.3 for the financial market and given an initial wealth $w \in \mathbb{R}$ it holds that*

$$\sup_{X_T \in A(w)} U(X_T) = \sup_{X_T \in A_\xi(w)} U(X_T), \quad (3.19)$$

⁵Such as the martingale approach (Cox and Huang (1989)) or dynamic programming (Carmona (2008)).

where $A(w)$ is the set of random wealths X_T that can be generated at maturity $T > 0$ with an initial wealth w , and where $A_\xi(w)$ is the subset of random wealths that are almost surely anti-comonotonic with ξ_T (in other words which are almost surely a non-increasing function of ξ_T).

Proposition 3.3.2 shows that any optimal investment strategy in the financial market has to be non-increasing in the state-price process at the horizon of the strategy T . This result first appeared in the pioneering work of Dybvig (1988a; 1988b) and was more recently extended by Bernard, Boyle and Vanduffel (2014). Various examples illustrating these findings can be found in Bernard, Boyle, and Vanduffel (2014) in a one-dimensional Black-Scholes market as well as in Bernard, Maj and Vanduffel (2011) in the n -dimensional Black-Scholes setting; see also Vanduffel et al. (2009) for the one-dimensional Lévy market.

Proposition 3.3.3 (Lower bound for insurance prices). *Let us assume that the objective functions satisfy the properties of Section 3.1.1 and are also such that the investor is risk-averse according to (3.1). Denote by p_\bullet either the bid price p_b or the ask price p_a ,*

- In the full insurance case ($h(C_T) = C_T$),

$$\mathbb{E}[\xi_T C_T] \leq p_\bullet. \quad (3.20)$$

- In the partial insurance case,

$$\mathbb{E}[\xi_T h(C_T)] \leq p_a.$$

Furthermore, we observe that the lower bound $\mathbb{E}[\xi_T C_T]$ is also the market price of the financial payoff $\mathbb{E}[C_T | \xi_T]$.

As discussed in Section 3.1.3, the state-price process at time T is not uniquely defined in an incomplete market. However, what we denote here by ξ_T is a specific construction for the state-price process which corresponds of the product of the pricing kernel for financial claims, times a random variable identically equal to 1 for the “non-financial risks”.

Inequality (3.20) essentially states that both the insured and insurer are not prepared to accept a price for the insurance claim C_T lower than the market price for $\mathbb{E}[C_T | \xi_T]$. Assuming the classical expected utility paradigm, this bound was already derived (using involved techniques adapted to his setting and presented in another context) in Hobson (2005). Here we show in a straightforward way that the bound is valid under a more general law-invariant preference framework and also that the lower bound holds for the ask price in case of partial insurance agreements. In particular, the bounds hold when assuming Yaari’s dual theory of choice under risk (Yaari (1987)). The following remark shows that the lower bound depends on the interaction between the claim C_T and “the market”.

Remark 3.3.4 (Other expression for lower bound). We have that

$$\mathbb{E}[\xi_T C_T] = e^{-rT} \mathbb{E}[C_T] + \text{Cov}[C_T, \xi_T]. \quad (3.21)$$

This follows from $\mathbb{E}[C_T \xi_T] = \text{Cov}[C_T, \xi_T] + \mathbb{E}[C_T] \mathbb{E}[\xi_T] = \text{Cov}[C_T, \xi_T] + \mathbb{E}[C_T] e^{-rT}$.

Remark 3.3.4 states that when the claim C_T and the state-price ξ_T are negatively correlated we find that $e^{-rT} \mathbb{E}[C_T]$ is no longer a lower bound for p_b and p_a , which contrasts with traditional (and intuitively appealing) wisdom stated in many actuarial textbooks on insurance pricing. The following remark clarifies further why this common belief is no longer valid. Mainly for ease of exposition, we restrict to the full insurance case.

Remark 3.3.5 (Relation with classical lower bound). Let us compare the lower bound (3.20) of Theorem 3.3.3 with the classical lower bound $e^{-rT}\mathbb{E}[C_T]$ that we obtained in a market without risky assets. Several cases can be identified.

- If C_T is independent of the market, the two lower bounds are equal

$$p_{\bullet} \geq e^{-rT}\mathbb{E}[C_T]. \quad (3.22)$$

Loosely speaking, the independence implies that “the financial market cannot help at all to hedge the insurance claim”. The lower bound is therefore also independent of the financial market and it appears therefore intuitive that it coincides with the classical bound.

- If C_T is positively correlated with the state-price ξ_T , the classical lower bound $e^{-rT}\mathbb{E}[C_T]$ is now strictly improved.

$$p_{\bullet} \geq e^{-rT}\mathbb{E}[C_T] + \text{Cov}[C_T, \xi_T] > e^{-rT}\mathbb{E}[C_T]. \quad (3.23)$$

- However if C_T is negatively correlated with the state-price ξ_T , the new lower bound is smaller

$$p_{\bullet} \geq e^{-rT}\mathbb{E}[C_T] + \text{Cov}[C_T, \xi_T]. \quad (3.24)$$

Following Remark 3.3.5, the lower bound for indifference prices of equity-linked insurance benefits will typically be *lower* than $e^{-rT}\mathbb{E}[C_T]$. Indeed equity-linked insurance payoffs are commonly increasing with some risky asset price S_T present in the market. We then observe that the covariance under the physical measure of the risky asset S_T and the state-price process ξ_T is given by

$$\begin{aligned} \text{Cov}(S_T, \xi_T) &= \mathbb{E}[S_T \xi_T] - \mathbb{E}[S_T]\mathbb{E}[\xi_T] \\ &= e^{-rT}(\mathbb{E}_Q[S_T] - \mathbb{E}[S_T]). \end{aligned} \quad (3.25)$$

Hence, $\text{Cov}(S_T, \xi_T) < 0$ usually holds (because the expected return of the risky asset in the real-world usually exceeds the risk-free rate). Therefore, for an equity-linked insurance contract it is possible that both the insured and the insurer agree on a premium lower than the discounted expected value.

Remark 3.3.5 also shows that the lower bound on the insurance price may be larger in the case when the insurance claim contains for instance a put option without upside participation in the financial market. The covariance between the claim and the stock price will then be negative and thus the covariance between C_T and ξ_T will be positive. In this case the actuarial bound $e^{-rT}\mathbb{E}[C_T]$ is not valid anymore and the price should be strictly bigger. Intuitively, the reason is that the states where the market is down are generally the most expensive states and this should be reflected in the price of the insurance claim.

Note also that when the financial market is reduced to a bond earning the risk-free rate, $\xi_T = e^{-rT}$ almost surely, and we find that (3.21) is the classical actuarial lower bound $e^{-rT}\mathbb{E}[C_T]$ again. Finally, we stress that the lower bound, calculated as the price of a financial payoff, is only valid at time $t = 0$. In fact, holding the financial payoff $\mathbb{E}[C_T|\xi_T]$ does not imply at all that its value at time t ($0 < t < T$) is a lower bound for the corresponding t -price of C_T .

3.3.3 Upper bound

In general, it is not possible to obtain an upper bound on indifference prices. The next theorem shows that such upper bound exists in the case of financial claims.

Proposition 3.3.4. *If C_T is hedgeable using the financial market (pure financial risk), then bid and ask prices for the partial insurance case $h(C_T)$ verify*

$$\ell = \mathbb{E} \left[\xi_T \cdot F_{C_T}^{-1}(1 - F_{\xi_T}(\xi_T)) \right] \leq p_{\bullet} \leq u := \mathbb{E} \left[\xi_T \cdot F_{C_T}^{-1}(F_{\xi_T}(\xi_T)) \right]$$

However, it is in general not possible to find such upper bound (independently of the preferences) and indifference prices may be unbounded when they involve non-financial risk. This feature already becomes obvious in the absence of financial market such as it appears in (3.11) when $h(C_T) = C_T$ and $\gamma \rightarrow +\infty$.

3.4 Example in the Black and Scholes model

Consider a one-dimensional Black Scholes model ($m = 1$). The risky asset price S_t evolves according to

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t, \quad (3.26)$$

where $\{W_t, t \geq 0\}$ is a standard \mathbb{P} -Brownian motion and $\mu > r$. The state price process $\{\xi_t, t \geq 0\}$ exists, is unique and is given by (1.8) in Chapter 1. For the ease of exposition, we recall its expression:

$$\xi_t = a \left(\frac{S_t}{S_0} \right)^{-\frac{\theta}{\sigma}}, \quad (3.27)$$

where $a = e^{\frac{\theta}{\sigma}(\mu - \frac{\sigma^2}{2})t - (r + \frac{\theta^2}{2})t}$ and $\theta = \frac{\mu - r}{\sigma}$. Also note that ξ_t is decreasing in S_t and for all $c \in \mathbb{R}$

$$\mathbb{P}(S_t > c) > \mathbb{Q}(S_t > c). \quad (3.28)$$

Finally, $\ln(\xi_t)$ is normally distributed with mean $M = -\frac{1}{2}\theta^2 t - rt$ and variance $\theta^2 t$.

Consider now the simplest possible insurance claim that pays at time $T = 1$ a payoff C_1 distributed as a Bernoulli distribution with parameter $p = 0.001$. We have that $\mathbb{P}(C_1 = 1) = p$ and $\mathbb{P}(C_1 = 0) = 1 - p$. We illustrate the behaviour of the lower bound of Theorem 3.3.3 hereby using the three cases presented in Remark 3.3.5.

First, we assume that the insurance claim may be linked to the death of a specific individual solely and is therefore non-tradable and independent of the financial market. In this case

$$\mathbb{E}[C_1 | \xi_1] = \mathbb{E}[C_1]. \quad (3.29)$$

Hence, applying Theorem 3.3.3, bid and ask prices, p_{\bullet} both satisfy

$$p_{\bullet} \geq \mathbb{E}[\xi_1 \mathbb{E}[C_1 | \xi_1]] = e^{-r} \mathbb{E}[C_1]. \quad (3.30)$$

as we could expect.

Second, we assume that C_1 pays 1 if a designated person dies and at the same time the level of the risky asset exceeds a given value H or, equivalently, the state-price process ξ_1 is lower than

some L (i.e. $\{\xi_1 < L\} = \{S_1 > H\}$). Then the insurance claim depends on both a non-tradable risk and on the financial market. The conditional expectation can still be calculated easily invoking the independence between death and the risky asset:

$$\begin{aligned} \mathbb{E}[C_1|\xi_1] &= \mathbb{E}[\mathbf{1}_{\text{death}}\mathbf{1}_{\xi_1 < L}|\xi_1] \\ &= \mathbb{P}(\text{death})\mathbf{1}_{S_1 > H} \end{aligned} \quad (3.31)$$

Hence, the market price of the claim $\mathbb{E}[C_1|\xi_1]$ is given by $e^{-r}\mathbb{P}(\text{death})\mathbb{Q}(S_1 > H)$ and thus from Theorem 3.3.3

$$p_\bullet \geq e^{-r} \mathbb{P}(\text{death}) \mathbb{Q}(S_1 > H). \quad (3.32)$$

Using (3.28) we note that

$$\mathbb{E}[C_1] = \mathbb{P}(\text{death}) \mathbb{P}(S_1 > H) > \mathbb{P}(\text{death}) \mathbb{Q}(S_1 > H) \quad (3.33)$$

so that the classical bound might be violated, i.e. $p_\bullet < e^{-r}\mathbb{E}[C_1]$ may hold.

Third, we assume that C_1 pays 1 if the designated person dies and the risky asset in the market is lower than H ($\{S_1 < H\} = \{\xi_1 > L\}$). One has $\text{Cov}(C_1, \xi_1) > 0$ and bid and ask prices satisfy

$$p_\bullet \geq \mathbb{E}[\xi_1 \mathbb{E}[C_1|\xi_1]] = e^{-r} \mathbb{P}(\text{death}) \mathbb{Q}(S_1 < H) > e^{-r} \mathbb{E}[C_1]. \quad (3.34)$$

The published version of this chapter contains an additional example of more realistic equity-linked life insurance contract. But the main ideas are already contained in the previous stylized example.

3.5 Conclusions

We have analyzed lower and upper bounds for prices of insurance claims obtained from indifference principles. The analysis is first conducted in the absence of investment opportunities (classical framework presented in actuarial textbook) and then in the presence of financial markets. Our approach extends the pricing of non-tradable claims in incomplete market (as presented in Hendersen and Hobson (2004) and Hobson (2005)) to general preferences (and not only the expected utility framework). While our results are valid in a quite general market setting, the presentation is adapted to the specific setting of insurance contracts and deals with partial insurance agreements in particular.

Taking into account the financial market when pricing insurance claims is important. In general, its presence influences indifference prices even in the case when the claim is independent of the financial market. Also, when the claim is hedgeable in the financial market its indifference price will be unique (independent of the preferences) and equal to the no-arbitrage price (market-consistency).

In the case of full insurance, explicit bounds are obtained for bid and ask prices and they correspond to the market prices of some financial payoffs that we determine explicitly. We also have bounds for the ask price in case of partial insurance agreements. The bounds are also robust in the sense that they are likely to be agreed on by all market participants. For financial claims, there exists an upper bound on indifference prices whereas it is not possible to find such upper bound when the risk is not hedgeable.

Part III

State-Dependent Preferences

Chapter 4

Mean-Variance Optimal Portfolios with Applications to Fraud Detection

So far, we have developed direct applications of cost-efficiency in Chapter 2 to expected utility theory and in insurance pricing in Chapter 3. In this chapter, we extend the concept to deal with a “state-dependent” problem. Specifically, the investment decision now involves a benchmark traded in the market that can be taken as a reference to make decisions. The chapter is organized as follows. We first study mean-variance efficient portfolios when there are no trading constraints and show that optimal strategies perform poorly in bear markets. We then assume that investors use a stochastic benchmark (linked to the market) as a reference portfolio. We derive mean-variance efficient portfolios when investors aim to achieve a given correlation (or a given dependence structure) with this benchmark. We also provide upper bounds on Sharpe ratios and show how these bounds can be useful for fraud detection. For example, it is shown that under some conditions it is not possible for investment funds to display a negative correlation with the financial market and to have a positive Sharpe ratio. All the results are illustrated in a Black-Scholes market. An extended version of this chapter is published in the *European Journal of Operational Research* (Bernard and Vanduffel (2014b)).

Markowitz (1952) and Roy (1952) were first in proposing a quantitative approach to determine the optimal trade-off between mean (return) and variance (risk). Their framework is nowadays known as mean-variance analysis and has become very influential as it combines algebraic simplicity with practical applicability. Markowitz pursued the study of optimal investment portfolios and his seminal works initiated a tremendous amount of research heading in several directions, ranging from the study of other notions for measuring risk, multi-period models (Mossin (1968), Cui et al. (2014)), non-negative final wealth (Korn and Trautmann (1995)) and imperfect markets (Xia and Yan (2006), Lim (2004; 2005)) to the inclusion of ambiguity on the returns (Goldfarb and Iyengar (2003)) or an uncertain horizon (Martellini and Urošević (2006)). See also Zhang et al. (2009), Leung et al. (2012) and Prigent (2007) for an excellent overview. More recently, several authors have been working on quadratic hedging or mean-variance hedging, which corresponds to the problem of approximating, with minimal mean squared error, a given payoff by the final value of a self-financing trading strategy in a financial market; see e.g., Lim (2005), Pham (2000) and Schweizer (1992; 2010), to cite only a few.

In an important contribution, Basak and Chabakauri (2010) have fully characterized time-consistent dynamic mean-variance optimal strategies. At any date prior to maturity, a time-

consistent optimal strategy is the best possible mean-variance efficient allocation of wealth, assuming that an optimal mean variance strategy is also selected at each later instant in time. Mean-variance optimal strategies that are derived in a static setting¹ violate time-consistency in the sense that it may become optimal for a mean-variance investor to deviate away from this optimal mean-variance strategy during the investment horizon. However, these optimal strategies (derived in the static setting) can still be justified by assuming that the investor is pre-committed at time 0 and thus executes the dynamic investment strategy that has been decided at time $t = 0$. While time-consistency is a natural requirement, the assumption of pre-commitment is compatible with an investment practice in which a (retail) investor purchases a financial contract (from a financial institution) and does not trade (herself) afterwards.² It also fits with the behavior of an investment manager who revisits (optimizes) his portfolio periodically and sticks to his strategy between two dates. In practice, managers and other investors may also have additional constraints when optimizing their portfolio. One motivation for having constraints is that optimal (unconstrained) strategies are typically long with the market index and perform poorly in poor economic situations (Bernard, Boyle, and Vanduffel (2014)). In this chapter, we show that the static setting is well suited to deal with a certain type of constraints that we motivate economically. See also Wang and Forsyth (2011) for a numerical approach of mean variance efficiency in a time-consistent framework and numerical comparisons of pre-committed strategies and time-consistent strategies.

Traditional mean-variance optimization consists in finding the best pre-committed allocation of assets assuming a buy-and-hold strategy or a constant-mix strategy (which requires a dynamic rebalancing to ensure a constant percentage invested in each asset). The question raises then how pre-committed mean-variance efficient portfolios can be derived when all strategies are allowed and available. Of course, allowing for more trading strategies and thus more degrees of freedom will further enhance optimality. The *first contribution* in this chapter is to derive optimal mean-variance strategies in this setting. We show that the optimal portfolio consists of a short position in the stochastic discount factor used for pricing derivatives and a long position in cash. We are also able to compute the maximum possible Sharpe ratio (Sharpe (1967)) of an optimal (mean-variance efficient) strategy. Bounds on the Sharpe ratio can be useful to regulators or other market participants for fraud detection, i.e., to assess whether the reported performance of a strategy is feasible or not. Recall for example that the Sharpe ratio of Madoff's strategy lied far above the maximum Sharpe ratio for plausible strategies (Bernard and Boyle (2009)).

In the second part of the chapter we extend our study to the case when there is additional information on the strategy, for example on the way it interacts with the financial market or any other benchmark asset as a source of background risk. Our *second contribution* is then to derive tighter bounds on the Sharpe ratio. This is useful for improved fraud detection or abnormal performance reporting. For example, it is shown that under some conditions it is not possible for investment funds to display negative correlation with the financial market and to have a positive Sharpe ratio.

Considering the interaction with a benchmark asset is also a natural way to make mean-variance efficient strategies more resilient against declining markets. Indeed, the mean-variance efficient portfolios derived in the first part of the chapter provide no protection against bear markets. In practice, many investors reward strategies that offer protection or, more generally, that exhibit some desired dependence with any other source of background risk (which we refer to as a benchmark).

¹By a "static setting", we mean that the strategy is derived at the initial time $t = 0$ as the mean-variance efficient optimum with respect to the terminal wealth W_T without consideration for its properties at intermediate dates.

²This is also consistent with the work of Goldstein, Johnson and Sharpe (2008) who propose a tool that allows consumers to specify their desired probability distribution of terminal wealth at maturity.

Our *third contribution* is to derive mean-variance optimal allocation schemes for investors who exhibit state-dependent preferences in the sense that they care about the first two moments of the strategy's distribution *and* additionally aim at obtaining a desired correlation or dependence with a benchmark asset.

The rest of the chapter is organized as follows. The optimal portfolio problem and the assumptions on the financial market are presented in Section 4.1. Section 4.2 provides explicit expressions for mean-variance efficient portfolios when there are no trading constraints as well as a first application to fraud detection. Sections 4.3 and 4.4 extend these preliminary results to the case when there are constraints on the correlation (respectively on the dependence) with a benchmark and illustrate how these results are particularly useful to improve fraud detection tools. Final remarks are presented in Section 4.5.

4.1 Market Setting

In this section, we provide our main assumptions and definitions. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space that describes a financial market. Assume that all market participants agree to use a (non-negative) stochastic discount process $(\xi_t)_t$ for pricing, i.e. the price at time 0 for a strategy with terminal payoff X_T (paid at time $T > 0$) writes as

$$c(X_T) = \mathbb{E}[\xi_T X_T]. \quad (4.1)$$

Note that the price of the unit cash-flow at time T is given by $c(1) = \mathbb{E}[\xi_T]$ and we define the risk-free rate r such that $e^{-rT} = \mathbb{E}[\xi_T]$. All payoffs X_T are assumed to be square integrable ensuring that $c(X_T) < +\infty$. In particular $\text{var}(\xi_T) < +\infty$. We remark that this practice is usually motivated by assuming a frictionless and arbitrage-free financial market where the usual definition of absence of arbitrage is employed. In particular, we do not take into account transaction costs (Pelsser and Vorst (1996)). When the market is complete (all payoffs can be replicated) the stochastic discount factor ξ_T is uniquely given, but in general an infinite number of choices is possible. However, using a milder notion of arbitrage, Platen and Heath (2006) argue that under some conditions, the stochastic discount factor ξ_T corresponds to the inverse of the so-called Growth Optimal Portfolio (GOP)³ and also that the latter can be proxied by a market index. This motivates why in the remainder of the chapter we refer to $1/\xi_t$ as “*the market index*” and we denote it by S_t^* . The pricing formula (4.1) can then be interpreted as the arithmetic average of the possible outcomes all expressed in units of the market index. Note how low values for the market index S_t^* correspond to high values for the discount factor ξ_t . This is consistent with economic theory in the sense that the states of a downturn are usually the most expensive states to insure and therefore correspond to the states ω where the highest values for the discount factor $\xi_T(\omega)$ are observed.

In the remainder of the chapter, we consider an investor with a fixed horizon $T > 0$ without intermediate consumption. We denote by $W_0 > 0$ her initial wealth. For convenience, we assume that all ξ_t ($t > 0$) are continuously distributed.

³The Growth Optimal Portfolio is a diversified strategy which ultimately outperforms all other strategies with probability one. In the literature it also appears as the *Kelly portfolio*.

4.2 Unconstrained Mean-Variance Optimal Portfolios

4.2.1 Mean-Variance Efficiency

Roy (1952) and Markowitz (1952) propose a quantitative approach to find mean-variance efficient allocation among risky assets assuming a buy-and-hold strategy. Their technique can also be applied in the context of constant-mix strategies. In this section we study mean-variance efficient portfolios when there are no restrictions on the possible strategies. Finding optimal policies turns out to be surprisingly simple. Indeed, let us first observe that an optimal mean-variance efficient final payoff X_T^* must necessarily be the cheapest possibility to generate a maximum mean for the given variance level.⁴ Otherwise it is easy to contradict the optimality of this payoff. Indeed, if the optimum is not the cheapest strategy then there is thus another strategy that is cheaper and also has maximum mean. The cost benefit can be invested in the risk free account and one obtains a strategy that has a higher mean for the same variance, which contradicts the mean-variance efficiency of the strategy.

Constructing a cheapest strategy amounts to minimizing the price (4.1). Observe that since $\text{std}(X_T^*)$ and $\text{std}(\xi_T)$ are both fixed and finite, minimizing the price (4.1) is equivalent to minimizing the correlation⁵ between X_T^* and the discount factor ξ_T . It is then a standard result in statistics that this occurs if and only if the optimal payoff X_T^* is linear in ξ_T (with a negative slope): the correlation is then minimized and equal to -1. The following proposition is now rather intuitive. A formal proof follows.

Proposition 4.2.1 (Mean-variance efficient portfolios). *Assume that the investor aims for a strategy that maximizes the expected return for a given variance s^2 for $s \geq 0$. The solution of the following mean-variance optimization problem*

$$\begin{cases} \max & \mathbb{E}[X_T] \\ \text{var}(X_T) = s^2 \\ c(X_T) = W_0 \end{cases}, \quad (4.2)$$

is denoted by X_T^* and given as

$$X_T^* = a - b\xi_T,$$

where a and b are non-negative and equal to

$$a = (W_0 + b\mathbb{E}[\xi_T^2]) e^{rT}, \quad b = \frac{s}{\sqrt{\text{var}(\xi_T)}}. \quad (4.3)$$

Proof. Let us choose a and $b \geq 0$ such that $X_T^* = a - b\xi_T$ satisfies the constraints of Problem (4.2), that is a and b verify $\text{var}(X_T^*) = s^2$ and $c(X_T^*) = W_0$. We find (4.3). Observe that $\text{corr}(X_T^*, \xi_T) = -1$ and X_T^* is thus the unique payoff that is perfectly negatively correlated with ξ_T while satisfying the stated constraints. Consider next any other strategy X_T which also verifies these constraints (but is not negatively linear in ξ_T). We find that

$$\text{corr}(X_T, \xi_T) = \frac{\mathbb{E}[\xi_T X_T] - \mathbb{E}[\xi_T]\mathbb{E}[X_T]}{\sqrt{\text{var}(\xi_T)}\sqrt{\text{var}(X_T)}} > -1 = \text{corr}(X_T^*, \xi_T). \quad (4.4)$$

⁴In other words, the optimal mean-variance portfolio must be cost-efficient in the sense defined in Chapter 1.

⁵ $\text{corr}(X_T, \xi_T) = \frac{\mathbb{E}[\xi_T X_T] - \mathbb{E}[\xi_T]\mathbb{E}[X_T]}{\text{std}(\xi_T)\text{std}(X_T)}$. As the moments of ξ_T are given, it follows that for given moments $\mathbb{E}[X_T]$ and $\text{std}(X_T)$ minimizing $\mathbb{E}[\xi_T X_T]$ is equivalent to minimizing correlation $\text{corr}(X_T, \xi_T)$.

Since $\text{var}(X_T) = s^2 = \text{var}(X_T^*)$ and $\mathbb{E}[\xi_T X_T] = W_0 = \mathbb{E}[\xi_T X_T^*]$ it follows from (4.4) that

$$\mathbb{E}[\xi_T] \mathbb{E}[X_T] < \mathbb{E}[\xi_T] \mathbb{E}[X_T^*],$$

which shows that X_T^* maximizes the expectation and thus solves Problem (4.2). \square

Note that the maximal expected return $\mathbb{E}[X_T^*]$ is given as

$$\mathbb{E}[X_T^*] = (W_0 + b \text{var}(\xi_T)) e^{rT} \geq W_0 e^{rT}, \quad (4.5)$$

because $b \geq 0$, reflecting that taking risk results in a positive risk premium. Proposition 4.2.1 shows that, in case all trading is permitted, mean-variance efficient portfolios amount to holding a cash amount ($a \geq 0$) combined with a short position ($b \geq 0$) in the stochastic discount factor ξ_T . While the benefits are limited to receiving a at the maximum, the investor suffers unlimited losses when the market index $S_T^* = 1/\xi_T$ goes to zero. The monotonicity of the optimal payoff with the market index may not suit investors who may wish to have some degree of protection for their portfolio in bad economic scenarios (i.e. in a declining market). Adding constraints on correlation or on the dependence can partly address this issue as we will see later in Section 4.3.

Rather than maximizing the expected return for a given (risk) variance, one can also consider the dual problem of minimizing the variance for a given expected return. We obtain the following result.

Proposition 4.2.2 (Mean-variance efficient portfolios - Dual setting). *Assume that the investor aims for a strategy that minimizes the variance for a given expected wealth. The solution of the following mean-variance optimization problem*

$$\begin{cases} \min & \text{var}(X_T) \\ \text{s.t.} & \mathbb{E}[X_T] = W_0 e^{mT} \\ & c(X_T) = W_0 \end{cases} \quad (4.6)$$

where $m \geq r$, is denoted by X_T^* and given as

$$X_T^* = a - b \xi_T,$$

where a and b are non-negative and equal to $a = W_0 e^{mT} + b e^{-rT}$, $b = \frac{W_0 (e^{(m-r)T} - 1)}{\text{var}(\xi_T)}$.

Note that the requirement $m \geq r$ is natural to ensure the problem is well-posed for mean-variance investors who prefer more to less. Otherwise, investing the initial wealth at the fixed risk-free rate r would always ensure higher final wealth with zero variance.

Mean-variance optimization is similar to maximizing quadratic utility. As the following result readily follows from Propositions 4.2.1 and 4.2.2, we omit its proof.

Corollary 4.2.3 (Optimizing quadratic utility). *Assume that the investor has a utility function $U(x) = x - \frac{\alpha}{2} x^2$ (where $0 < \alpha \leq \frac{e^{-rT}}{W_0}$ and $x \leq \frac{1}{\alpha}$)⁶. The solution of the following expected utility maximization problem*

$$\max_{c(X_T)=W_0} \mathbb{E}[u(X_T)] \quad (4.7)$$

⁶Note that imposing $\alpha \leq \frac{e^{-rT}}{W_0}$ is natural to ensure the investor is also considering (utility of) wealth levels $x > W_0 e^{rT}$, when optimizing her expected utility. In the opposite case, optimizing quadratic utility gives rise to portfolios that exhibit variability for a lower expected return than what can be obtained by simply investing at the fixed rate r . In addition, the requirement $x \leq \frac{1}{\alpha}$ is needed to ensure the utility function is non-decreasing on admissible wealth levels.

is denoted by X_T^* and is given by $X_T^* = a - b\xi_T$, where a and b are non-negative and given by $a = \frac{1}{\alpha}$ and $b = \frac{e^{-rT} - \alpha W_0}{\alpha \mathbb{E}[\xi_T^2]}$.

4.2.2 Example in the Black-Scholes Setting

Let us show how the results apply in a Black-Scholes market (Chapter 1, Section 1.2.2). While this setting does not always allow to accurately reflect true market behavior, its tractability makes it a reference model used in practice and a traditional work-horse employed in the finance literature to develop new ideas and get insights. The unique stochastic discount factor process $(\xi_t)_t$ in the Black-Scholes setting is given in (1.8) and recalled here

$$\xi_t = e^{-rt} e^{-\theta W_t - \frac{1}{2}\theta^2 t}, \quad \theta = \frac{\mu - r}{\sigma}. \quad (4.8)$$

Furthermore, we also find that $\frac{dS_t^*}{S_t^*} = \left(\frac{\theta}{\sigma}\mu + (1 - \frac{\theta}{\sigma})r\right) dt + \theta dW_t$ so that the market index S_t^* amounts to a constant mix strategy, where at time t a fraction $\frac{\theta}{\sigma}$ is invested in the risky asset and the remaining fraction $1 - \frac{\theta}{\sigma}$ in the bank account. The solution to the *mean-variance* optimization problem (4.2) is thus given by

$$X_T^* = a - \frac{b}{S_T^*}, \quad (4.9)$$

where a, b are non-negative and given by $a = W_0 e^{rT} + s \frac{e^{\theta^2 T}}{\sqrt{e^{\theta^2 T} - 1}}$ and $b = \frac{s e^{rT}}{\sqrt{e^{\theta^2 T} - 1}}$. It is clear that a similar result can be derived for Problem (4.6).

4.2.3 Maximum Sharpe Ratio and Application to Fraud Detection

The Sharpe ratio is a well-known measure balancing risk (variance) and return (mean) of a portfolio X_T (see Sharpe (1967)). It is defined as

$$SR(X_T) = \frac{\mathbb{E}[X_T] - W_0 e^{rT}}{\text{std}(X_T)}.$$

It is clear that mean-variance optimality of a portfolio is tied to the maximality of the Sharpe ratio.

Proposition 4.2.4 (Maximal Sharpe Ratio). *All mean-variance efficient portfolios X_T^* have the same maximal Sharpe ratio given by*

$$SR^* = \sqrt{\frac{\text{var}(\xi_T)}{\mathbb{E}^2[\xi_T]}} = e^{rT} \text{std}(\xi_T). \quad (4.10)$$

Furthermore, a portfolio is mean-variance efficient if and only if it has maximal Sharpe ratio.

In the specific context of a Black-Scholes market, Proposition 4.2.4 can also be found in Goetzmann et al. (2002)⁷ but our result is more general and holds in a fairly general market setting. This maximum Sharpe ratio applies for all mean-variance efficient strategies as well as any admissible

⁷In the appendix in Goetzmann et al. (2002) the maximal Sharpe ratio is proved to be $\exp(\theta^2 \Delta t) - 1$, where in the context of this chapter Δt is the investment horizon. Note that this result does not appear in the published version of Goetzmann et al. (2007).

trading strategy in the financial market. Hence it can be used for fraud detection in the sense that it allows to trace Sharpe ratio of fund managers which are too high to be “feasible”. When the Sharpe ratio of a strategy violates the upper bound SR^* , this can be interpreted as a signal that there might be a fraud, or a mistake regarding the reported returns. In practice, to apply fraud detection one needs to construct an estimator for the Sharpe ratio which can be achieved using the reported returns of the fund at hand. The upper bound itself requires the estimation of the standard deviation of the stochastic discount factor ξ_T . While this variable is not really observable in the market, several methods exist to estimate its distributional properties; see for example Aït-Sahalia and Lo (2001). We now pursue the basic idea of fraud detection in some more detail in the Black-Scholes setting.

In the Black-Scholes market one finds from (4.8) that $\mathbb{E}[\xi_T] = e^{-rT}$ and $\mathbb{E}[\xi_T^2] = e^{-2rT}e^{\theta^2 T}$. Hence, in this setting the expression (4.10) for the maximal Sharpe-ratio SR^* can be made explicit and we find that

$$SR^* = \sqrt{e^{\theta^2 T} - 1}. \quad (4.11)$$

Using observed market data one can now readily estimate SR^* and next compare Sharpe ratios of funds (derived from reported data) with this maximum. This was used by Bernard and Boyle (2009) to show that the returns from the option strategy pursued by Madoff were too good to be true.

The Black-Scholes setting does not always comply with real markets which implies that fraud detection based on the expression (4.11) is prone to some model error. A possible way to address this drawback consists in a better use of the available market data allowing to obtain a non-parametric estimator for SR^* . The following proposition exhibits the Sharpe ratio in terms of observable option prices.

Proposition 4.2.5 (Fraud detection). *Assume that $\xi_T = f(S_T)$ (where f is typically decreasing and S_T is the risky asset) and that all European call options on the underlying S_T maturing at $T > 0$ are traded.⁸ Let $C(K)$ denote the price of a call option on S_T with strike K . Then, the Sharpe ratio $SR(X_T)$ of any admissible strategy with payoff X_T satisfies the following upper bound*

$$SR(X_T) \leq \sqrt{e^{2rT} \int_0^{+\infty} f(K) \frac{\partial^2 C(K)}{\partial K^2} dK} - 1. \quad (4.12)$$

Note that in an incomplete market, the pricing kernel ξ_T is not unique anymore. The upper bounds (4.10) and (4.12) are then still valid upper bounds for the Sharpe ratio of a given strategy X_T (with respect to the pricing kernel ξ_T that is used to price this strategy). However, these bounds may not be attainable anymore because the market is incomplete. In other words, a dynamic strategy that achieves the maximum Sharpe ratio may not exist (if $a - b\xi_T$ is not attainable). The fact that $a - b\xi_T$ is not attainable is not a real problem for the application to fraud detection since it is still an upper bound.

It is important to understand that the violation of the upper bound as derived in Proposition 4.2.5, or as given in (4.11) for the specific case of a Black-Scholes model, has to be seen as an indication (a signal) that there could be a fraud, but not as a formal proof of it. We already

⁸There is often a large (but finite) number of options available in the market, in which case assuming a continuum of strikes could be seen as a reasonable approximation of reality. Carr and Chou (1997) note that it is “analogous to the continuous trading assumption permeating the continuous time literature”. See also Breeden and Litzenberger (1978).

explained that the calculation of the upper bound is subject to model error, thus a violation of this upper bound does not always imply a fraud. There might be other reasons that explain why an observed manager's Sharpe ratio can be higher than the upper bound even if there is no fraud.

Firstly, the Sharpe ratio of a strategy is based on the average and the standard deviation. In practice, one has to estimate these two moments and hence only confidence intervals for the Sharpe ratio can be obtained but not the true value. Hence, the observed Sharpe ratio may violate the upper bound but this fact is not necessarily (statistically) meaningful.⁹ Secondly, the upper bound that we propose has been derived in a continuous time framework by optimizing over all self-financing strategies that are adapted to the market information on the security prices. In practice, the manager may be able to capture information outside the financial market and to use this extra information to optimize his portfolio. In particular, a manager with insider information has a strategy that is also (partially) based on this “future information” and thus not adapted to the prices filtration. In this case, the Sharpe ratio of such strategy may lie above the theoretical bound (that is derived among all adapted self-financing strategies). A violation of the upper bound may thus also serve as a potential signal for detecting insider trading. Similarly, a manager who is rebalancing his portfolio on a daily basis may use an extra source of information to decide if he goes long or short in the risky asset for the next business day. There is then always a chance that he consistently makes the right decision and thus achieves a higher Sharpe ratio (by luck). In the context of a multidimensional financial market with n stocks, there is a small probability to always pick the best stocks while rebalancing the portfolio. Being an outstanding stock picker was one of the arguments used by Madoff to explain the exceptionally high Sharpe ratio of his investment strategy.

To conclude, it is important to understand that the violation of the upper bound derived in Proposition 4.2.5 or given in (4.11) for the specific case of a Black-Scholes model should be seen as a signal that there could be a fraud, but not as a proof.

4.3 Mean-Variance Efficiency with a Correlation Constraint

The fraud detection mechanism described in the previous section can be greatly improved by taking into account additional information available in the market. Indeed, the maximum Sharpe ratios derived in Propositions 4.2.4 and 4.2.5 do not take into account the dependence features between the investment strategy and the financial market. Regulators, however, can estimate the Sharpe ratio of a hedge fund but can also investigate correlations of the fund with indices in the market, and this additional source of information may be useful for refining the process of fraud detection. A popular strategy amongst hedge funds is the so-called “market-neutral” strategy. One of its key properties is that it typically ensures very low correlation with market indices. We show that in this instance the maximum possible Sharpe ratio is significantly reduced compared to the unconstrained case.

Furthermore, recall that unconstrained mean-variance efficient payoffs have bounded gains but no protection against a market crash giving rise to unlimited losses (Propositions 4.2.1 and 4.2.2). An investor may then be interested in choosing the dependence with a benchmark of her choice so

⁹It is well-documented in the literature that optimal portfolio choice is subject to parameter uncertainty, in that a small perturbation of the inputs may lead to a large change in the optimal portfolio. Robust portfolio selection techniques have been developed to deal with this issue (Ben-Tal and Nemirovski (1998), Goldfarb and Iyengar (2003), Tütüncü and Koenig (2004)). For a complete discussion of robust portfolio optimization and the associated solution methods, see Fabozzi et al. (2007) and the references therein.

that, for instance, her investment strategy is no longer decreasing when the market index decreases. It appears as a suitable reference point for her investment, and it allows to better control the states in which cash-flows are received. The worst outcomes for the strategy do no longer necessarily coincide with bad scenarios for the market index.

Hence, in this section we focus on investors who care about the first two moments of the distribution of a strategy *and* additionally aim at obtaining a desired correlation (Section 4.3) or dependence (Section 4.4) with a benchmark asset B_T .

4.3.1 Mean-Variance Efficiency

The following proposition gives a mean-variance optimal allocation policy for an investor with a fixed correlation constraint.

Proposition 4.3.1 (Optimal portfolio with a correlation constraint). *Let B_T be a benchmark which is linearly independent from ξ_T with $\text{var}(B_T) < +\infty$. Let $|\rho| < 1$ and $s > 0$. Consider the following mean-variance optimization problem*

$$\begin{aligned} & \max \mathbb{E}[X_T] \\ & \begin{cases} \text{var}(X_T) = s^2 \\ c(X_T) = W_0, \\ \text{corr}(X_T, B_T) = \rho \end{cases} \end{aligned} \quad (4.13)$$

Let a , b and c be uniquely determined by the set of equations

$$\begin{aligned} \rho &= \text{corr}(cB_T - \xi_T, B_T) \\ s &= b\sqrt{\text{var}(\xi_T - cB_T)} \\ W_0 &= ae^{-rT} - b(E[\xi_T^2] - cE[\xi_TB_T]), \end{aligned}$$

then X_T^* given as $X_T^* = a - b(\xi_T - cB_T)$, is a solution to Problem (4.13).

Proof. Consider the function $f(c) := \text{corr}(cB_T - \xi_T, B_T)$ and observe that $\lim_{c \rightarrow -\infty} f(c) = -1$ and $\lim_{c \rightarrow +\infty} f(c) = 1$. Since $f(c)$ is continuous it follows from the intermediate value theorem that the equation $\rho = f(c)$ has a solution when solving for c . Moreover, this solution is unique as

$$f'(c) = \frac{\text{var}(B_T)\text{var}(\xi_T) - (\text{cov}(\xi_T, B_T))^2}{(\text{std}(cB_T - \xi_T))^3 \text{std}(B_T)} > 0.$$

It is now clear that $X_T^* = a - b(\xi_T - cB_T)$ is the unique payoff that is linear in $\xi_T - cB_T$ while satisfying all constraints and note that $b > 0$. It remains to show that it is the optimal solution to Problem (4.13). Hence, consider any other payoff X_T that satisfies the constraints and which is non-linear in $\xi_T - cB_T$. We have that

$$\text{corr}(X_T, \xi_T - cB_T) = \frac{\mathbb{E}[X_T(\xi_T - cB_T)] - \mathbb{E}[\xi_T - cB_T]\mathbb{E}[X_T]}{\text{std}(\xi_T - cB_T)\text{std}(X_T)} > -1 = \text{corr}(X_T^*, \xi_T - cB_T)$$

Since both X_T and X_T^* satisfy the constraints we have that $\text{std}(X_T) = \text{std}(X_T^*)$, $\mathbb{E}[X_T\xi_T] = \mathbb{E}[X_T^*\xi_T]$ and $\text{cov}(X_T, B_T) = \text{cov}(X_T^*, B_T)$. Hence the inequality holds true if and only if $\mathbb{E}[X_T^*] > \mathbb{E}[X_T]$. This ends the proof. \square

Remark 4.3.1. In the Problems (4.2), (4.6) and (4.13), the equalities in the constraints can sometimes be replaced by inequalities without impacting the solution. First, the solution to Problem (4.2) is also the portfolio, which maximizes the expectation such that the variance is less than s^2 instead of being equal to s^2 (as we do). This feature clearly appears from the solution $a - b\xi_T$, as b decreases if s decreases and thus the expected final value of the portfolio decreases as well. Similarly, in the dual problem (4.6), the solution that minimizes the variance of the terminal value of the portfolio for a given expected value is also the solution to a variance minimization with minimum expected value. However, for the correlation constraint in Problem (4.13), replacing the equality by an inequality usually affects the solution. For example, let us assume that we replace the constraint “ $\text{corr}(X_T, B_T) = \rho$ ” by the constraint “ $\text{corr}(X_T, B_T) \leq \rho$ ” and also assume that the optimum $a - b\xi_T$ that was derived in the unconstrained case (as a solution to Problem (4.2)) satisfies this inequality constraint. Clearly, the unconstrained optimum $a - b\xi_T$ (for Problem (4.2)) will then also be the optimum for the constrained Problem (4.13). However, it is clear that one can also have that $\text{corr}(a - b\xi_T, B_T) > \rho$ and in this instance the unconstrained solution fails to be the optimum of the constrained Problem (4.13). Note also that it is possible to solve Problem (4.13) in presence of inequality constraints on the correlation, i.e. when $\rho_1 \leq \text{corr}(X_T, B_T) \leq \rho_2$ for some $-1 \leq \rho_1 \leq \rho_2 \leq 1$. Finding the optimal strategy consists of two-step optimization in which we derive first the optimum, say X_T^ρ , of Problem (4.13) for each $\rho_1 \leq \rho \leq \rho_2$ and next we find the optimal solution by optimizing $\mathbb{E}[X_T^\rho]$ on the interval $[\rho_1, \rho_2]$.

Note that in the case when the benchmark B_T is not a function of the stochastic discount factor ξ_T (or equivalently of the market index S_T^*), the worst case outcomes for the optimal strategy $X_T^* = a - b(\xi_T - cB_T)$ do not necessarily occur when the market index is low. Bad outcomes for the market index S_T^* might be compensated by good outcomes for the benchmark asset. In contrast, when B_T and thus also X_T^* depends on S_T^* only, then the optimum may or may not provide protection against downturns (depending on its precise functional relationship with the market index). We now provide an illustration of the theoretical results.

Consider a Black-Scholes market. Let us solve Problem (4.13) when the benchmark B_T is the market index, that is $B_T = S_T^*$. From Proposition 4.3.1, the optimal solution is of the form $X_T^* = a - b(\xi_T - cS_T^*)$, where c is computed from the equation $\rho = \text{corr}(cS_T^* - \xi_T, S_T^*)$, b is derived from $b = \frac{s}{\sqrt{\text{var}(\xi_T - cS_T^*)}}$ and $a = W_0 e^{rT} + b(e^{-2rT + \theta^2 T} - c)e^{rT}$.

Figure 4.3.1 shows that constraining the strategy to be (sufficiently) negatively correlated can improve the returns during a crisis (when S_T^* is low). However this goes at the cost of a lower performance when the market index increases significantly. Note that the unconstrained optimal strategy also appears as an optimal strategy in presence of a (redundant) correlation constraint, that is when $\rho = \text{corr}(a - b\xi_T, S_T^*)$ (which is close to 1 in this example).

4.3.2 Maximum Sharpe Ratio and Application to Fraud Detection

It is clear that adding constraints reduces degrees of freedom for trading strategies and thus lowers the maximum possible Sharpe ratio. We obtain the following proposition.

Proposition 4.3.2 (Constrained Maximal Sharpe Ratio). *Any mean-variance efficient portfolio X_T^* which satisfies the additional constraint $\text{corr}(X_T^*, B_T) = \rho$ with a benchmark asset B_T (that is not linearly dependent to ξ_T) has the same maximum Sharpe ratio SR_ρ^* given by*

$$SR_\rho^* = e^{rT} \frac{\text{cov}(\xi_T, \xi_T - cB_T)}{\text{std}(\xi_T - cB_T)} \leq SR^* = e^{rT} \text{std}(\xi_T). \quad (4.14)$$

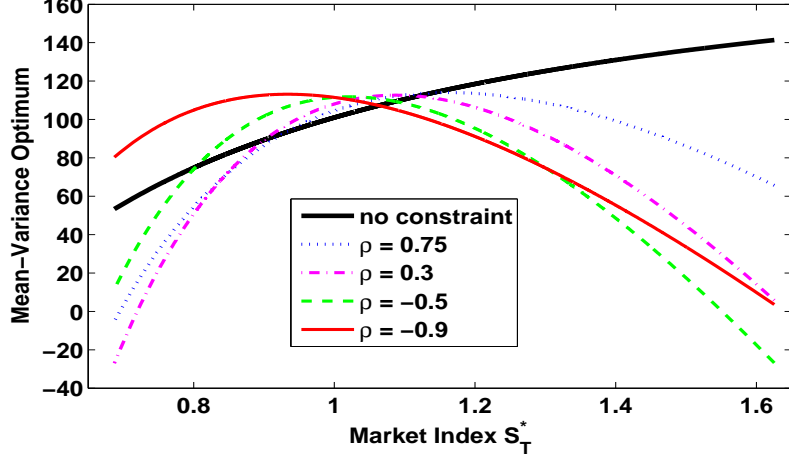


Figure 4.3.1: Optimal payoffs as a function of the market index S_T^* , for different values of the correlation ρ with the benchmark S_T^* using the following parameters: $W_0 = 100$, $r = 0.05$, $\mu = 0.07$, $\sigma = 0.2$, $T = 1$, $S_0 = 100$, $s = 10$.

where SR^* is the unconstrained Sharpe ratio found in Proposition 4.2.4 and where c is determined uniquely by the equation $\text{corr}(\xi_T - cB_T, B_T) = \rho$.

Note that the constrained strategy exhibits a strictly lower Sharpe ratio unless $c = 0$ which just means that the unconstrained optimal strategy happens to satisfy the correlation constraint. The proposition shows that for fraud detection it is useful to incorporate correlation features of displayed returns. We illustrate this point in a numerical example in the next paragraph in the Black-Scholes setting. In particular, we observe that maximum Sharpe ratios can be significantly lower than in the unconstrained case.

Figure 4.3.2 displays the maximum Sharpe ratio for the unconstrained case (which is approximately equal to 0.1 for the parameter set used in Figure 4.3.2) and constrained Sharpe ratio for different levels of correlation constraints when the market index S_T^* is the benchmark.

Observe that for low correlation levels ($\rho \in [-0.1, 0.1]$), the maximum Sharpe ratio is 0.02 only, which is five times lower than when there are no constraints. Adding the information on the dependence between the strategy's performance and the market index can thus greatly improve fraud detection. Observe also that for negative correlation levels, the maximum Sharpe ratio can be negative¹⁰, thus if hedge fund returns display a negative correlation with the financial market and a positive Sharpe ratio, then there could be some suspicion about these returns. This is strongly different from what we observed in the unconstrained case as the maximum Sharpe ratio is always positive in this case because of (4.5).

Finally, observe that the constrained case reduces to the unconstrained one when ρ is equal to the correlation of the unconstrained optimum $a - b\xi_T$ with the market index S_T^* (which happens when ρ is close to 1). The constraint is thus redundant in that case.

¹⁰At first, it might look counter intuitive that an optimal strategy has a lower return than what can be achieved risk-free. However, enforcing a negative dependence with the market comes at some cost. A similar observation can be drawn for the put option: it has a low expected return but provides protection when markets fall.

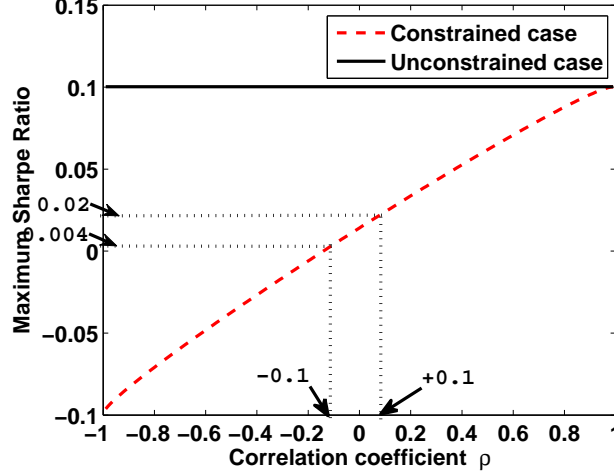


Figure 4.3.2: Maximum Sharpe ratio SR_ρ^* given by (4.14) for different values of the correlation ρ when the benchmark is $B_T = S_T^*$. We use the following parameters: $W_0 = 100$, $r = 0.05$, $\mu = 0.07$, $\sigma = 0.2$, $T = 1$, $S_0 = 100$.

In practice, investors may consider more than one benchmark when making investment decisions. Likewise, a fund manager's portfolio returns are observed in conjunction with the returns of many other market indices. In particular, one may have at hand a correlation matrix between a given fund and n market indices, and having this information allows to develop an improved upper bound on the Sharpe ratio. In this section, we derive optimal strategies (and thus also the maximum Sharpe ratio) amongst strategies that satisfy the correlation matrix. Using not only information about the marginal distribution (computed as the Sharpe ratio of the portfolio) but also the correlation matrix significantly improves the upper bound, and thus may facilitate fraud detection. More details on the case of multiple correlation constraints can be found in the published version of this chapter in the *European Journal of Operational Research*.

4.4 Mean-Variance Optimal Portfolios with a Dependence Constraint

Correlation is only one property related to dependence. It measures the linear relationship between strategies but falls short in depicting dependence fully. A useful device for reflecting the interaction between the strategy's payoff X_T and B_T is the copula. Indeed Sklar's theorem shows that the joint distribution of (B_T, X_T) can be decomposed as

$$\mathbb{P}(B_T \leq y, X_T \leq x) = C(F_{B_T}(y), F_{X_T}(x)), \quad (4.15)$$

where C is the joint distribution (also called the copula) for a pair of uniform random variables U and V over $(0, 1)$ and where F_{B_T} and F_{X_T} denote respectively the (marginal) cdf of B_T and X_T . Hence, the copula $C : [0, 1]^2 \rightarrow \mathbb{R}$ fully describes the interaction between the strategy's payoff X_T and the benchmark B_T . A constraint on the copula is much more informative than a correlation constraint. A copula is a function and contains full information about the interaction between two variables. A correlation is a single number and cannot describe the complex nature of dependence

fully.¹¹ Similarly, Value-at-Risk is only a single number and will never describe the risk as well as knowing the exact loss distribution.

The partial derivative $c_u(v) := \frac{\partial}{\partial u} C(u, v)$ has an interesting property, namely it can be interpreted as a conditional probability:

$$c_u(v) = \mathbb{P}(V \leq v | U = u). \quad (4.16)$$

Property (4.16) is extremely useful for constructing payoffs with desired dependence properties. For example, $c_U(V)$ is a uniform variable that depends on U and V and which is independent of U . Conversely, if U and T are independent uniform random variables then $c_U^{-1}(T)$ is a uniform variable (depending on U and T) that has copula C with U .

The following propositions give mean-variance optimal allocation schemes in the presence of a benchmark. We use the notation $C(\cdot, \cdot)$ to reflect the desired dependence for the couple (B_T, X_T^*) that is between the benchmark B_T and the final value of the investment strategy X_T^* at the investment horizon T .

The next proposition provides mean-variance efficient portfolios when the market index S_T^* , and thus ξ_T , is used as the benchmark.

Proposition 4.4.1 (Optimal portfolio when the market index is the benchmark). *Let $B_T = \xi_T$. For $t \in (0, T)$, define the variable \mathcal{A}_t as*

$$\mathcal{A}_t = \left(c_{F_{\xi_T}(\xi_T)} \right)^{-1} \left[j_{F_{\xi_T}(\xi_T)}(F_{\xi_t}(\xi_t)) \right],$$

where the functions $j_u(v)$ and $c_u(v)$ are defined as the first partial derivative for $(u, v) \rightarrow J(u, v)$ and $(u, v) \rightarrow C(u, v)$ respectively, and where J denotes the copula for the random pair (ξ_T, ξ_t) . Assume that $\mathbb{E}[\xi_T | \mathcal{A}_t]$ is strictly decreasing in \mathcal{A}_t . For $s > 0$, a solution of the following constrained mean-variance optimization problem

$$\begin{cases} \max & \mathbb{E}[X_T] \\ \text{var}(X_T) = s^2 \\ c(X_T) = W_0 \\ C : \text{copula between } X_T \text{ and } B_T \end{cases} \quad (4.17)$$

is given by X_T^* ,

$$X_T^* = a - b\mathbb{E}[\xi_T | \mathcal{A}_t]. \quad (4.18)$$

Here a, b are non-negative and given by $a = (W_0 + b\mathbb{E}[\xi_T \mathbb{E}[\xi_T | \mathcal{A}_t]]) e^{rT}$, $b = \frac{s}{\text{std}(\mathbb{E}[\xi_T | \mathcal{A}_t])}$.

Note that the maximal expected return of X_T^* is given as

$$\begin{aligned} \mathbb{E}[X_T^*] &= (W_0 + b\text{var}(\mathbb{E}[\xi_T | \mathcal{A}_t])) e^{rT} \\ &> W_0 e^{rT}. \end{aligned} \quad (4.19)$$

Actually, comparing (4.19) with (4.5), one observes that adding dependence constraints decreases the expected return of the optimal strategy and this decrease is directly proportional to $\text{var}(\xi_T) - \text{var}(\mathbb{E}[\xi_T | \mathcal{A}_t])$.

¹¹Note that even when two variables are normally distributed their dependence is not described by their correlation coefficient. Knowledge of the correlation coefficient is sufficient for depicting the dependence when the variables follow a bivariate normal distribution, but this is an assumption which actually imposes a lot of structure on the interaction between both variables.

We now formulate a few important remarks. (i) In general, the Mean-Variance optimal portfolios discussed in this section are not unique as the choice of t is arbitrary. (ii) One may wonder what happens if in Proposition 4.4.1 the non-decreasingness property for $\mathbb{E}[\xi_T|\mathcal{A}_t]$ is not fulfilled. Then, it still follows that mean-variance efficient portfolios, provided they exist, must write as $f(\mathcal{A}_t)$ for some non-decreasing f . Unfortunately, it is then not clear how to find a function f that minimizes $\text{corr}(f(\mathcal{A}_t), \xi_T)$, or equivalently, that minimizes $\text{corr}(f(\mathcal{A}_t), \mathbb{E}[\xi_T|\mathcal{A}_t])$. (iii) When $\mathbb{E}[\xi_T|\mathcal{A}_t]$ is increasing in \mathcal{A}_t then it holds for the mean-variance efficient portfolio $f(\mathcal{A}_t)$ (provided it exists) that $\mathbb{E}[f(\mathcal{A}_t)] \leq W_0 e^{rT}$, meaning that the imposed benchmark constraint comes at significant cost.

We remark from the previous proposition that fixing the dependence with ξ_T does not generally result in unique mean-variance efficient portfolios. However, when the benchmark is not functionally dependent with ξ_T , the optimal allocation can become unique as the following proposition shows.

Proposition 4.4.2 (Constrained Mean-Variance Efficiency). *Let $s > 0$. Assume that the benchmark B_T has a joint density with ξ_T . Define the variable \mathcal{A} as*

$$\mathcal{A} = \left(c_{F_{B_T}(B_T)} \right)^{-1} \left[j_{F_{B_T}(B_T)}(1 - F_{\xi_T}(\xi_T)) \right], \quad (4.20)$$

where the functions $j_u(v)$ and $c_u(v)$ are defined as the first partial derivative for $(u, v) \rightarrow J(u, v)$ and $(u, v) \rightarrow C(u, v)$ respectively, and where J denotes the copula for the random pair (B_T, ξ_T) . Furthermore, assume that $\mathbb{E}[\xi_T|\mathcal{A}]$ is decreasing in \mathcal{A} . Then, the solution to the problem

$$\begin{cases} \max & \mathbb{E}[X_T] \\ \text{var}(X_T) = s^2 \\ c(X_T) = W_0 \\ C : \text{copula between } X_T \text{ and } B_T \end{cases} \quad (4.21)$$

is uniquely given as

$$X_T^* = a - b\mathbb{E}[\xi_T|\mathcal{A}],$$

where a, b are non-negative and equal to $a = (W_0 + b\mathbb{E}[\xi_T\mathbb{E}[\xi_T|\mathcal{A}]])e^{rT}$ and $b = \frac{s}{\text{std}(\mathbb{E}[\xi_T|\mathcal{A}])}$.

Similarly as the case when the benchmark is the market index, the maximum expected return verifies

$$\mathbb{E}[X_T^*] = (W_0 + b\text{var}(\mathbb{E}[\xi_T|\mathcal{A}]))e^{rT} \geq W_0 e^{rT}.$$

Let us consider a Black-Scholes Market and specialize the results by considering a Gaussian dependence¹² and by taking the market index as the benchmark.

Proposition 4.4.3 (Case $B_T = S_T^*$). *Let $B_T = S_T^*$. Assume that $\rho_0 \geq 0$. Then, the solution to the problem (4.17) when the copula C is the Gaussian copula with correlation ρ_0 , i.e. $C_{\rho_0}^{\text{Gauss}}$, is given by X_T^* ,*

$$X_T^* = a - bG_T^c. \quad (4.22)$$

Here G_T is a weighted average of the benchmark and the market index. It is given as $G_T = (S_t^*)^\alpha S_T^*$ with α ,

$$\frac{1}{\alpha} = \frac{t}{T} \left(\rho_0 \sqrt{\frac{T-t}{t} \frac{1}{1-\rho_0^2}} - 1 \right),$$

¹²We say that two variables X and Y have a gaussian dependence with correlation coefficient ρ_0 if (X, Y) is distributed as $(f(N), g(M))$, where f and g are increasing, and where N and M are (bivariate) standard normally distributed random variables that exhibit a correlation coefficient equal to ρ_0 .

and where the coefficients a , b and c are given as

$$a = W_0 e^{rT} + b e^{rT} \mathbb{E}[\xi_T G_T^c], \quad b = \frac{s}{\sqrt{\text{var}(G_T^c)}}, \quad c = -\frac{\alpha t + T}{(\alpha + 1)^2 t + (T - t)}.$$

Proposition 4.4.4 (Case $B_T = S_t^*$). *Let $B_T = S_t^*$ ($0 < t < T$) be the benchmark. Assume that $\rho_0 \geq -\sqrt{1 - \frac{t}{T}}$. Then, the solution to Problem (4.21) when the copula C is the Gaussian copula with correlation ρ_0 , $C_{\rho_0}^{\text{Gauss}}$ is given by X_T^* ,*

$$X_T^* = a - b G_T^c. \quad (4.23)$$

Here G_T is a weighted average of the benchmark and the market index. It is given as $G_T = (S_t^*)^\alpha S_T^*$ with α ,

$$\alpha = \rho_0 \sqrt{\frac{T-t}{t} \frac{1}{1 - \rho_0^2}} - 1,$$

where the coefficients a , b and c are given as

$$a = W_0 e^{rT} + b e^{rT} \mathbb{E}[\xi_T G_T^c], \quad b = \frac{s}{\sqrt{\text{var}(G_T^c)}}, \quad c = -\frac{\alpha t + T}{(\alpha + 1)^2 t + (T - t)}.$$

The Black-Scholes setting allows to derive an explicit expression for the Maximal Sharpe ratio. We obtain the following result.

Proposition 4.4.5 (Constrained Maximal Sharpe Ratio). *All mean-variance efficient portfolios X_T^* which satisfy the additional constraint that the copula between their final value X_T^* and the market portfolio S_t^* (for all $t \leq T$) is the Gaussian copula, have the same maximal Sharpe ratio $SR_{\rho_0, G}^*$ given by*

$$SR_{\rho_0, G}^* = \frac{e^{rT} \mathbb{E}[\xi_T G_T^c] - \mathbb{E}[G_T^c]}{\text{std}(G_T^c)} \leq SR^*. \quad (4.24)$$

Here $\mathbb{E}[G_T^c] = e^{M + \frac{V}{2}}$ and $\text{var}(G_T^c) = (e^V - 1)e^{2M + V}$, with $M := \mathbb{E}[\ln(G_T^c)] = c(r + \frac{\theta^2}{2})(\alpha t + T)$ and $V := \text{var}(\ln(G_T^c)) = c^2 \theta^2 (\alpha^2 t + T + 2\alpha t)$. Moreover $\mathbb{E}[\ln(\xi_T) + \ln(G_T^c)] = M - rT - \frac{\theta^2}{2}T$ and $\text{var}(\ln(\xi_T) + \ln(G_T^c)) = \theta^2 (c^2 \alpha^2 t + (c - 1)^2 T + 2c(c - 1)\alpha t)$ so that $\mathbb{E}[\xi_T G_T^c]$ reflects the expectation of a lognormal which can be computed from these two first log-moments similarly as we did for G_T^c .

Proof. We omit the proof as it is a straightforward calculation. \square

Remark 4.4.1. Proposition 4.4.5 is fundamentally different from the result in Proposition 4.3.2, which only derived the maximum Sharpe ratio with a given correlation. Imposing that the dependence between the portfolio and a benchmark is specified by a Gaussian copula with given correlation coefficient is more informative than imposing a correlation constraint as done in Proposition 4.3.2. There are indeed infinitely many copulas that can result in the same level of correlation and the Gaussian copula is just one example. In practice it is easier to estimate the correlation than the copula and Proposition 4.3.2 might thus be more useful for practical application. However, if one has a good idea of the desired copula, the result in Proposition 4.4.5 will be strictly better in the sense that the upper bound will be smaller and thus the fraud detection scheme will be improved.

Observe that the constrained case reduces to the unconstrained maximum Sharpe ratio when the correlation in the Gaussian copula is $\rho_0 = \sqrt{t/T}$. The reason is that the copula between the unconstrained optimum and S_t^* is the Gaussian copula with correlation $\rho_0 = \sqrt{t/T}$. The constraint is thus redundant in that case. Let us illustrate this last proposition by a numerical example in the Black-Scholes setting.

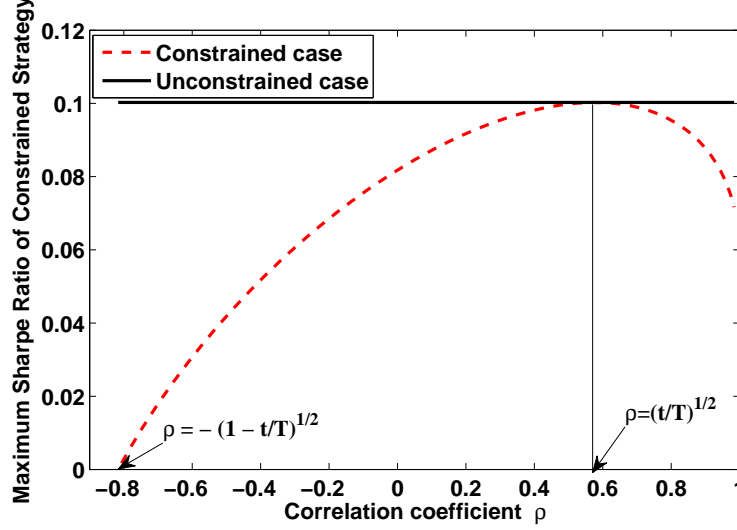


Figure 4.4.1: Maximum Sharpe ratio $SR_{\rho_0, G}^*$ given by (4.24) for different values of the correlation ρ_0 when the benchmark is $B_T = S_t^*$. We use the following parameters: $t = 1/3$, $\sqrt{t/T} = 0.577$, $-\sqrt{1-t/T} = -0.816$, $W_0 = 100$, $r = 0.05$, $\mu = 0.07$, $\sigma = 0.2$, $T = 1$, $S_0 = 100$.

4.5 Final Remarks

In this chapter we first analyze mean-variance optimal portfolios when there are no constraints on trading and dependence. We show that optimal portfolios consist in having a short position in the stochastic discount factor. Next, we depart from the classical setting and assume the investor also seeks for portfolios satisfying dependence constraints. We are able to provide optimal portfolios and also derive bounds for maximum possible Sharpe ratios. Throughout the chapter we explore how the results can be useful for fraud detection.

Throughout the chapter we dealt with several examples of state-dependent constraints by looking at the dependence between the portfolio and the benchmark. The static setting considered in the chapter allows to solve for explicit optimal strategies in the presence of these constraints. It is not straightforward to handle such constraints in a time-consistent setting (Basak and Chabakauri (2010)) and we leave it for future research. Another interesting future direction is to improve optimal mean-variance hedging using our results.

Chapter 5

Optimal Portfolios under Worst-Case Scenarios

This chapter is an extract from Bernard, Chen, and Vanduffel (2014) published in *Quantitative Finance*. This paper was awarded the 2015 Redington prize from the Society of Actuaries. All proofs can be found in the appendices of the published paper.

The Expected Utility Theory (EUT) introduced by von Neumann and Morgenstern (1947) has had a long-standing influence on research in economics and finance. As a descriptive theory for making choices under uncertainty, however, it has proven problematic. Indeed, many experiments (e.g., the Friedman and Savage puzzle (1948), the paradox of Allais (1953)) and numerous empirical studies have shown discrepancies between the predictions of the theory regarding optimal and therefore anticipated investor behavior, and observed behavior (i.e., what people actually do). In attempts to better capture systematic behavioral departures from the expected utility framework, economists and psychologists have proposed a variety of alternatives to EUT, including Dual Theory (Yaari (1987)), Rank Dependent Utility Theory (Quiggin (1993)), Cumulative Prospect Theory¹ (Tversky and Kahneman (1992)), and SP/A Theory (Lopes (1987), Lopes and Oden (1999), Shefrin and Statman (2000)).

While these theories differ significantly from each other with respect to their basic motivations and assumptions, they are all consistent with respect to *law-invariant* and *increasing* preferences.² In all of the aforementioned theories, for an investor with a fixed horizon without intermediate consumption, the optimal strategy is always decreasing in the state price process at maturity (Beare (2011), Bernard, Boyle, and Vanduffel (2014), Carlier and Dana (2011), Dybvig (1988b), Rieger (2011)). This is equivalent to saying that optimal strategies are increasing in the Growth Optimal Portfolio (GOP) when using the GOP as numeraire (Long (1990), Luenberger (1997)). The GOP is a portfolio with the property that it will almost surely accumulate more wealth than any other strategy at an infinite horizon, and hence it can be seen as a representation of “the

¹Kahneman and Tversky (1979) develop Prospect Theory, which offers a psychophysical explanation for decision making. Since Prospect Theory violates first-order stochastic dominance (“an assumption that many theorists were reluctant to give up”, Tversky and Kahneman ((1992), p. 299)), the authors propose Cumulative Prospect Theory, which can be seen as a technical modification of Prospect Theory.

²Law invariant preferences means that investors care only about the distribution of final wealth and not about the states in which the flows are received. Preferences are increasing when x dollars are preferred over y dollars ($x > y$). Most decision theories satisfy these assumptions. Note also that preferences are law-invariant and increasing if and only if first-order stochastic dominance is respected (Chapter 2, Bernard, Chen, and Vanduffel (2015)).

market”. For instance, Platen and Heath (2006) show that the GOP can usually be approximated by a well diversified portfolio. This finding allows us to conclude that in all decision theories with law-invariant and increasing preferences, optimal portfolios are always (approximately) long with a diversified market portfolio. This conclusion elucidates the role of diversification in optimal portfolio selection and shows that traditional diversification strategies, such as constant-mix strategies and buy-and-hold portfolios, are (approximately) optimal for rational investors. Nevertheless, we argue that optimal portfolios do not satisfy the needs of all investors.

First, investors may need to consume more in a crisis than in other states of the economy. Indeed, a recession increases chances of unemployment, higher taxes and less social security. Consuming more in times of scarcity provides protection against such potential costs, which often accompany worst-case scenarios. However, the lowest outcomes for an optimal strategy exactly occur when the market is at the lowest (e.g., during a recession). The recent financial crisis drew attention to this problem. As a *New York Times* (2011) article explained, “(...) an increasing number of investors now want protection for financial end times”. Many of the surveyed fund managers for the article are seeking to offer investors protection against financial doomsdays. To fill this need, they are working to design new strategies that optimize outcomes in worst states of the economy (e.g., “As the stock markets fell, a tail risk or black swan fund would profit (...”). In the financial literature, diversification is often seen as the best way to guard against market crashes. However, when a crisis occurs, seemingly unrelated assets are usually far more correlated than predicted and fall short of providing downside protection. This is consistent with the observations of Amenc and Martellini (2011) who are adamant that diversification and risk management are not the same. In response, hedge funds have created so-called Armageddon funds, which bet on financial disaster. As this strategy is profitable when the market is crashing but not otherwise, it cannot be explained using standard theories. Nonetheless, these funds have attracted significant resources, suggesting that they are useful for some investors.

Secondly, the 2008 financial crisis has led regulators and governments to set additional capital requirements for banks, in part to reflect the impact that their failure would have on the rest of the economy. For example, Basel III imposes extra capital requirements on financial institutions for their contribution to systemic risk.³ Such actions, however, raise the awareness of (institutional) investors about worst states of the economy, as their capital requirements are then linked to how their portfolios perform in such states. Unfortunately, optimal strategies for decision theories consistent with “law-invariant” and “increasing” preferences share the same states in which the outcomes are the lowest. As many (traditional) investment portfolios are often constructed as (proxies for) optimal portfolios, they are strongly mutually dependent. In this sense, optimal strategies may contribute to the simultaneous failures of companies, and thus to systemic risk. The fact that joint failures of corporates can arise from the correlation of returns on the asset side of their balance sheets is also acknowledged in Acharya (2009), Wagner (2009) and Ibragimov et al. (2009). Acharya (2009) cautions against a blanket call for more capital in the financial industry and instead recommends that “regulating each bank as a function of both its joint (correlated) risk with other banks [and] its individual (bank-specific) risk”. The basic principle here is that while the risk of a failure of a business can never be eliminated, it is not acceptable to have multiple large-scale failures occurring simultaneously.

The two arguments described above correspond to situations in which investors seek “security” in specific states of the economy. To the best of our knowledge, this feature was first addressed

³Basel III was supposed to come into force on January 1, 2013, but both European and U.S. regulators delayed application of the new capital requirements. January 1, 2014, is now widely seen as the target date for the new rules to take effect.

in the safety-first approach developed by Roy (1952) and next, more fundamentally, in the SP/A theory (Lopes (1987), Lopes and Oden (1999), Shefrin and Statman (2000)). The “safety-first” criterion refers to a risk management technique that allows selecting one portfolio over another based on the criterion that the probability of the portfolio return falling below a minimum desired threshold is minimized. In the SP/A theory, there is a balance between the desires for security (S) and potential (P), as well as with the aspiration (A) associated with each of these aims. Security relates to a concern (fear) of avoiding low levels of wealth, and potential relates to a desire (hope) to reach high levels of wealth. As explained in Shefrin and Statman (2000), SP/A Theory is well-suited to developing a behaviorally-based framework for portfolio selection, a framework that is also referred to as Behavioral Portfolio Theory (BPT).⁴

A formal analysis of BPT is provided in section 27.3 of Shefrin (2008). It consists of finding the optimal consumption c that takes value c_1, c_2, \dots, c_n in n states of the world. The optimum must maximize⁵ $\sum_i v_i c_i$, subject to a budget constraint and an aspiration constraint. The aspiration constraint is closely related to a Value-at-Risk constraint as the aspiration is modeled by a minimum probability A of meeting a certain threshold L , $\mathbb{P}(c \geq L) \geq A$. Let us recall the example in Chapter 27 (Table 27.1) of Shefrin (2008). Consider an economy with eight possible equiprobable states (each has probability 0.125) and assume that the cost of consumption in each state is different. For an initial budget (wealth to spend) of \$1 the investor wants to maximize the expected consumption subject to a constraint on the minimum probability that her consumption meets the aspiration level of 0.9. Since the objective is the expectation and all probabilities are equal, it is clear that one unit of consumption in any state has the same value for the investor. It is then optimal to spend the budget of \$1 to obtain consumption in the cheapest state only, since this state allows the investor to achieve the highest consumption for the initial budget, and thus the highest expected payoff. In Shefrin’s example (Shefrin (2008), Table 27.1 in Chapter 27), the cheapest state costs 0.28; thus, the optimal strategy is to consume $1/0.28$ in this particular state. However, for this particular allocation of final wealth, the probability of getting at least 0.9 (aspiration level) is only 0.125 as there is no consumption in any other state. Hence, if one adds the constraint on the probability of meeting the aspiration level of 0.9, then the investor may need to spread her consumption over several states. If the minimum probability is 0.5, then she needs to consume at least 0.9 in four different states (she will automatically choose the four cheapest states), so that the aggregate probability of getting 0.9 is 0.5, and, if she still has budget left, she consumes the rest in the cheapest state. It is clear that in the SP/A model it is never optimal to allocate consumption to the most expensive state (worst state of the economy) unless the aspiration level A is strictly higher than $7/8$ (because there are eight states in the example).

As Shefrin (2008) explains, the SP/A-based BPT also has the important property that optimal portfolios contain both very risky (lottery stocks) and risk-free investments. Lottery stocks, like lottery tickets, offer many people the only hope of reaching aspirations, whereas bonds allow protection from poverty. These findings are consistent with the empirical work of Polkovnichenko (2005) and Kumar (2009). In particular, BPT is also consistent with the demand for structured products that combine downside protection (floor) with some upside potential (Shefrin (2008)). While SP/A theory explains the use of capital guaranteed products, CPPI strategies or other structured products to satisfy the needs of investors, one observes that optimal portfolios are still long with the market. They provide security in absolute levels (by controlling the probability of meeting some aspiration level), but they also perform poorly in the most expensive states of the economy (which

⁴Under some conditions Rieger (2010) shows that SP/A theory is consistent with expected utility theory.

⁵The v_i are (possibly distorted) probability weights.

typically correspond to a crisis).⁶ Similarly, like the other traditional frameworks for portfolio selection with law-invariant and increasing preferences, SP/A theory is not designed to model the desire for security in specific states of the world. Here, we provide a framework that maintains the stylized features of the SP/A theory while dealing with the aim of security in a more flexible way.

Finally, our approach is different from traditional portfolio theories in that we do not specify the value function to optimize but instead choose a target distribution of terminal wealth. This idea builds on recent work of Goldstein, Johnson and Sharpe (2008), who propose an approach based on questionnaires allowing customers to construct their desired distribution of wealth at maturity. Thus, we consider investors who target some distribution for their final wealth and who also seek to obtain some degree of security in specific states of the economy (e.g., crises). We model a crisis as an event in which the market index is lower than its Value-at-Risk at some high confidence level. This approach is consistent with the CoVaR approach proposed by Adrian and Brunnermeier (2011) to model a market under stress. The desire for security is dealt with by allowing the investor to choose the interaction of his strategy with the market under crisis regimes. For example, the investor may want his strategy to behave independently of the market when there is a crisis, but the framework that we propose can handle any desired (tail) dependence.

Our contributions can be summarized as follows. First, we construct the cheapest possible strategy that delivers a desired wealth distribution at the end of the investment horizon and, additionally, preserves a desired tail dependence with the market. This construction is technically related to finding improved Fréchet-Hoeffding bounds for copulas in the presence of additional information. Secondly, while the new strategies that we propose are suboptimal for expected utility maximizers, we point out that they no longer incur their largest losses in the worst states of the economy. They thus outperform traditional optimal strategies during a crisis and are useful for investors who seek security in economic downturns. In particular, they appear to be useful in reducing systemic risk. Our results shed new light on the role of diversification in investment portfolios.

Section 5.1 presents the financial framework. In Section 5.2 we revisit some traditional investment strategies used to diversify risk and show that they perform poorly in crises. Section 5.3 provides a theory that allows for the construction of strategies that are resilient to worst-case scenarios. Examples are given in Section 5.4. Section 5.5 concludes.

5.1 Setting

In this section we describe the financial market and clarify the intimate link between pricing and the GOP. To facilitate its understanding, this chapter is developed in the context of the multi-dimensional Black-Scholes model. However, the results can easily be extended to include more general market models (see Section 5.5). Most of the material presented here can also be found in Chapter 15 of Luenberger (1997).

5.1.1 Financial Market

The financial market is equipped with a physical measure \mathbb{P} and contains a (risk-free) bond with price process $\{B_t = B_0 e^{rt}, t \geq 0\}$, as well as risky assets S^1, S^2, \dots, S^n governed by price processes

⁶See the example described in the preceding paragraph.

$\{S_t^i, t \geq 0, i = 1, 2, \dots, n\}$. We also assume that trading can be done continuously and that the market is frictionless and arbitrage-free. Finally, all investors agree on the state-price process $\{\xi_t, t \geq 0\}$ used to value derivatives in this market. Hence, the initial price, $c(X_T)$, of a given strategy with payoff X_T maturing at the fixed horizon $T > 0$ is given by

$$c(X_T) = E[\xi_T X_T]. \quad (5.1)$$

In particular, it holds that $c(1) = E[\xi_T] = e^{-rT}$. Note that $c(X_T)$ can also be presented as the discounted expectation under the risk-neutral measure \mathbb{Q} defined through $\xi_t = e^{-rt} \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_t$.

Note that the above description is still rather general: it includes the classical Black-Scholes setting as a special case, in which case the process $\{\xi_t, t \geq 0\}$ is known unambiguously. This framework is used in the remainder of the chapter (see also Bernard, Maj, and Vanduffel (2011)). It assumes that prices S_t^i ($i = 1, 2, \dots, n$) of the risky assets evolve according to

$$\frac{dS_t^i}{S_t^i} = \mu_i dt + \sigma_i dW_t^i, \quad i = 1, 2, \dots, n, \quad (5.2)$$

where the processes $\{W_t^i, t \geq 0\}$ ($i = 1, \dots, n$) are (correlated) standard Brownian motions, with constant correlation coefficients ρ_{ij} given as

$$\forall t, s \geq 0, \quad \rho_{ij} = \text{Corr} \left(W_t^i, W_{t+s}^j \right).$$

In what follows, there is at least one asset S^i having a non-zero risk premium ($\mu_i > r$). The log-returns are also assumed to be linearly independent (no redundant assets).

5.1.2 Growth Optimal Portfolio (GOP) and Pricing

In this section we recall why the GOP can be used as the numeraire portfolio for pricing. To this end, we first consider portfolios in which the initially invested fractions π_i in the risky assets S^i are kept constant during the entire investment horizon (through rebalancing). Let $1 - \sum_{i=1}^n \pi_i$ be the proportion that is invested in the risk-free asset. The value process $\{S_t^\pi, t \geq 0\}$ of such a constant-mix portfolio with a vector of initial proportions $\pi := (\pi_1, \pi_2, \dots, \pi_n) \neq \vec{0}$ follows the equation

$$\frac{dS_t^\pi}{S_t^\pi} = \mu_\pi dt + \sigma_\pi dW_t^\pi, \quad (5.3)$$

where $\{W_t^\pi, t \geq 0\}$ is the standard Brownian motion defined through $W_t^\pi = \frac{\sum_{i=1}^n \pi_i \sigma_i W_t^i}{\sqrt{\pi^T \cdot \Sigma \cdot \pi}}$. Here, Σ is the (positive definite) variance-covariance matrix with elements $(\Sigma)_{ij} := \sigma_{ij} = \rho_{ij} \sigma_i \sigma_j$, and

$$\mu_\pi = r + \pi^T \cdot (\mu - r \mathbf{1}) \text{ and } \sigma_\pi^2 = \pi^T \cdot \Sigma \cdot \pi, \quad (5.4)$$

and $\mathbf{1}$ denotes a vector of size n of ones. It is well-known that the solution to equation (5.3) is

$$S_t^\pi = S_0^\pi \exp(X_t^\pi), \quad (5.5)$$

with $X_t^\pi = (\mu_\pi - \frac{1}{2}\sigma_\pi^2) t + \sigma_\pi W_t^\pi$. Next, we consider in particular the constant-mix portfolio that maximizes the expected growth rate $\mu_\pi - \frac{1}{2}\sigma_\pi^2$. It is easily shown that this portfolio, denoted by π^* , is given by

$$\pi^* = \Sigma^{-1} \cdot (\mu - r \mathbf{1}). \quad (5.6)$$

This constant-mix strategy is also the optimal portfolio for an expected log-utility maximizer (Luenberger (1997)). In the literature, π^* (and S_{π^*}) is typically referred to as the growth-optimal portfolio (GOP). For convenience, we denote S_{π^*} by S^* ; we also denote μ_{π^*} , σ_{π^*} and $W_T^{\pi^*}$ by μ_* , σ_* and W_T^* , respectively.

The following proposition shows that there is an intimate connection between the GOP and the state-price process $\{\xi_t, t \geq 0\}$ for the multidimensional Black-Scholes market. The original proof can be found in Chapter 15 (“Optimal Portfolio Growth”) of Luenberger (1997).

Proposition 5.1.1 (State-price Process). *In the multi-dimensional Black-Scholes market, the state-price process $\{\xi_t, t \geq 0\}$ is given by*

$$\xi_t = \frac{S_0^*}{S_t^*}. \quad (5.7)$$

Proposition 5.1.1 establishes a bijection between the GOP and the state-price process. In particular, the pricing formula (5.1) actually states that the price of a payoff X_T paid at T equals the arithmetic average of all the possible outcomes expressed in units of the GOP. In other words, the GOP appears here as a numeraire portfolio that is useful for pricing.

5.1.3 Market Crises

The purpose of this chapter is to investigate optimal portfolio selection when the agent is particularly concerned about his final wealth during a crisis. As the GOP S^* can also be interpreted as a major market index (see also Platen and Heath (2006)), it appears intuitive to define a stressed market (or crisis) at time T as an event in which *the market* - materialized through S^* - drops below its Value-at-Risk at some high confidence level. The corresponding states of the economy satisfy

$$\{S_T^* < q_\alpha\}, \quad (5.8)$$

where q_α is such that $\mathbb{P}(S_T^* < q_\alpha) = 1 - \alpha$ and α is typically high (e.g., $\alpha = 0.98$). This approach is consistent with the CoVaR approach proposed by Adrian and Brunnermeier (2011) to model a market under stress. Note that the event (5.8) also corresponds to the states of the economy for which

$$\{\xi_T > c\}, \quad (5.9)$$

with c such that $\mathbb{P}(\xi_T > c) = 1 - \alpha$. The states of the crisis are thus the highest values of ξ_T that correspond to the most expensive Arrow-Debreu states.

5.2 Traditional Diversification Strategies

In this section, we discuss some traditional diversification strategies and show how they fall short under adverse events.

5.2.1 Buy-and-Hold Strategies

The buy-and-hold strategy is the simplest investment strategy. The initial amount is $V_0 = w_0 + \sum_{i=1}^n w_i$, where w_0 is invested in the bank account and w_i is invested in stock S^i ($i = 1, 2, \dots, n$) at

time 0 and no further action is undertaken. The absence of a dynamic rebalancing implies that the proportion of (historically) good performing assets will increase in the asset mix. The final wealth at time $T > 0$ is obtained by

$$V_T = w_0 e^{rT} + \sum_{i=1}^n w_i \frac{S_T^i}{S_0^i}.$$

Note that V_T writes as the sum of a fixed floor value provided by the risk-free asset augmented with the random values for the different assets.

5.2.2 Constant-Mix Strategies

In contrast to the static buy-and-hold strategy, constant-mix strategies need dynamic rebalancing to preserve the initial target allocation. This rebalancing is problematic in the sense that the preservation of the initial proportions requires better-performing assets to be sold in exchange for the purchase of worse-performing ones. For an initial investment of V_0 , it follows that at maturity one holds the following final wealth, V_T :

$$V_T = V_0 \frac{S_T^\pi}{S_0^\pi},$$

where π is the vector of proportions and S_T^π is as in (5.5). It is well-known that constant-mix portfolios of the form $\pi = \alpha \pi^*$, with $\alpha > 0$ and π^* being the optimal vector of proportions for the GOP (see (5.6)), are optimal strategies for CRRA expected utility maximizers. Specifically, CRRA investors with a constant relative risk aversion coefficient $\eta > 0$ have utility

$$U(x) = \begin{cases} \frac{x^{1-\eta}}{1-\eta} & \text{when } \eta \neq 1 \\ \log(x) & \text{when } \eta = 1 \end{cases},$$

and the optimal wealth satisfying an initial budget V_0 is then calculated as $[U']^{-1}(k\xi_T)$, where k is chosen to meet the budget constraint. After some straightforward calculations,

$$V_T^* = V_0 (e^{rT})^{1-\frac{1}{\eta}} \left(\frac{S_T^*}{S_0} \right)^{\frac{1}{\eta}}. \quad (5.10)$$

Hence, the terminal wealth V_T^* is the result of a constant-mix strategy, in which one invests a $\frac{1}{\eta}$ proportion of the initial wealth V_0 in the GOP and the remaining $1 - \frac{1}{\eta}$ proportion in a bank account.

5.2.3 Performance of Investment Strategies during a Crisis

Bernard, Boyle, and Vanduffel (2014) characterize optimal strategies for investors with increasing law-invariant preferences (i.e., investors who care only about the final distribution and who prefer more to less). In this context it is clear that an optimal strategy must be the cheapest possible one that provides some distribution F for terminal wealth (also called a cost-efficient strategy). Using Fréchet-Hoeffding bounds, it can be shown that strategies are cost-efficient if and only if they are almost surely non-decreasing in the GOP S_T^* . In fact, a strategy V_T with distribution F is cost-efficient if and only if, almost surely,

$$V_T = F^{-1}(F_{S_T^*}(S_T^*)).$$

This monotonicity result has also been proved by Beare (2011) and by Rieger (2011), who formulates the optimal portfolio choice problem as a Monge-Kantorovich optimal transportation problem. Cost-efficient strategies have been extensively studied in Bernard, Boyle, and Vanduffel (2014); See also Bernard, Maj, and Vanduffel (2011) and Vanduffel et al. (2009; 2012) for more examples and some related results. As the optimal strategies are non-decreasing in the GOP, they provide their worst outcomes in downturns. In other words, while the comonotonicity⁷ between V_T and S_T^* is key to optimality for agents with law-invariant preferences, it comes at the cost of having little protection against unfavorable market conditions. To clarify this point, let us consider two payoffs V_T^1 and V_T^2 with the same distribution F . Furthermore, V_T^1 is assumed to be comonotonic with S_T^* , while V_T^2 is not. It is straightforward to show that $(V_T^1 | S_T^* < q_\alpha)$ is smaller in first-order stochastic dominance than $(V_T^2 | S_T^* < q_\alpha)$. This explains why investors who also seek value when S_T^* is low tend to prefer V_T^2 over V_T^1 .

All constant-mix strategies give rise to a lognormally distributed terminal wealth V_T , but not all of them are optimal. The requirement of comonotonicity between V_T and the GOP, S_T^* (which itself can be seen as the terminal wealth of a constant-mix strategy), implies that optimality can only be achieved through rebalancing the GOP with the bank account.

Nevertheless, constant-mix strategies and buy-and-hold strategies are typically strongly correlated with the GOP and thus suffer from their positive dependence with market movements during crises. This shortcoming appears clearly in the numerical illustration in Figure 5.4.1 page 87 (Strategies 1, 2 and 3). Hence, there is a real need to propose new strategies that are as cost-efficient as possible but that do not suffer from the same drawbacks. New strategies of this kind are developed in Section 5.3. In Section 5.4, we illustrate our findings with numerical examples and show how the traditional investment strategies presented in this section (i.e., the buy-and-hold and constant-mix portfolio strategies) are outperformed by the newly proposed “optimal” strategies, which provide protection during market crises.

5.3 Optimal Tail Diversification

We depart here from the traditional setting and study the design of optimal strategies for investors whose preferences encompass both a certain distribution of terminal wealth and a desire for extra security in stressed market environments. Investors may obtain “security” by meeting a minimum (global) aspiration level as in SP/A theory; however, this may not be sufficient to satisfy their security needs in certain states of the economy. The existence of Armageddon funds, which yield advantages only in times of scarcity, is a clear signal that some investors are interested in receiving protection against worst-case scenarios. As another indication that buying protection against worst-case scenarios is not unreasonable, we mention that many people buy property insurance despite the existence of cheaper financial contracts that provide the same distribution (Chapter 3, Bernard and Vanduffel (2014a)). The basic intuition is that insurance provides protection when it is needed (when the property burns down) whereas the financial contract does not and investors are prepared to pay an extra premium for this insurance feature. This section focuses on the design of optimal strategies meeting the preferences of investors for a certain distribution of terminal wealth⁸ and

⁷For a comprehensive account on the topic of comonotonicity in the context of decision making we refer to Dhaene et al. (2006).

⁸Under some conditions, it can be shown that imposing a given cdf for the optimal terminal wealth is equivalent to stating the utility function of an investor who maximizes expected utility (Chapter 2, Bernard, Chen, and Vanduffel (2015)).

additionally, the extra desire for security in stressed market environments.

5.3.1 Optimal Strategies with Constraints in the Tail

The next theorem constitutes the key result in this chapter. As it is expressed with full generality, it is rather technical. All of the subsequent results apply this main theorem in different contexts. Theorem 5.3.1 provides the cheapest (and thus optimal) investment strategy for an investor who wants to achieve a fixed distribution F for his final wealth at investment horizon T as well as a desired interaction with the market under a crisis regime. Cheapest strategies that exhibit a fixed distribution and satisfy one or more additional constraints are called “constrained cost-efficient” strategies (Chapter 1, Bernard, Boyle, and Vanduffel (2014)).

Proposition 5.3.1 (Optimal strategies under a crisis regime). *Let F be the cumulative distribution function (cdf) that the agent wants to achieve for his terminal wealth V_T at investment horizon T (i.e., $\forall x, \mathbb{P}(V_T \leq x) = F(x)$). We further assume that the investor chooses the joint distribution G of his final wealth V_T with the financial market S_T^* when there is a financial crisis (that is, $S_T^* \leq q_\alpha$):*

$$\mathbb{P}(S_T^* \leq s, V_T \leq x \mid S_T^* \leq q_\alpha) = G(s, x). \quad (5.11)$$

Then, an optimal wealth V_T^* is given by

$$V_T^* = \begin{cases} F^{-1}\left(h(F_{S_T^*}(S_T^*) - \alpha)\right) & \text{when } S_T^* > q_\alpha \\ F^{-1}\left(g\left(1 - F_{S_T^*}(S_T^*), j_{F_{S_T^*}}(S_T^*)(F_{Z_T}(Z_T))\right)\right) & \text{when } S_T^* \leq q_\alpha \end{cases}, \quad (5.12)$$

where Z_T is any random variable such that (S_T^*, Z_T) is continuously distributed and where $j_u(v), g(u, v)$ and $h(x)$ are defined as follows: Since the random pair (S_T^*, Z_T) is continuously distributed, there is a unique (copula) function $J(\cdot, \cdot)$ such that $\mathbb{P}(S_T^* \leq s, Z_T \leq x) = J(F_{S_T^*}(s), F_{Z_T}(x))$. We then denote by $j_u(v)$ its first partial derivative, i.e., $j_u(v) = \frac{\partial}{\partial u} J(u, v)$. Similarly, we can recast equation (5.11) as

$$\forall s \in [0, q_\alpha], x \in \mathbb{R}, \mathbb{P}(S_T^* \leq s, V_T \leq x) = C^*(F_{S_T^*}(s), F(x)) \quad (5.13)$$

for the appropriate C^* determined through G . The functions $h(\cdot)$ and $g(\cdot)$ are then defined as $h(x) = \inf\{c \mid c - C^*(\alpha, c) \geq x\}$ and $g(u, v) = c_u^{-1}(v)$, where $c_u(v) = \frac{\partial C^*}{\partial u}(1 - u, v)$.

The strength of Theorem 5.3.1 is the explicit expression (5.12) for the optimal strategy under the general (tail) dependence structure assumed in (5.11). The optimal terminal wealth V_T^* , given by (5.12), can be interpreted as the cheapest strategy with cdf F that verifies the constraint on the dependency in the tail (5.11).

Note that unlike cost-efficient strategies, there is no general characterization for constrained cost-efficient strategies, and even their existence is not guaranteed. Nevertheless, in the setting described in Theorem 5.3.1 the optimum exists, and we are able to provide an explicit construction (5.12). This requires some careful work, such as proving the existence of a lower bound on a set of copulas.

Furthermore, there is no uniqueness to the optimal strategy (5.12) in general. It depends largely on the particular choice of the market variable Z_T used in its construction, and the only requirement is that (S_T^*, Z_T) is continuously distributed. Hence, we can identify Z_T with another asset in the financial market, or we can take it as the value of the GOP at an earlier time $t < T$.

In Section 5.3.2 we provide explicit expressions for optimal strategies that exhibit independence from the GOP during a crisis. Other types of dependence are available in the published version of the chapter. In Section 5.4 we provide numerical experiments and comparisons between traditional diversification strategies and optimal strategies constructed as in Theorem 5.3.1.

5.3.2 Independence in the Tail

In the published version (Bernard, Chen and Vanduffel 2014), two theorems are established in the most general market setting. Here we only evaluate the properties of the new strategies in the specific context of a two-dimensional Black-Scholes model. In this setting, the expressions for the optimal strategies can be further simplified and are given here in two corollaries. There are two assets, S^1 and S^2 , such that

$$\begin{cases} \frac{dS_t^1}{S_t^1} = \mu_1 dt + \sigma_1 dW_t^1, \\ \frac{dS_t^2}{S_t^2} = \mu_2 dt + \sigma_2 dW_t^2, \end{cases} \quad (5.14)$$

where W^1 and W are two correlated Brownian motions under the physical probability measure \mathbb{P} . W can be decomposed into two independent standard \mathbb{P} -Brownian motions through $dW_t = \rho_{12}dW_t^1 + \sqrt{1 - \rho_{12}^2}dW_t^2$. We further assume that $\mu_1 > r$ and $\mu_2 > r$.

Corollary 5.3.2 (Tail Diversification Using a Path-Dependent Strategy when $n = 2$). *In a two-dimensional Black-Scholes market, the cheapest path-dependent strategy with a cumulative distribution F but that is independent of S_T^* when $S_T^* \leq q_\alpha$ can be constructed as*

$$V_T^* = \begin{cases} F^{-1} \left(\frac{F_{S_T^*}(S_T^*) - \alpha}{1 - \alpha} \right) & \text{when } S_T^* > q_\alpha, \\ F^{-1}(\Phi(A)) & \text{when } S_T^* \leq q_\alpha, \end{cases} \quad (5.15)$$

where

$$A = \frac{\frac{1}{\sigma_1} \left[\ln \left(\frac{S_t^1}{(S_t^1)^{t/T}} \right) - \left(1 - \frac{t}{T}\right) \ln(S_0^1) \right] + \frac{1}{\sigma_2} \left[\ln \left(\frac{S_t^2}{(S_t^2)^{t/T}} \right) - \left(1 - \frac{t}{T}\right) \ln(S_0^2) \right]}{\sqrt{2(1 + \rho_{12}) \left(t - \frac{t^2}{T} \right)}}.$$

Corollary 5.3.3 (Tail Diversification Using a Path-Independent Strategy when $n = 2$). *In a two-dimensional Black-Scholes market, the cheapest path-independent strategy with a cumulative distribution F but that is independent of S_T^* when $S_T^* \leq q_\alpha$ can be constructed as*

$$V_T^* = \begin{cases} F^{-1} \left(\frac{F_{S_T^*}(S_T^*) - \alpha}{1 - \alpha} \right) & \text{when } S_T^* > q_\alpha, \\ F^{-1}(\Phi(A)) & \text{when } S_T^* \leq q_\alpha, \end{cases} \quad (5.16)$$

where

$$A = \frac{\frac{1}{\sigma_1} \left[\ln \left(\frac{S_T^1}{S_0^1} \right) - \left(\mu_1 - \frac{\sigma_1^2}{2} \right) T \right] - \frac{1}{\sigma_2} \left[\ln \left(\frac{S_T^2}{S_0^2} \right) - \left(\mu_2 - \frac{\sigma_2^2}{2} \right) T \right]}{\sqrt{2(1 - \rho_{12}) T}}.$$

So far Theorem 5.3.1 has been applied to a situation in which the investor wants to achieve independence with the GOP during a crisis situation. This corresponds to the case where the copula C^* in Theorem 5.3.1 is simply given as $C^*(u, v) = uv$. The published version of this chapter

contains more general dependence structures. For example, it is possible to apply our theory to any type of dependence. We study in particular the case of a Gaussian dependence. It is particularly interesting when one seeks to achieve a negative correlation with the financial market when there is a crisis. Such a strategy is then able to offer some protection against worst-case scenarios. The disadvantage of the Gaussian dependence is that there are no closed-form expressions. Other dependence structures such as the Clayton or the Frank copula can be solved explicitly. Formulas can be found in Bernard, Chen, and Vanduffel (2014).

5.4 Applications

We present some numerical results in a two-dimensional Black-Scholes market. We assume that the risk-free rate is $r = 0.05$, and the instantaneous rate of return and volatility under the physical measure for the two stocks are given respectively by $\mu_1 = 0.07$, $\sigma_1 = 0.20$ and $\mu_2 = 0.08$, $\sigma_2 = 0.30$. Finally, the correlation ρ_{12} is set to be equal to 0.25.

We consider six strategies – Strategies 1, 2 and 3 are standard diversified strategies, whereas each of the remaining strategies corresponds to a corollary or a proposition in the chapter.

Strategy 1: Investment in the GOP. Strategy 1 invests the entire initial wealth of $V_0^1 = 100$ in the Growth Optimal Portfolio. Hence, we implement a constant-mix strategy with weights π_i^* ($i = 1, 2$) given in (5.6). The final payoff V_T^1 , paid at time T , writes as $V_T^1 = V_0^1 \frac{S_T^*}{S_0^*}$.

Strategy 2: Naive Buy-and-Hold Strategy. Strategy 2 is a naive buy-and-hold strategy in which the initial wealth of $V_0^2 = 100$ is split into one third in each risky asset and one third in the bank account. The final payoff, paid at time T , is equal to

$$V_T^2 = \frac{100}{3} \frac{S_T^1}{S_0^1} + \frac{100}{3} \frac{S_T^2}{S_0^2} + \frac{100}{3} e^{rT}.$$

Strategy 3: Naive Constant-Mix Strategy. Strategy 3 consists of investing $V_0^3 = 100$ according to a constant-mix strategy with proportions $\pi_1 = \pi_2 = 1/3$ in the risky assets and $(1 - \pi_1 - \pi_2)$ in the risk-free asset.

Strategy 4: Independent Tail (Corollary 5.3.2). Strategy 4 is a constrained strategy with the same distribution as the investment in the GOP at maturity (Strategy 1), but that is independent of the GOP when the GOP drops below q_α , where $\alpha = 5\%$. We assume that it is path-dependent and construct it as in Corollary 5.3.2.

Strategy 5: Independent Tail (Corollary 5.3.3). Strategy 5 is also a constrained strategy with the same distribution as the investment in the GOP at maturity, but it is independent of the GOP when the GOP drops below q_α , where $\alpha = 5\%$. We assume that it is path-independent and construct it as in Corollary 5.3.3.

Strategy 6: Gaussian Tail with Negative Dependence Strategy 6 is also a constrained strategy with the same distribution as the investment in the GOP at maturity, but the dependence with the GOP is Gaussian with correlation coefficient $\rho = -0.5$ during a crisis.

Figure 5.4.1 represents the outcomes for the different strategies as a function of the GOP when $T = 1$. This figure clearly shows the major difference between traditional diversification strategies (graphs 1, 2 and 3), which exhibit a strong positive correlation with the GOP, and the newly proposed “constrained cost-efficient strategies” (graphs 4, 5 and 6), which do not depict this feature

in a crisis regime (i.e., for low values of the GOP). We observe in the last three graphs that the new strategies depend not only on the GOP, but also on another source of uncertainty (either the GOP at an earlier date, or another asset in the market). Indeed, when the GOP is below its Value-at-Risk at 95% confidence, the graphs no longer result in curves, but instead generate clouds of points. This observation reflects some degree of diversification when the market is down. Note also that, by construction, Strategies 1, 4, 5 and 6 all lead to the same final distribution of wealth. In a crisis regime, the average outcome for Strategy 6 is higher than the average outcome obtained using Strategies 4 and 5. This is because we required the former strategies to display a negative dependence with the GOP during a stressed period, whereas independence was assumed for the latter strategies.

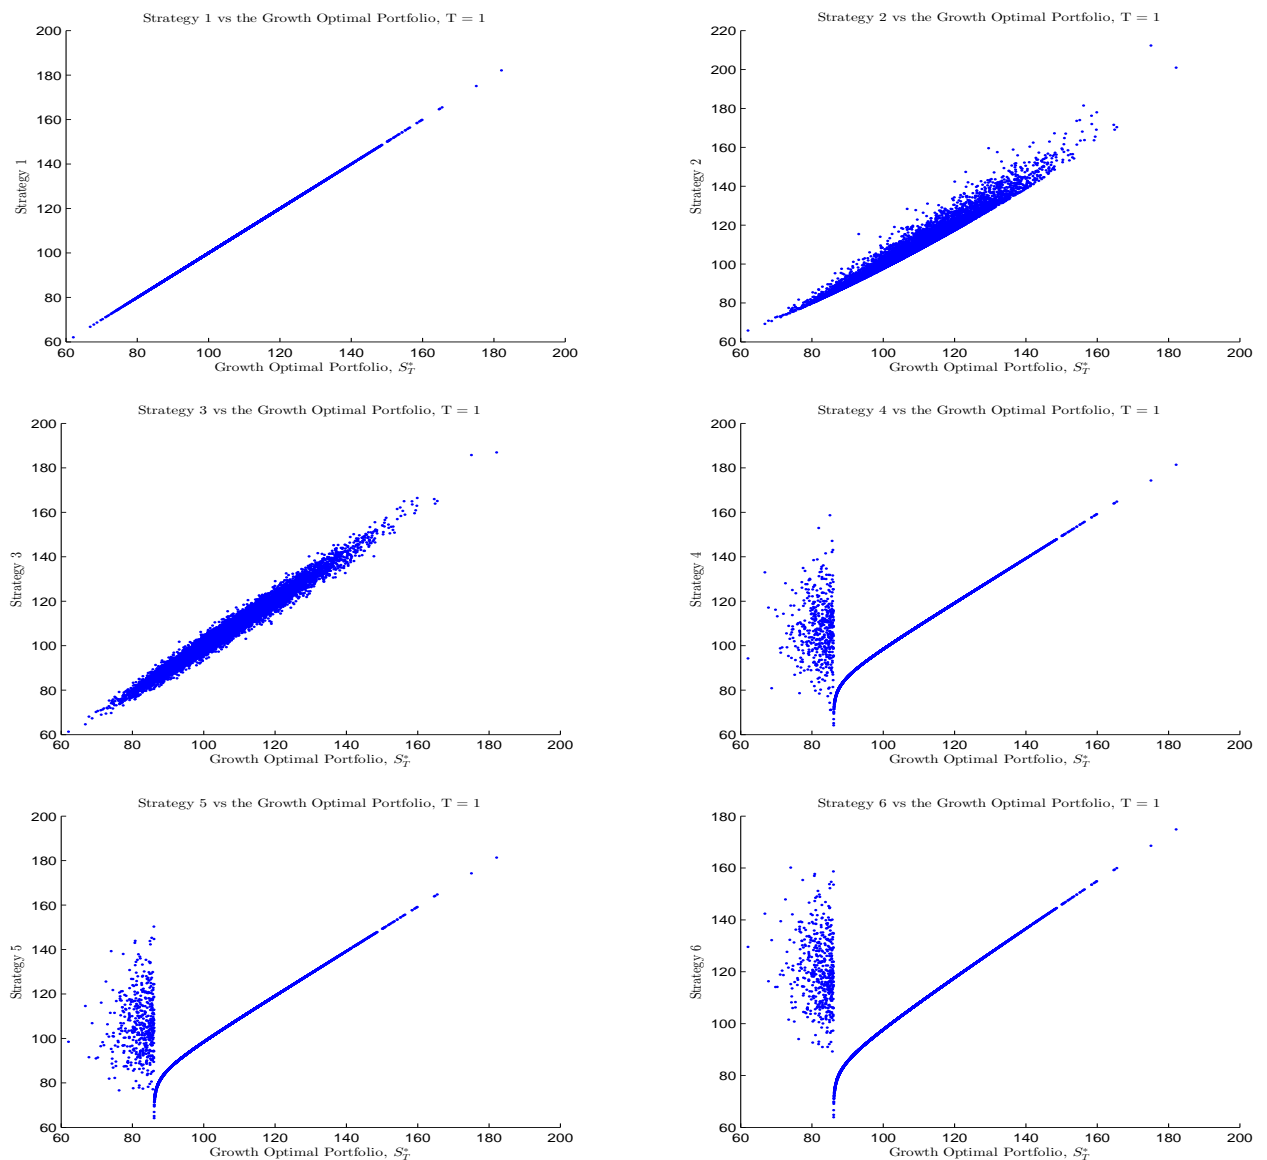


Figure 5.4.1: Plot of the Payoffs of Strategies 1–6 compared to that of the GOP.

For each strategy, we further investigate several measures reflecting performance and risk. We

consider *cost*, *Sharpe ratio* and various relevant *conditional probabilities*. The numerical results in Table 5.4.1 are obtained using the set of parameters described above.

The column headed “Cost” reflects the cost of each strategy. It is equal to 100 for strategies 1, 2 and 3 given that the initial investment V_0^i is exactly 100 by construction for $i = 1, 2$ and 3. Strategies 4 to 6 have the same distribution of final wealth as the first strategy, but they are strictly more expensive. This is because Strategy 1 is non-decreasing in the GOP, which implies that it is the cheapest possible way to obtain this distribution (cost-efficient strategy). In contrast, due to the presence of a constraint on the tail dependence, Strategies 4 to 6 are *not* non-decreasing in the GOP and thus are more expensive than Strategy 1.

Note also that Strategy 1 is the only cost-efficient strategy among the eight strategies. Although Strategies 2 and 3 have the same cost as Strategy 1, both are suboptimal in the sense that it is possible to construct a non-decreasing strategy in the GOP that has the same distribution as Strategy 2 (resp. Strategy 3) but that is strictly cheaper (Bernard, Boyle, and Vanduffel (2014)).

The Sharpe ratio of a strategy has always been a standard measure of performance. It is defined as

$$\frac{E[V_T] - V_0 e^{rT}}{\text{std}(V_T)}.$$

We observe that there is a moderate (and, for long term horizon investors, even a rather strong) decrease in the Sharpe ratios of Strategies 4 to 6 compared to that of Strategy 1 due to the efficiency loss for the latter strategies. For instance, in the case of the optimal Strategy 4 constructed from Corollary 5.3.2, we have

$$V_T^4 \sim V_T^1, \quad (5.17)$$

but the cost of V_T^4 is strictly higher than the cost of V_T^1 . Hence, $E[V_T^i] = E[V_T^1]$ and $\text{Var}(V_T^i) = \text{Var}(V_T^1)$, but $V_0^i > V_0^1 = 100$ ($i = 4, 5, 6$). For Strategies 4 to 6, the loss in Sharpe ratio associated with the additional constraint on the dependence in the tail can be thought of as the price of security. The “stronger” the need for security, the higher the loss in the Sharpe ratio.

In our newly proposed setting, there is a clear trade-off between the gain in protection (or performance) in stressed states of the market and the loss in Sharpe ratio (and, more generally, the loss in any law-invariant objective). The loss in Sharpe ratio is less pronounced for the naive buy-and-hold and constant-mix Strategies (2 and 3) because they show strong correlations with the first strategy, which corresponds to investing in the GOP and does not provide protection in stressed markets.

The interpretation of changes in the Sharpe ratios in Table 5.4.1 shows that constraining the dependence of the strategies during a crisis comes at a price. Indeed, we note that the measures of risk and performance deteriorate slightly when each of the newly proposed strategies is implemented. However, each of these measures is obtained through a stand-alone evaluation of the strategy, an approach that completely ignores the important issue of protection during a crisis. Effectively, for the new strategies, the largest losses no longer occur in the worst states in the economy. Therefore, a fair comparison of the eight strategies also needs to incorporate an analysis of what happens in the worst states of the financial market.

To this end, we have performed an analysis of conditional probabilities, and these are reported in the last three columns of Table 5.4.1. It is clear that the four types of tail dependencies significantly reduce the probability of underperforming a risk-free investment (event A) or losing 25% of initial wealth (event B) when a financial crisis occurs (event C). The probability of losing 25% of initial wealth is 1 when investing in the GOP, but it is reduced to less than 8% for strategy 6, all of

	T	Cost	Sharpe	$\mathbb{P}(A C)$	$\mathbb{P}(A D)$	$\mathbb{P}(B C)$
1	1	100	0.125	1.00	1.00	0.19
	5	100	0.266	1.00	1.00	1.00
	20	100	0.444	1.00	1.00	1.00
2	1	100	0.121	1.00	0.96	0.16
	5	100	0.239	0.9998	0.965	0.99
	20	100	0.275	1.00	0.96	0.996
3	1	100	0.123	1.00	0.95	0.27
	5	100	0.261	1.00	0.95	1.00
	20	100	0.432	1.00	0.95	1.00
4	1	100.28	0.104	0.48	0.95	0.01
	5	101.67	0.214	0.46	0.94	0.13
	20	108.86	0.340	0.45	0.93	0.27
5	1	100.30	0.103	0.48	0.94	0.01
	5	101.63	0.215	0.47	0.94	0.14
	20	108.88	0.341	0.44	0.93	0.25
6	1	100.56	0.083	0.12	0.91	0.000
	5	103.40	0.159	0.12	0.90	0.01
	20	121.06	0.199	0.13	0.89	0.05

Table 5.4.1: Numerical Results: We define two events related to *the market portfolio* S^* , i.e., the market crisis $C = \{S_T^* < q_\alpha\}$ and a decrease in the market $D = \{S_T^* < S_0^* e^{rT}\}$. We further define two events for the portfolio value by $A = \{V_T < V_0 e^{rT}\}$ and $B = \{V_T < 75\% V_0 e^{rT}\}$

which exhibit negative dependence during a crisis. Strategies 4 and 5, which are unrelated with the financial market during a crisis, perform better in terms of Sharpe Ratio than the strategies that exhibit negative dependence (because the latter strategies are more expensive). Their probability of losing more than 25% of initial wealth is less than 1% for a one year horizon and less than 27% over a 20-year period. Finally, we emphasize that Strategies 2 and 3 do suffer from the same drawbacks as Strategy 1 in the sense that they offer little protection during a crisis and other bad events. While their Sharpe-ratios are near optimal, their reported conditional probabilities indicate that they are not resilient against crises.

5.5 Conclusions

In this chapter we studied optimal portfolios using the Black-Scholes setting as a reference framework. When preferences are law-invariant, optimal strategies are necessarily non-decreasing in the Growth Optimal Portfolio (GOP). Since the GOP can be generically identified with “*the market*”, this leads to the conclusion that the worst outcomes for optimal strategies occur in bear markets (e.g., during a financial crisis). Arguably, this is at odds with the aspirations of most investors, who may desire some resilience for their wealth against crisis situations.

Hence, in this chapter, we depart from the traditional setting and study optimal strategies for investors who are only concerned about final distribution of wealth but who also impose constraints on its interaction with stressed financial markets. We model a stressed situation as an event in which the GOP is lower than its Value-at-Risk at some high confidence level. This approach is consistent

with the CoVaR approach proposed by Adrian and Brunnermeier (2011) to model a market under stress. Furthermore, the idea of assessing risk and performance of a portfolio not only by looking at its final distribution but also by looking at its interaction with economic conditions is related to the increasing concern with evaluating systemic risk. Acharya (2009) recommends “regulating each bank as a function of both its joint (correlated) risk with other banks [and] its individual (bank-specific) risk”. Capital requirements based on the risk of a company considered in isolation are not a satisfactory measure of risk.

This chapter also extends the SP/A Theory and Behavioral Portfolio Theory in the sense that we allow the desire for security to become state-dependent. More precisely, we construct the cheapest possible strategy that delivers a given wealth distribution at a given horizon and, additionally, preserves a desired dependence with the market during a financial crisis. Hence, “security” is preserved by specifying a tail dependence, whereas “potential” is maximized by finding the cheapest possible payoff. The

The results can be generalized to more involved market settings as follows. Using a milder notion of arbitrage, Platen and Heath (2006) argue that, in general, the price of (non-negative) payoffs could be achieved using the pricing rule (5.1), where the role of the stochastic discount factor is played by the inverse of the GOP.⁹ Hence, in this generalized setting, all of our results still hold in an obvious way. Under the standard notion of arbitrage, the results also hold true if we identify the inverse of the stochastic discount factor with “*the financial market*” (so that the definition of a crisis given by (5.9) is still valid). This definition is very intuitive in the sense that the states of a financial crisis are probably the most expensive states to insure, and therefore this corresponds to the states in which the highest values for the state-price process are observed. Moreover, in addition to the Black-Scholes setting, such identification can also be made in Lévy markets, in which the participants use Esscher pricing.

⁹Note that the alternative representation, in which the price is reflected as a risk-neutral expectation, no longer necessarily holds. This is because a risk-neutral measure does not always exist in this setting.

Chapter 6

Optimal Payoffs under State-dependent Preferences

This last chapter makes use of cost-efficiency but goes beyond the original setting in which investors only care about the distribution of wealth (law invariant preferences). In this chapter, we propose a setting to incorporate state-dependent preferences.

Most decision theories, including expected utility theory, rank dependent utility theory and cumulative prospect theory, assume that investors are only interested in the distribution of returns and not in the states of the economy in which income is received. Optimal payoffs have their lowest outcomes when the economy is in a downturn, and this feature is often at odds with the needs of many investors. We introduce a framework for portfolio selection within which state-dependent preferences can be accommodated. Specifically, we assume that investors care about the distribution of final wealth *and* its interaction with some benchmark. In this context, we are able to characterize optimal payoffs in explicit form. Furthermore, we extend the classical expected utility optimization problem of Merton to the state-dependent situation. Some applications in security design are discussed in detail and we also solve some stochastic extensions of the target probability optimization problem. Proofs and details can be found in the corresponding published version of this chapter in *Quantitative Finance* (Bernard, Moraux, Rüschendorf, and Vanduffel (2015)).

Studies of optimal investment strategies are usually based on the optimization of an expected utility, a target probability or some other (increasing) *law-invariant* measure. Assuming that investors have law-invariant preferences is equivalent to supposing that they care only about the distribution of returns and not about the states of the economy in which the returns are received. This is, for example, the case under expected utility theory, Yaari's dual theory, rank-dependent utility theory, mean-variance optimization and cumulative prospect theory. Clearly, an optimal strategy has some distribution of terminal wealth and must be the cheapest possible strategy that attains this distribution. Otherwise, it is possible to strictly improve the objective and to contradict its optimality. Dybvig (1988a) was the first to study strategies that reach a given return distribution at lowest possible cost. Bernard, Boyle, and Vanduffel (2014) (Chapter 1) call these strategies cost-efficient and study their properties. In a fairly general market setting these authors show that the cheapest way to generate a given distribution is obtained by a contract whose payoff is decreasing in the pricing kernel (see also Carlier and Dana (2011)). The basic intuition is that investors consume less in states of economic recession because it is more expensive to insure returns under these conditions. This feature is also explicit in a Black-Scholes framework, in which optimal

payoffs at time horizon T are shown to be an *increasing* function of the price of the risky asset (as a representation of the economy) at time T . In particular, such payoffs are path-independent.

An important issue with respect to the optimization criteria and the resulting payoffs under most standard frameworks, is that their worst outcomes are obtained when the market declines. Arguably, this property of optimal payoffs does not fit with the aspirations of investors, who may seek protection against declining markets or, more generally, may consider a benchmark when making investment decisions. In other words, two payoffs with the same distribution do not necessarily present the same “value” for a given investor. Bernard and Vanduffel (2014a) show that insurance contracts can usually be substituted by financial contracts that have the same payoff distribution but are cheaper. The existence of insurance contracts that provide protection against specific events shows that these instruments must present more value for an investor than financial payoffs that lack this feature. This observation supports the general observation that investors are more inclined to receive income in a “crisis” (for example when their property burns down or when the economy is in recession) than under “normal” conditions.

This chapter makes several theoretical contributions to the study of optimal investment strategies and highlights valuable applications of its findings in the areas of portfolio management and security design. First, we clarify the setting under which optimal investment strategies necessarily exhibit path-independence. These findings complement Cox and Leland (1982; 2000) and Dybvig’s (1988a) seminal results and underscore the important role of path-independence in traditional optimal portfolio selection. Thereafter, as our main contribution, we introduce a framework for portfolio selection that makes it possible to consider the states in which income is received. More precisely, it is assumed that investors target some distribution for their terminal wealth and *additionally* aim for a certain (desired) interaction with a random benchmark.¹ For example, the investor may want his strategy to be unrelated to the benchmark when it decreases but to follow this benchmark when it performs well. Using our framework, we can characterize optimal payoffs explicitly (Theorems 6.3.1 and 6.3.3) in this setting. Furthermore, we show that optimal strategies in this setting become *conditionally increasing* functions of the terminal value of the underlying risky asset.

A further main contribution in the published version of the chapter is the extension of the classical result of portfolio optimization under expected utility (Cox and Huang (1989)). Specifically, we determine the optimal payoff for an expected utility maximizer under a dependence constraint, reflecting a desired interaction with the benchmark. The proof builds on isotonic approximations and their properties (Barlow *et al.* (1972)). We also solve two stochastic generalizations of Browne (1999) and Spivak and Cvitanić (1999) classical target optimization problem in the given state-dependent context.

Finally, we show how these theoretical results are useful in security design and can help to simplify (and improve) payoffs commonly offered in the financial markets. We show how to substitute highly path-dependent products by payoffs that depend only on two underlying assets, which we refer to as “twins”. This result is illustrated with an extensive discussion of the optimality of Asian options. We also construct alternative payoffs with appealing properties.

The last chapter is organized as follows. Section 6.1 outlines the setting of the investment problem under study. In Section 6.2, we restate basic optimality results for path-independent

¹The chapter draws its inspiration from the last section of Chapter 1 (Bernard, Boyle and Vanduffel (2014)), in which a constrained cost-efficiency problem is solved when the joint distribution between the wealth and some benchmark is determined in some specific area (local dependence constraint).

payoffs for investors with law-invariant preferences. We also discuss in detail the sufficiency of path-independent payoffs when allocating wealth. In Section 6.3, we point out drawbacks of optimal path-independent payoffs and introduce the concept of state-dependence used in the following sections. We show that “twins”, defined as payoffs that depend only on two underlying asset values, are optimal for state-dependent preferences. In Section 6.4, we discuss applications to improve security designs. In particular, we propose several improvements in the design of geometric Asian options. Final remarks are presented in Section 6.5.

6.1 Framework and Notation

Consider investors with a given finite investment horizon T and no intermediate consumption. We model the financial market on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, in which \mathbb{P} is the real-world probability measure. The market consists of a bank account B paying a constant risk-free rate $r > 0$, so that B_0 invested in a bank account at time 0 yields $B_t = B_0 e^{rt}$ at time t . Furthermore, there is a risky asset (say, an investment in stock) whose price process is denoted by $S = (S_t)_{0 \leq t \leq T}$. We assume that S_t ($0 < t < T$) has a continuous distribution F_{S_t} . The no-arbitrage price² at time 0 of a payoff X_T paid at time $T > 0$ is given by

$$c_0(X_T) = \mathbb{E}[\xi_T X_T], \quad (6.1)$$

where $(\xi_t)_t$ is the state-price density process³ ensuring that $(\xi_t S_t)_t$ is a martingale. Moreover, based on standard economic theory, we assume throughout this chapter that state prices are decreasing with asset prices,⁴ i.e.,

$$\xi_t = g_t(S_t), \quad t \geq 0, \quad (6.2)$$

where g_t is decreasing (in markets where $\mathbb{E}[S_T] > S_0 e^{rT}$). There is empirical evidence that this relationship may not hold in practice, which is called the pricing kernel puzzle (Brown and Jackwerth (2004), Grith et al. (2013)). Many explanations have been provided in the literature (Brown and Jackwerth (2004), Hens and Reichlin (2013)), including state-dependence of preferences (Chabi-Yo et al. (2008)). Therefore, (6.2) is not consistent with a market populated by investors with state dependent preferences. However, we do not tackle the problem of equilibrium and instead study the situation of a small investor whose state-dependent preferences do not influence the pricing kernel that is exogenously given in the market. This is a commonly studied situation since the work of Karatzas et al. (1987).

The functional form (6.2) for $(\xi_t)_t$ allows us to present our results regarding optimal portfolios using $(S_t)_t$ as a reference, which is practical. We will explain in Section 6.5 how the results and characterizations of the optimality of a payoff X_T are tied to its (conditional) anti-monotonicity with ξ_T and do not depend on the functional form (6.2) per se. Note that assumption (6.2) is satisfied by many popular pricing models, including the CAPM, the consumption-based models

²The payoffs we consider are all tacitly assumed to be square integrable, to ensure that all expectations mentioned in the chapter exist. In particular, $c_0(X_T) < +\infty$ for any payoff X_T considered throughout this chapter.

³The process is commonly so designated. However, strictly speaking, it is not a density that is at issue, but rather the product of a discount factor (generally strictly less than 1) and the Radon-Nikodym derivative between the physical measure and the risk-neutral measure.

⁴See e.g., Cox, Ingersoll and Ross (1985) and Bondarenko (2003), who shows that property (6.2) must hold if the market does not allow for statistical arbitrage opportunities, where a statistical arbitrage opportunity is defined as a zero-cost trading strategy delivering at T , a positive expected payoff unconditionally, and non-negative expected payoffs conditionally on ξ_T .

and by exponential Lévy markets in which the market participants use Esscher pricing (Vanduffel et al. (2009)). It is also possible to use a market model in which prices are obtained using the Growth Optimal portfolio (GOP) as numéraire (Platen and Heath (2006)), as is discussed further in Section 6.5. The Black–Scholes model can be seen as a special case of this latter setting (Chapter 1, Section 1.2.2).

6.2 Law-invariant Preferences and Optimality of Path-independent Payoffs

In this section, it is understood that investors have *law-invariant* (state-independent) preferences. This means that they are indifferent between two payoffs having the same payoff distribution (under \mathbb{P}). In this case, any random payoff X_T (that possibly depends on the path of the underlying asset price) admits a path-independent alternative with the same price, which is at least as good for (i.e., desirable in the eyes of) these investors. Recall that a payoff is *path-independent* if there exists some function f such that $X_T = f(S_T)$ holds almost surely. Hence, investors with law-invariant preferences only need to consider path-independent payoffs when making investment decisions. Under the additional (typical) assumption that preferences are *increasing*, any path-dependent payoff can be strictly dominated by a path-independent one that is increasing in the risky asset.⁵

Note that results in this section are related closely to the original work of Cox and Leland (1982; 2000), Dybvig (1988a), Bernard, Boyle and Vanduffel (2014) (see Chapter 1) and Carlier and Dana (2011). These overview results are recalled here to facilitate the exposition of the extensions that are developed in the following sections.

Proposition 6.2.1 shows that for any given payoff there exists a path-independent alternative with the same price that is at least as good for investors with law-invariant preferences. Thus, such an investor needs only to consider path-independent payoffs. All other payoffs are indeed redundant in the sense that they are not needed to optimize the investor’s objective.

Proposition 6.2.1 (Sufficiency of path-independent payoffs). *Let X_T be a payoff with price c and having a cdf F . Then, there exists at least one path-independent payoff $f(S_T)$ with price $c := c_0(f(S_T))$ and cdf F .*

Proposition 6.2.1, however, does not conclude that a given path-dependent payoff can be strictly dominated by a path-independent one. However, if we assume that preferences are increasing, the conclusion was already established in Chapter 1. We recall the result in the following corollary.

Corollary 6.2.2 (Optimal payoffs for increasing law-invariant preferences). *For any payoff Y_T at price c that is not almost surely increasing in S_T there exists a path-independent payoff Y_T^* at price c that is a strict improvement for any investor with increasing and law-invariant preferences.*

6.3 Optimal payoffs under state-dependent preferences.

Many of the contracts chosen by law-invariant investors do not offer protection in times of economic hardship. In fact, due to the observed monotonicity property with S_T , the lowest outcomes for an

⁵This dominance can easily be implemented in practice, as all path-independent payoffs can be replicated statistically with European call and put options as shown e.g., by Carr and Chou (1997) and by Breeden and Litzenberger (1978).

optimal (thus, cost-efficient) payoff occur when the stock price S_T reaches its lowest levels. More specifically, denote by $f(S_T)$ a cost-efficient payoff (with an increasing function f) and by X_T another payoff such that both are distributed with F at maturity. Then, $f(S_T)$ delivers low outcomes when S_T is low and it holds for all $a \geq 0$ that

$$\mathbb{E}[f(S_T)|S_T < a] \leq \mathbb{E}[X_T|S_T < a]. \quad (6.3)$$

Let F be the distribution of a put option with payoff $X_T := (K - S_T)^+ = \max(K - S_T, 0)$. Bernard, Boyle and Vanduffel (2014) (see Chapter 1) show that the payoff of the cheapest strategy with cdf F can be computed explicitly. It is given by $X_T^* = (K - a S_T^{-1})^+$ with $a := S_0^2 \exp(2(\mu - \sigma^2/2)T)$ and is a power put option (with power -1). X_T^* is the cheapest way to achieve the distribution F , whereas the first “ordinary” put strategy (with payoff X_T) is actually the most expensive way to do so. These payoffs interact with S_T in fundamentally different ways, as one payoff is increasing in S_T while the other is decreasing in it. A put option protects the investor against a declining market, in which consumption is more expensive than is otherwise typical, whereas the cost-efficient counterpart X_T^* provides no protection but rather emphasizes the effect of a market deterioration on the wealth received.

As mentioned in the introduction, the use of put options and the demand for insurance (Bernard and Vanduffel (2014a)) are signals that many investors care about states of the economy in which income derived from investment strategies is received. In particular, they may seek strategies that provide protection against declining markets or, more generally, that exhibit a desired dependence with some benchmark.

Hence, in the remainder of this chapter, we consider investors who exhibit state-dependent preferences in the sense that they seek a payoff X_T with a desired distribution *and* a desired dependence with a benchmark asset A_T . In other words, they fix the joint distribution G of the random couple (X_T, A_T) . The optimal state-dependent strategy is the one that solves for

$$\min_{(X_T, A_T) \sim G} c_0(X_T). \quad (6.4)$$

Note that the setting also includes law-invariant preferences as a special (limiting) case when A_T is deterministic. In this case, we effectively revert to the framework of state-independent preferences that we discussed in the previous section. In what follows, we consider as benchmark the underlying risky asset or any other asset in the market, considered at final or intermediate time(s). Moreover, to ensure that the impact of state-dependent preferences on the structure of optimal payoffs is clear, we have organized the rest of the present section along similar lines to those of Section 6.2.

6.3.1 Sufficiency of Twins

In this chapter, any payoff that writes as $f(S_T, A_T)$ or $f(S_T, S_t)$ is called a *twin*. We show first that, in our state-dependent setting, for any payoff there exists a twin that is at least as good. When also assuming that preferences are increasing, we find that optimal payoffs write as twins, and we are able to characterize them explicitly. Conditionally on A_T , optimal twins are increasing in the terminal value of the risky asset S_T .

The following theorems show that for any given payoff there is a twin that is at least as good for investors with state-dependent preferences.

Theorem 6.3.1. (Twins as payoffs with a given joint distribution with a benchmark A_T and price c). *Let X_T be a payoff with price c having joint distribution G with some benchmark A_T , where (S_T, A_T) is assumed to have a joint density with respect to the Lebesgue measure. Then, there exists at least one twin $f(S_T, A_T)$ with price $c = c_0(f(S_T, A_T))$ having the same joint distribution G with A_T .*

Theorem 6.3.1 does not cover the case in which S_T plays the role of the benchmark (because (S_T, S_T) has no density). This interesting case is considered in the following theorem (Theorem 6.3.2).

Theorem 6.3.2 (Twins as payoffs with a given joint distribution with S_T and price c). *Let X_T be a payoff with price c having joint distribution G with the benchmark S_T . Assume that (S_T, S_t) for some $0 < t < T$ has a joint density with respect to the Lebesgue measure. Then, there exists at least one twin $f(S_t, S_T)$ with price $c = c_0(f(S_t, S_T))$ having a joint distribution G with S_T . An example is given by*

$$f(S_t, S_T) := F_{X_T|S_T}^{-1}(F_{S_t|S_T}(S_t)). \quad (6.5)$$

Theorems 6.3.1 and 6.3.2 imply that investors who care about the *joint distribution of terminal wealth with some benchmark A_T* need only consider the twins in both cases, i.e., when (A_T, S_T) is continuously distributed, as in Theorem 6.3.1, or when A_T is equal to S_T , as in Theorem 6.3.2. These results extend Proposition 6.2.1 to the presence of a benchmark and state-dependent preferences. All other payoffs are useless in the sense that they are not needed for these investors per se.⁶

Note that in Theorem 6.3.2, t can be chosen freely in $(0, T)$ and the dependence with respect to S_t is not fixed. So, for instance, replacing $F_{S_t}(S_t)$ with $1 - F_{S_t}(S_t)$ in (6.5) would also lead to the appropriate properties. Hence, there is an infinite number of twins $f(S_t, S_T)$ having the joint distribution G with S_T . All of them have the same price.⁷ The question then arises: how does one select one among them. A natural possibility is to determine the optimal twin $X_T = f(S_t, S_T)$ by imposing an additional criterion. For example, one could define the best twin X_T as the one that minimizes

$$\mathbb{E}[(X_T - H_T)^2], \quad (6.6)$$

where H_T is another payoff that is not a function of S_T . This approach appears natural in the context of simplifying the design of contracts. For instance, start with a geometric Asian option and compute its joint distribution G with S_T . Then, all twins as in (6.5) have the same price but one of them may be closer to the original Asian derivative (in the sense of minimizing the distance, as in (6.6)). Note that since all marginal distributions are fixed, the criterion (6.6) is equivalent to maximizing the correlation between X_T and H_T .

6.3.2 Optimality of Twins

Next, we investigate the cost optimality of twins. As discussed above, if the benchmark A_T coincides with S_T , then all twins that satisfy $(X_T, A_T) \sim G$ have the same cost and the problem of searching

⁶This finding is consistent with the result obtained by Takahashi and Yamamoto (2013), who apply it to replicate a joint distribution in the hedge fund industry.

⁷To see this, recall that the joint distribution between the twin $f(S_t, S_T)$ and S_T is fixed and thus also the joint distribution between the twin and ξ_T (as ξ_T is a decreasing function of S_T due to (6.2)). All twins $f(S_t, S_T)$ with such a property have the same price $\mathbb{E}[\xi_T f(S_t, S_T)]$.

for the cheapest one is not meaningful. However, this observation is no longer true when the benchmark A_T has a density with S_T . In this case, the cheapest twin is determined by Theorem 6.3.3 that extends Corollary 1.2.3 in Chapter 1 to the state-dependent case. Corollary 1.2.3 finds that among the infinite number of payoffs with a given distribution F , the cheapest one is increasing in S_T . In the state-dependent setting one has that optimal payoffs are increasing in S_T , conditionally on A_T .

Theorem 6.3.3 (Cost optimality of twins). *Assume that (S_T, A_T) has joint density with respect to the Lebesgue measure. Let G be a bivariate cumulative distribution function. The optimal state-dependent strategy determined by*

$$\min_{(X_T, A_T) \sim G} c_0(X_T) \quad (6.7)$$

has an almost surely unique solution X_T^ which is a twin of the form $f(S_T, A_T)$. X_T^* is almost surely increasing in S_T , conditionally on A_T , and given by*

$$X_T^* := F_{X_T|A_T}^{-1}(F_{S_T|A_T}(S_T)). \quad (6.8)$$

Recall from Section 6.2 that when preferences are law-invariant, optimal payoffs are path-independent and increasing in S_T . When preferences are state-dependent, we observe from expression (6.8) that optimal state-dependent payoffs may become path-dependent, and are increasing in S_T , conditionally on A_T . We end this section with a corollary derived from Theorem 6.3.3. The result echoes the one established for investors with law-invariant preferences in the previous section (Corollary 6.2.2)

Corollary 6.3.1 (Cheapest twin). *Assume that (S_T, A_T) has joint density with respect to the Lebesgue measure. Let G be a bivariate cumulative distribution function. Let X_T be a payoff such that $(X_T, A_T) \sim G$. Then, X_T is the cheapest payoff if and only if, conditionally on A_T , X_T is (almost surely) increasing in S_T .*

6.4 Improving Security Design

In this section, we show that the results above are useful in designing balanced and transparent investment policies for retail investors as well as financial institutions:

1. If the investor who buys the financial contract has law-invariant preferences and if the contract is not increasing in S_T , then there exists a strictly cheaper derivative (cost-efficient contract) that is strictly better for this investor. We find its design by applying Corollary 1.2.3 from Chapter 1.
2. If the investor buys the contract because of the interaction with the market asset S_T , and the contract depends on another asset, then we can apply Theorem 6.3.2 to simplify its design while keeping it “at least as good.” The contract then depends, for example, on S_T and S_t for some $t \in (0, T)$.
3. If the investor buys the contract because he likes the dependence with a benchmark A_T , which is not S_T , and if the contract does not only depend on A_T and S_T , then we use Theorem 6.3.1 to construct a simpler one that is “at least as good” and that writes as a function of S_T and A_T . Finally, if the obtained contract is not increasing in S_T conditionally on A_T , then it is also possible to construct a strictly cheaper alternative using Theorem 6.3.3 and Corollary 6.3.1.

We now use the Black–Scholes market to illustrate these three situations. We begin with the example of an Asian option with fixed strike. The published version of this chapter contains an example with floating strike. Consider a fixed strike (continuously monitored) geometric Asian call with payoff given by

$$Y_T := (G_T - K)^+. \quad (6.9)$$

Here, K denotes the fixed strike and G_T is the geometric average of stock prices from 0 to T , defined as

$$\ln(G_T) := \frac{1}{T} \int_0^T \ln(S_s) ds. \quad (6.10)$$

We can now apply the results derived above to design products that improve upon Y_T .

a. Use of cost-efficiency payoff for investors with increasing law-invariant preferences. By applying Corollary 1.2.3 to the payoff Y_T (6.9), one finds that the cost-efficient payoff associated with a fixed strike (continuously monitored) geometric Asian call is

$$Y_T^* = d \left(S_T^{1/\sqrt{3}} - \frac{K}{d} \right)^+, \quad (6.11)$$

where $d = S_0^{1-\frac{1}{\sqrt{3}}} e^{\left(\frac{1}{2}-\sqrt{\frac{1}{3}}\right)\left(\mu-\frac{\sigma^2}{2}\right)T}$.

While the above result can also be found in Bernard, Boyle and Vanduffel (2014) (see Chapter 1), they are worth considering here for the purpose of comparison with what follows. Note that letting K go to zero provides a cost-efficient payoff that is equivalent to the geometric average G_T .

b. A twin that is useful for investors who care about the dependence with S_T . By applying Theorem 6.3.2 to the payoff G_T , we can find a twin payoff $R_T(t) = f(S_t, S_T)$ such that

$$(S_T, R_T(t)) \sim (S_T, G_T). \quad (6.12)$$

By definition, this twin preserves existing dependence between G_T and S_T . However, compared to the original contract it is simpler and “less” path-dependent, as it depends only on two values of the path of the stock price. Interestingly, the call option written on $R_T(t)$ and the call option written on G_T have the same joint distribution with S_T . Consequently,

$$(S_T, (R_T(t) - K)^+) \sim (S_T, (G_T - K)^+). \quad (6.13)$$

$(R_T(t) - K)^+$ is therefore a twin equivalent to the fixed strike geometric Asian call (as in Theorem 6.3.2). We can compute $R_T(t)$ by applying Theorem 6.3.2, and we find that

$$R_T(t) = S_0^{\frac{1}{2}-\frac{1}{2\sqrt{3}}}\sqrt{\frac{T-t}{t}} S_t^{\frac{T}{t}\frac{1}{2\sqrt{3}}}\sqrt{\frac{t}{T-t}} S_T^{\frac{1}{2}-\frac{1}{2\sqrt{3}}}\sqrt{\frac{t}{T-t}}, \quad (6.14)$$

where t is freely chosen in $(0, T)$.⁸ The equality of joint distributions exposed in (6.13) implies that the call option written on $R_T(t)$ has the same price as the original fixed strike (continuously monitored) geometric Asian call (6.9). The time-0 price of both contracts is therefore

$$c_0((R_T(t) - K)^+) = S_0 e^{-\frac{rT}{2} - \frac{\sigma^2 T}{12}} \Phi(\tilde{d}_1) - K e^{-rT} \Phi(\tilde{d}_2), \quad (6.15)$$

⁸Formula (6.14) is based on the expression (6.5) for a twin dependent on S_t and S_T . Note that there is no uniqueness. For example, $1 - F_{S_t|S_T}(S_t)$ is also independent of S_T , and we can thus also consider $H_T(t) := F_{X_T|S_T}^{-1}(1 - F_{S_t|S_T}(S_t))$ as a suitable twin ($0 < t < T$) satisfying the joint distribution, as in (6.12). In this case, one obtains $H_T(t) = S_0^{\frac{1}{2}+\frac{1}{2\sqrt{3}}}\sqrt{\frac{T-t}{t}} S_t^{-\frac{T}{t}\frac{1}{2\sqrt{3}}}\sqrt{\frac{t}{T-t}} S_T^{\frac{1}{2}+\frac{1}{2\sqrt{3}}}\sqrt{\frac{t}{T-t}}$.

where $\tilde{d}_1 = \frac{\ln(S_0/K) + rT/2 + \sigma^2 T/12}{\sigma\sqrt{T/3}}$ and $\tilde{d}_2 = \tilde{d}_1 - \sigma\sqrt{T/3}$ (see Kemna and Vorst (1990)).

c. Choosing among twins. The construction in Theorem 6.3.2 depends on t . Maximizing the correlation between $\ln(R_T(t))$ and $\ln(G_T)$ is nevertheless a possible way to select a specific t . The covariance between $\ln(R_T(t))$ and $\ln(G_T)$ is provided by

$$\text{cov}(\ln(R_T(t)), \ln(G_T)) = \frac{\sigma^2}{2} \left(\frac{T}{2} + \frac{\sqrt{t}\sqrt{T-t}}{2\sqrt{3}} \right)$$

and, by construction of $R_T(t)$, the standard deviations of $\ln(R_T(t))$ and $\ln(G_T)$ are both equal to $\sigma\sqrt{\frac{T}{3}}$. Maximizing the correlation coefficient is therefore equivalent to maximizing the covariance, and thus of $f(t) = (T-t)t$. This maximum is obtained for $t^* = \frac{T}{2}$, and the maximal correlation ρ_{\max} between $\ln(R_T(t))$ and $\ln(G_T)$ is

$$\rho_{\max} = \frac{3}{4} + \frac{\sqrt{3}\sqrt{(T-t^*)t^*}}{4T} = \frac{3}{4} + \frac{\sqrt{3}}{8} \approx 0.9665,$$

which shows that the optimal twin is highly correlated to the initial Asian, while being considerably simpler. Note that both the maximum correlation and the optimum $R_T(\frac{T}{2})$ are robust to changes in market parameters.

In the published version of this chapter (Bernard, Moraux, Rüschenendorf, and Vanduffel (2015)), we discuss several additional contributions to the field of portfolio management. We first derive the optimal investment for an expected utility maximizer who has a constraint on the dependence with a given benchmark. Next, we revisit optimal strategies for target probability maximizers (see Browne (1999) and Spivak and Cvitanic (1999)), and we extend this problem in two directions by adding dependence constraints and by considering a random target. In both cases, we derive analytical solutions that are given by twins.

6.5 Conclusions

In this chapter, we introduce a state-dependent version of the optimal investment problem. We deal with investors who target a known wealth distribution at maturity (as in the traditional setting) and *additionally* desire a particular interaction with a random benchmark. We show that optimal contracts depend at most on two underlying assets, or on one asset evaluated at two different dates, and we are able to characterize and determine them explicitly. Our characterization of optimal strategies allows us to extend the classical expected utility optimization problem of Merton to the state-dependent situation. Throughout the chapter, we have assumed that the state-price density process ξ_T is a decreasing functional of the risky asset price S_T and that there is a single risky asset. It is possible to relax these assumptions and yet still to provide explicit representations of optimal payoffs. However, the optimality is then no longer related to path-independence properties.

Throughout the chapter, we assumed that ξ_T is decreasing in S_T (in (6.2)). Moreover, we use the one-dimensional Black-Scholes model to illustrate our findings. However, the case of multidimensional markets described by a price process $(S_t^{(1)}, \dots, S_t^{(d)})_t$ is essentially included in the results presented in this chapter, assuming that the state-price density process $(\xi_t)_t$ of the risk-neutral measure chosen for pricing is of the form $\xi_t = g_t(h_t(S_t^{(1)}, \dots, S_t^{(d)}))$ with some real functions g_t, h_t (as

in Bernard, Maj and Vanduffel (2011) who considered the state-independent case). All results in the chapter apply by replacing the one-dimensional stock price process S_t by the one-dimensional process $h_t(S_t^{(1)}, \dots, S_t^{(d)})$. In addition, we have assumed that asset prices are continuously distributed, which amounts essentially to assuming that the state-price density process ξ_t is continuously distributed at any time. An extension to the case in which ξ_t may have atoms is possible but not in the scope of the present chapter.

A straightforward extension of the results presented in this chapter is to consider the market model of Platen and Heath (2006) using the Growth Optimal Portfolio (GOP). It consists of replacing the state-price density process ξ_t by $1/S_t^*$, where S_t^* denotes the value of the GOP at time t . Then our results are also valid in their setting, where the GOP is taken as the reference (see Bernard, Chen, and Vanduffel (2014) or Chapter 5). Other dependence constraints can be considered, e.g. a constraint on the correlation between the terminal wealth and a benchmark (Bernard and Vanduffel (2014b) or Chapter 4).

Chapter 7

Conclusions and Research Directions

We have presented the concept of cost-efficiency in Part I and some direct applications in Part II. This concept is a novel way to look at investment problems. It allows to characterize the properties of optimal portfolios or optimal wealth allocation in settings when the agent is concerned by the distribution of his final wealth only. We have seen in the later applications in Part III that cost-efficiency can be extended to more realistic settings with constraints, which make the problem “state-dependent.” Again, we characterize the properties of optimal portfolio in a context when the decisions are subject to a “state-dependent” context in which not only the distribution of final wealth matters but also in which states the cash-flows are received. It is well-known that agents typically do not make their decision in isolation by looking only at the distribution of their wealth. Preferences change as the conditions change. If the happiness of the agent depends not only on the amount of money he receives but also on some extra reference information, such as the amount of money her friends receive, her health status or how the market index is doing, then the agent has state-dependent preferences. If it is clear that agents have state-dependent preferences, it is not so clear how to model them. Chapter 6 proposes a general way to model state-dependent preferences.

Most of economic theories (even the most recent behavioural theories such as cumulative prospect theory of Tversky and Kahneman (1992)) are law-invariant. However, it is clear that the world is more and more integrated and that the optimization faced by companies and agents are more and more “state-dependent.” Risk and performance are assessed in a relative way by comparing to benchmarks rather than by considering absolute quantities (e.g., when comparing portfolio managers (Lo (2008); Pojarliev and Levich (2008)) who are usually compared under the same market conditions and with access to the same assets).

Another example is the recent debate on capital requirements to cover systemic risk (Adrian and Brunnermeier (2011), Acharya (2009), Bernard, Brechmann, and Czado (2013), Bernard and Czado (2015)). It is well documented in the literature that companies tend to be more dependent in a crisis, i.e., under stress conditions in the economy. During the financial crisis of 2007-2009, taxpayers worldwide had to bail out numerous financial institutions. Governments are still trying to understand why regulation failed, why capital requirements were insufficient and how a guaranty fund can be best structured to withstand future crises. To do so, one needs to understand the risk that each institution represents to the financial system. In the financial and insurance industry, capital requirements have the following common properties. First, they depend solely on the distribution of the institution’s risk and not on the outcomes in different states of the world. Second, capital requirements and margins calculations treat each institution in isolation. An important element is missing in the above risk assessment: the dependence between the individual institution

and the financial system. Specifically, Acharya (2009) cautions against a blanket call for more capital in the financial industry and instead recommends “regulating each bank as a function of both its joint (correlated) risk with other banks [and] its individual (bank-specific) risk.” The basic principle here is that while the risk of failure of a business can never be eliminated, it is not acceptable to have multiple large-scale failures occurring simultaneously. Risk measures and decision models should assess risk not only on distributional properties but also on the interaction with other economic variables and on the states of the economy in which cash-flows are exchanged (scenario constraints). Regulators should assess each bank using both its individual risk and its joint risk with other banks (systemic risk). The most popular proposals are based on capital requirements computed conditionally on the economy being under stress or conditionally on one of the big market players going bankrupt. By definition, these capital requirements are state-dependent. The impact of imposing such capital requirements on the global economy is still uncertain. This is one of the avenue of research that I will pursue in the coming years, i.e. to understand the impact on the pricing kernel and on the equilibrium in the economy when agents have state-dependent preferences. Basak and Shapiro (2001) show that Value-at-Risk requirements applied to all players in the industry impact the pricing kernel and may have adverse effects on the financial market. I want to extend their study to encompass cases in which capital requirements are state-dependent. The ultimate goal is to understand the potential impact on the industry of a regulatory framework based on state-dependent risk assessment or scenarios. To do so, I plan to use a theoretical approach and exploit characterization of optimal portfolio choice from previous work (Chapters 5 and 6 corresponding to Bernard, Chen, and Vanduffel (2014) and Bernard, Moraux, Rüschendorf, and Vanduffel (2015)).

We also note that we have solved portfolio optimization problems in the context of maximizing the agent’s objective function given a pricing kernel in the economy. Most examples are pursued in the Black-Scholes setting in which the pricing kernel is a direct function of the stock price. The assumptions behind this optimization are that the investor is a small investor and that there is no feedback effect, i.e. that his investment decisions do not change market prices. If investors are large players or if all investors in the market have state-dependent preferences, the equilibrium in the market will be affected. This is also left for future research. Note that Chapter 2 suggests that the equilibrium will have similar properties when all agents have increasing law-invariant preferences and when all agents maximize expected utility of terminal wealth with concave utility. This is also an interesting research direction.

In Chapter 4 (Bernard and Vanduffel (2014b)) we find dynamic optimal portfolios using mean variance optimization, equivalently by maximizing the portfolio Sharpe ratio (Sharpe (1994)). The Sharpe ratio is the ratio between the excess expected return (against a risk-free asset) and the standard deviation of the portfolio. In the current market environment, assets do not display symmetric distributions of returns. They typically have fat tails, and therefore the two first moments (mean and variance) are not sufficient to choose optimal portfolios. The sole criteria of maximizing Sharpe ratios may thus appear limited. It is necessary to consider higher-order moments such as the Skewness and the Kurtosis to fully understand the properties of an investment portfolio. Chapter 2 establishes a correspondence between the distribution of final wealth and the utility function that explains the demand for this distribution. It is now of interest to further study this connection to better link properties of the utility function and properties of the distribution of final wealth. Another limitation is that consumption only occurs at maturity. It is possible to extend our study to intermediary consumption (Bernard and Kwak (2016a)) and we are working on designing cost-efficient dynamic strategies. This may, for example, help ultimately pension funds to manage their funds dynamically and in a cost-efficient way.

We also plan to consider more advanced performance measures that go beyond the Sharpe ratio

and that are used in the industry. A typical example is the Omega risk measure (Keating and Shadwick (2002b,a)) or the Sortino ratio (Sortino and Price (1994), Kaplan and Knowles (2004)). The Omega risk measure or the Sortino ratio account for the asymmetry of the distribution and reward the presence of skewness (Bertrand and Prigent (2011)). It is of interest to derive the optimal portfolio for a manager that maximizes the Omega risk measure and the Sortino ratio of a portfolio in a dynamic setting. As these measures are law-invariant (i.e. they only depend on the probability distribution of the portfolio), we expect that the cost-efficiency theory will be useful in characterizing the optimal portfolios (Bernard, Boyle, and Vanduffel (2014)). Specifically, cost-efficiency allows to characterize the shape of the optimal wealth. This characterization can be used to design algorithms to solve optimal portfolio selection for any law-invariant objective function and thus also in the special case of the Omega or Sortino ratios.

A recent alternative to existing performance risk measures has been proposed by Aumann and Serrano (Aumann and Serrano (2008)). Their risk measure is now known as the Aumann-Serrano index and so far it has not been studied in the context of portfolio selection. However, this index is consistent with first order and second order stochastic dominance (Homm and Pigorsch (2012)), which are important features in portfolio selection. Hence, it is of great interest to better understand the index and to study if it can be useful in portfolio selection. Precisely, we plan to study optimal portfolio derived by optimizing this criteria, and compare them to the ones obtained with more traditional performance measures that are used in the industry (Sharpe ratio, Sortino ratio, Omega risk measure, etc.). Furthermore, we also intend to closely investigate the hedging properties of the proposed strategies, and to assess their risk using stochastic ordering concepts that we borrow from actuarial science.

The above research directions built directly on the concept of cost-efficiency and its extension to include state-dependent constraints. I also plan to further work in the three other research areas presented in the introduction. I thus conclude with two additional research directions. The first one builds on my work in insurance and hedging financial derivatives (first and second area of research listed in the introduction). The second is about recent research progress in risk management and dependence modelling (third area of research).

I have ongoing research on variable annuities (VA) that are insurance products. A typical VA consists of two phases. First, the policyholder makes regular payments into a fund managed by the insurer (accumulation phase). Then, she receives income from the insurance company with some minimum guarantees for some given period (payment phase). As the Solvency II regulatory framework requires a market consistent value of liabilities, VA issuers face new challenges. Reserve requirements for VA products are highly “scenario dependent” and are not known at issuance (technical provisions and solvency capital requirements change over time depending on market conditions, e.g., interest rates, volatility, equity market...). The issuer’s income from most VAs is often computed as a fixed percentage fee of the fund under management, which means that it is high/low when the market goes up/down. Unfortunately, the market value of embedded guarantees moves in the opposite direction. When equity goes down, guarantees become expensive. Moreover, volatility typically increases so that delta hedging programs need to be rebalanced more often (more costly), and hedging programs that require the purchase of options become expensive, as options prices increase when volatility increases. Thus, the VA issuers’ income becomes smaller when it is most needed. Our motivation comes from VAs that have recently been offered in the industry. We cite, for instance, the SunAmerica GLWB issued in 2011, with 8% rollup and fees linked to the Volatility Index (VIX) reported by the Chicago Board Options Exchange, and the latest variable annuity contracts offered by American General Life in 2014 (see the prospectus dated May 1, 2014 of the “Polaris Variable Annuities Choice IV”). In this latter contract, the fee rate adjustment is

tied to changes in the VIX index. We propose to model and study the benefits of charging state-dependent fees in VAs that depend on the volatility level. We expect that the fees will be lower on average for the policyholder and that they will provide better matching for the insurer between the actual value of the guarantee and the premium collected by the insurer. We thus expect to improve the hedging program for VAs using state-dependent fees. Such an approach challenges the existing VA modeling; see e.g., Coleman, Li, and Patron (2006), Milevsky and Salisbury (2006), Bauer, Kling, and Russ (2008), Chen, Vetzal, and Forsyth (2008), Dai, Kuen Kwok, and Zong (2008)... Almost all of the research to date has focused on a fixed fee rate, i.e., the policyholder has to pay a constant percentage of the fund. VAs with state-dependent fees linked to volatility, however, are realistic and attractive in the insurance industry. Linking the issuers income to the volatility index VIX is good for insurers, as it allows for a better match between the hedging cost and the values of the guarantees in periods of high volatility. It is also good for policyholders, as they will pay lower fees during long periods of low volatility markets. By accepting the idea of sharing the hedging difficulties of the insurer when markets are volatile, policyholders contribute to reducing their risk and thus to the probability that their benefits will ultimately be paid. The proposed design has a state-dependent fee, as it depends on the traded VIX index. State-dependent fees have recently appeared in the literature with the work of Bernard, Hardy, and MacKay (2014), MacKay, Augustyniak, Bernard, and Hardy (2016) and Bernard and MacKay (2015). These papers consider fees linked to the fund value: fees are paid only when the fund goes above some threshold value. However, such a state-dependent fee might create moral issues as the insurer manages the fund and has reduced incentives to increase the fund value. In particular, if the fund value is close to the threshold, the insurer has incentives to misreport the fund value or to create a last minute loss in order to move the fund value to be just below the threshold. It is thus important to link the VA fee to variables that are not subject to manipulation risk. State-dependent fees can also help to decrease the surrender incentives in equity-linked insurance products with guarantees, and to facilitate the risk management of the guarantees for insurers.

We recently made significant progress in the third area of research presented in the introduction, i.e. in risk management and dependence modelling. We have received the 2014 *PRMIA Annual Frontiers in Risk Management Award* for the paper “A new approach to assessing model risk,” which is now published in *Journal of Banking and Finance* (Bernard and Vanduffel (2015)). Building on this paper, we aim at developing new risk management tools and novel risk indicators to assess dependence among market variables under specific market conditions in a forward-looking way. Specifically, we want to capture changes in the dependence among market variables under certain scenarios (for instance, during a crisis). We are particularly interested in quantifying changes in dependence in the tail. To measure tail risk, Kelly and Jiang (2014) apply extreme value theory to stock price returns. Their new tail risk index is linked to the probability of observing an extreme return and measures the fatness of the distributions tail. The authors show that this index has a strong predictive power and that it is highly correlated with other indicators, such as option implied kurtosis and implied skewness (Bakshi, Kapadia, and Madan (2003)) for the S&P 500 index, and is thus closely related to the tail risks perceived by option market participants. We propose a new approach to infer the dependence among stock prices using option prices only and in particular the dependence in the tails. This approach is related to the popular forward-looking measure for volatility, the VIX index, which measures the implied volatility of the S&P 500 index using index options. It is clear that the volatility of the S&P 500 index is influenced both by changes in the correlations among the 500 institutions that comprise the index and by changes in the respective volatilities of its components. We suggest to design a dependence measure that is not influenced by changes in volatility and that is driven solely by the information contained in option prices.

Several measures of dependence and comovement have been developed using stock prices. Here, the difficulty is that we want to use option prices as the only source of available information, and only a few proposals in the literature have investigated this possibility.

One approach that has been taken consists of building an appropriate multi-asset model consistent with observed index options and individual stock options. Cont and Deguest (2013), for instance, propose a mixture of models that can reproduce some sets of multivariate option prices and individual options. They then derive a notion of implied correlation as the correlation matrix to fit their model. The authors are able to quantify model risk on implied correlation by studying the set of such matrices. Avellaneda and Boyer-Olson (2002) and Jourdain and Sbair (2012) have conducted related studies that include model assumptions. The second approach is to measure the implied dependence among assets without making model assumptions. The CBOE S&P 500 Implied Correlation Indexes represent attempts to measure the implied correlation in a model-free way. In July 2009, the Chicago Board Options Exchange (CBOE) began disseminating such indexes, which are now well accepted measures based on individual implied volatilities and index implied volatility and thus driven by option prices only. Driessen, Maenhout, and Vilkov (2009, 2012) use a stochastic correlation model to extract an implied correlation. The authors define implied correlation similarly to the CBOE implied correlation index formula, although the volatilities are not implied volatilities, but model free implied variances from index and individual options. However, they constrain the model by assuming that all pairwise correlations are driven by the same time-varying correlation as a mean reverting process (in the same spirit as Cochrane, Longstaff, and Santa-Clara (2008)). Extensions have also been proposed by Buss and Vilkov (2012) and by Backus, Chernov, and Martin (2011) who infer disaster risk premia from options.

We plan to study the shortcomings of the CBOE S&P 500 Implied Correlation Indexes and propose a new approach to measure implied correlation that overcomes these deficiencies. For example, the CBOE implied correlation is driven by at-the-money option prices only, depends on marginal distributions and is a global correlation measure. We propose a measure that does not depend on margins and that can be computed conditionally on certain events, not just globally. Also, the CBOE implied correlation seeks to be model-free. However, we can show that the correlation parameter that is obtained from this procedure is consistent with a very strong assumption that each component of the index, as well as the index itself, is modelled by a LogNormal distribution. Our method uses the Rearrangement Algorithm (RA) of Puccetti and Rüschendorf (2012), which was initially developed as a method to construct the dependence among the variables X_j such that the distribution of the variance of the sum $S = X_1 + \dots + X_n$ is as small as possible, and which was applied later successfully to finding bounds on the Value-at-Risk of portfolios (Embrechts, Puccetti, and Rüschendorf (2013); Bernard, Rüschendorf, and Vanduffel (2016)). Once we have adapted this algorithm, it can be used to infer a dependence structure (copula) that is consistent with the marginal distributions and with knowledge of the distribution of the sum, of the difference or, more generally, of a weighted sum (Bernard, Rüschendorf, and Vanduffel (2016), Bernard and McLeish (2016)). We started a collaboration with O. Bondarenko (Bernard, Bondarenko, and Vanduffel (2016)) that we plan to pursue in the coming years. This last project is timely and relevant to the current economic environment. Many practitioners misinterpret and potentially misuse the CBOE implied correlation index. First, our proposal will contribute to providing a better understanding of what the CBOE implied correlation index is. Second, we will propose new indices that compute tail correlations. These are more refined measures of correlation and, as such, they have the potential to provide a significant improvement over trading strategies based on the CBOE implied correlation index. In addition, our approach is quite generic and can be applied in a variety of contexts, including in developing systemic risk indicators.

References

- ACHARYA, V. (2009): “A theory of systemic risk and design of prudential bank regulation,” *Journal of Financial Stability*, 5, 224–255.
- ADRIAN, T., AND M. BRUNNERMEIER (2011): “CoVaR,” Discussion paper, National Bureau of Economic Research.
- AHMED, A. (2011): “New Investment Strategy: Preparing for End Times,” *New York Times*.
- AÏT-SAHALIA, Y., AND A. LO (2001): “Nonparametric Estimation of State-Price Densities implicit in Financial Asset Prices,” *Journal of Finance*, 53(2), 499–547.
- ALLAIS, M. (1953): “Le Comportement de l’Homme Rationnel Devant le Risque: Critique des Axiomes et Postulats de l’École Américaine,” *Econometrica*, 21(4), 503–546.
- AMENC, N., AND L. MARTELLINI (2011): “In Diversification we trust?,” *Journal of Portfolio Management*, Winter, 1–2.
- AMIN, G. S., AND H. M. KAT (2003): “Hedge Fund Performance 1990-2000: Do the “Money Machines” Really Add Value?,” *Journal of Financial and Quantitative Analysis*, 38(2), 251–274.
- ARTZNER, P., F. DELBAEN, J.-M. EBER, AND D. HEATH (1999): “Coherent measures of risk,” *Mathematical finance*, 9(3), 203–228.
- AUMANN, R. J., AND R. SERRANO (2008): “An economic index of riskiness,” *Journal of Political Economy*, 116(5), 810–836.
- AVELLANEDA, M., AND D. BOYER-OLSON (2002): “Reconstruction of volatility: pricing index options by the steepest descent approximation,” *Courant Institute-NYU Working Paper*.
- BACKUS, D., M. CHERNOV, AND I. MARTIN (2011): “Disasters implied by equity index options,” *The journal of finance*, 66(6), 1969–2012.
- BAGNOLI, M., AND T. BERGSTROM (2005): “Log-concave probability and its applications,” *Economic Theory*, 26(2), 445–469.
- BAKSHI, G., N. KAPADIA, AND D. MADAN (2003): “Stock return characteristics, skew laws, and the differential pricing of individual equity options,” *Review of Financial Studies*, 16(1), 101–143.
- BARBARIN, J., AND P. DEVOLDER (2005): “Risk measure and fair valuation of an investment guarantee in life insurance,” *Insurance: Mathematics and Economics*, 37(2), 297–323.

- BARLOW, R. E., D. J. BARTHOLOMEW, J. BREMNER, AND H. D. BRUNK (1972): *Statistical inference under order restrictions: the theory and application of isotonic regression*. Wiley New York.
- BASAK, S., AND G. CHABAKAURI (2010): “Dynamic mean-variance asset allocation,” *Review of Financial Studies*.
- BASAK, S., AND A. SHAPIRO (2001): “Value-at-Risk Based Risk Management: Optimal Policies and Asset Prices,” *Review of Financial Studies*, 14(2), 371–405.
- BAUCELLS, M., AND F. H. HEUKAMP (2006): “Stochastic dominance and cumulative prospect theory,” *Management Science*, 52(9), 1409–1423.
- BAUER, D., A. KLING, AND J. RUSS (2008): “A universal pricing framework for guaranteed minimum benefits in variable annuities,” *Astin Bulletin*, 38(02), 621–651.
- BEARE, B. (2011): “Measure preserving derivatives and the pricing kernel puzzle,” *Journal of Mathematical Economics*, 47(6), 689–697.
- BEGHIN, L., AND E. ORSINGER (1999): “On the maximum of the generalized Brownian bridge,” *Lithuanian Mathematical Journal*, 39(2), 157–167.
- BEN-TAL, A., AND A. NEMIROVSKI (1998): “Robust convex optimization,” *Mathematics of Operations Research*, 23(4), 769–805.
- BERKELAAR, A., R. KOUWENBERG, AND T. POST (2004): “Optimal Portfolio Choice Under Loss Aversion,” *Review of Economics & Statistics*, 86(4), 973–987.
- BERNARD, C. (2013): “Risk sharing and pricing in the reinsurance market,” in *Handbook of Insurance*, pp. 603–626. Springer.
- BERNARD, C., O. BONDARENKO, AND S. VANDUFFEL (2016): “Rearrangement Algorithm and Maximum Entropy,” *Available at SSRN*.
- BERNARD, C., AND P. BOYLE (2009): “Mr. Madoff’s amazing returns: An analysis of the split-strike conversion strategy,” *Journal of Derivatives*, 17(1).
- (2011a): “Monte Carlo methods for pricing discrete Parisian options,” *European Journal of Finance*, 17(3), 169–196.
- (2011b): “A natural hedge for equity indexed annuities,” *Annals of Actuarial Science*, 5(02), 211–230.
- BERNARD, C., P. BOYLE, AND W. GORNALL (2009): “Locally-capped investment products and the retail investor,” *Journal of Derivatives*, 18(4), 72–88.
- BERNARD, C., P. BOYLE, AND S. VANDUFFEL (2014): “Explicit representation of cost-efficient strategies,” *Finance*, 35(2), 5–55.
- BERNARD, C., AND P. P. BOYLE (2010): “Explicit representation of cost-efficient strategies,” *AFFI december meeting*.
- BERNARD, C., E. BRECHMANN, AND C. CZADO (2013): “Statistical assessments of systemic risk measures,” *Chapter 6 in the Handbook on Systemic Risk, edited by JP Fouque and J Langsam*.

- BERNARD, C., AND A. CHEN (2009): “On the regulator–insurer interaction in a structural model,” *Journal of computational and applied mathematics*, 233(1), 3–15.
- BERNARD, C., J. S. CHEN, AND P. BOYLE (2016): “Power Options in Executive Compensation,” *Journal of Derivatives*, 23, 9–20.
- BERNARD, C., J. S. CHEN, AND S. VANDUFFEL (2014): “Optimal portfolios under worst-case scenarios,” *Quantitative Finance*, 14(4), 657–671.
- (2015): “Rationalizing investors choice,” *Journal of Mathematical Economics*, 59, 10–23.
- BERNARD, C., AND Z. CUI (2011): “Pricing timer options,” *Journal of Computational Finance*, 15(1), 69–104.
- (2014): “Prices and asymptotics for discrete variance swaps,” *Applied Mathematical Finance*, 21(2), 140–173.
- BERNARD, C., Z. CUI, AND D. MCLEISH (2012): “Nearly exact option price simulation using characteristic functions,” *International Journal of Theoretical and Applied Finance*, 15(7).
- (2014): “Convergence of the discrete variance swap in time-homogeneous diffusion models,” *Quantitative Finance Letters*, 2(1), 1–6.
- (2016): “On the martingale property in stochastic volatility models based on time-homogeneous diffusions,” *Mathematical Finance*, *forthcoming*.
- BERNARD, C., AND C. CZADO (2013): “Multivariate option pricing using copulae,” *Applied Stochastic Models in Business and Industry*, 29(5), 509–526.
- (2015): “Conditional quantiles and tail dependence,” *Journal of Multivariate Analysis*, 138, 104–126.
- BERNARD, C., M. DENUIT, AND S. VANDUFFEL (2015): “Measuring portfolio risk under partial dependence information,” *Available at SSRN 2406377*.
- BERNARD, C., AND M. GHOSOUB (2010): “Static portfolio choice under cumulative prospect theory,” *Mathematics and Financial Economics*, 2(4), 277–306.
- BERNARD, C., M. HARDY, AND A. MACKAY (2014): “State-dependent fees for variable annuity guarantees,” *Astin Bulletin*, 44(03), 559–585.
- BERNARD, C., X. HE, J.-A. YAN, AND X. Y. ZHOU (2015): “Optimal Insurance Design Under Rank-Dependent Expected Utility,” *Mathematical Finance*, 25, 154–186.
- BERNARD, C., S. JI, AND W. TIAN (2013): “An optimal insurance design problem under Knightian uncertainty,” *Decisions in Economics & Finance*, 36(2), 99–124.
- BERNARD, C., X. JIANG, AND S. VANDUFFEL (2012): “Note on “Improved Fréchet bounds and model-free pricing of multi-asset options”,” *Journal of Applied Probability*, 49(3), 866–875.
- BERNARD, C., X. JIANG, AND R. WANG (2014): “Risk aggregation with dependence uncertainty,” *Insurance: Mathematics and Economics*, 54, 93–108.
- BERNARD, C., A. KOLKIEWICZ, AND J. TANG (2016): “Variable Annuities with Fees Tied to VIX,” *Working paper*.

- BERNARD, C., AND M. KWAK (2016a): “Dynamic preferences for popular investment strategies in pension funds,” *Scandinavian Actuarial Journal*, 1(5), 398–419.
- (2016b): “Semi-static Hedging of Variable Annuities,” *Insurance: Mathematics and Economics*, 67, 173–186.
- BERNARD, C., AND O. LE COURTOIS (2012a): “Asset risk management of participating contracts,” *Asia-Pacific Journal of risk & insurance*, 6(2).
- (2012b): “Performance regularity: A new class of executive compensation packages,” *Asia-Pacific Financial Markets*, 19(4), 353–370.
- BERNARD, C., O. LE COURTOIS, AND F. QUITTARD-PINON (2005a): “Market value of life insurance contracts under stochastic interest rates and default risk,” *Insurance: Mathematics and Economics*, 36(3), 499–516.
- (2005b): “A new procedure for pricing Parisian options,” *Journal of Derivatives*, 12(4), 45–53.
- (2005c): “A study of mutual insurance for bank deposits,” *The Geneva Risk and Insurance Review*, 30(2), 129–146.
- (2006): “Development and pricing of a new participating contract,” *North American Actuarial Journal*, 10(4), 179–195.
- (2008): “Pricing derivatives with barriers in a stochastic interest rate environment,” *Journal of Economic Dynamics and Control*, 32(9), 2903–2938.
- (2010): “Protection of a Company Issuing a Certain Class of Participating Policies in a Complete Market Framework,” *North American Actuarial Journal*, 14(1), 131–149.
- BERNARD, C., AND C. LEMIEUX (2008): “Fast simulation of equity-linked life insurance contracts with a surrender option,” in *Proceedings of the 40th Conference on Winter Simulation*, pp. 444–452. Winter Simulation Conference.
- BERNARD, C., AND W. V. LI (2013): “Pricing and hedging of cliquet options and locally capped contracts,” *SIAM Journal on Financial Mathematics*, 4(1), 353–371.
- BERNARD, C., Y. LIU, N. MACGILLIVRAY, AND J. ZHANG (2013): “Bounds on capital requirements for bivariate risk with given marginals and partial information on the dependence,” *Dependence Modeling*, 1, 37–53.
- BERNARD, C., AND M. LUDKOVSKI (2012): “Impact of counterparty risk on the reinsurance market,” *North American Actuarial Journal*, 16(1), 87–111.
- BERNARD, C., AND A. MACKAY (2015): “Reducing surrender incentives through fee structure in variable annuities,” in *Innovations in Quantitative Risk Management*, pp. 209–223. Springer.
- BERNARD, C., A. MACKAY, AND M. MUEHLBEYER (2014): “Optimal surrender policy for variable annuity guarantees,” *Insurance: Mathematics and Economics*, 55, 116–128.
- BERNARD, C., M. MAJ, AND S. VANDUFFEL (2011): “Improving the Design of Financial Products in a Multidimensional Black-Scholes Market,” *North American Actuarial Journal*, 15(1), 77–96.

- BERNARD, C., AND D. MCLEISH (2016): “Algorithms for Finding Copulas Minimizing Convex Functions of Sums,” *arXiv preprint arXiv:1502.02130*.
- BERNARD, C., F. MORAUX, L. RÜSCHENDORF, AND S. VANDUFFEL (2015): “Optimal Payoffs under State-dependent Preferences,” *Quantitative Finance*, 15(7), 1157–1173.
- BERNARD, C., L. RÜSCHENDORF, AND S. VANDUFFEL (2014): “Optimal claims with fixed payoff structure,” *Journal of Applied Probability*, 51, 175–188.
- (2016): “Value-at-Risk bounds with variance constraints,” *Journal of Risk and Insurance*, *forthcoming*.
- BERNARD, C., L. RÜSCHENDORF, S. VANDUFFEL, AND J. YAO (2016): “How robust is the Value-at-Risk of credit risk portfolios?,” *European Journal of Finance*, *forthcoming*.
- BERNARD, C., AND W. TIAN (2009): “Optimal reinsurance arrangements under tail risk measures,” *Journal of risk & insurance*, 76(3), 709–725.
- (2010): “Optimal insurance policies when insurers implement risk management metrics,” *Geneva Risk & Insurance Review*, 35, 47–80.
- BERNARD, C., AND S. VANDUFFEL (2014a): “Financial bounds for insurance claims,” *Journal of risk & insurance*, 81(1), 27–56.
- (2014b): “Mean–variance optimal portfolios in the presence of a benchmark with applications to fraud detection,” *European Journal of Operational Research*, 234(2), 469–480.
- (2015): “A new approach to assessing model risk in high dimensions,” *Journal of Banking and Finance*, 58, 166–178.
- BERTRAND, P., AND J.-L. PRIGENT (2001): “Portfolio insurance strategies: OBPI versus CPPI,” *Finance*, 26, 5–32.
- (2011): “Omega performance measure and portfolio insurance,” *Journal of Banking & Finance*, 35(7), 1811–1823.
- BIRNBAUM, M., AND J. NAVARRETTE (1998): “Testing descriptive utility theories: Violations of stochastic dominance and cumulative independence,” *Journal of Risk & Uncertainty*, 17, 49–78.
- BJÖRK, T. (2004): *Arbitrage theory in continuous time*. Oxford university press.
- BLACK, F. (1988): *Individual investment and consumption under uncertainty, Portfolio Insurance: A Guide to Dynamic Hedging*. D. L. Luskin (Ed.), New York, John Wiley and Sons, pages 207–225.
- BONDARENKO, O. (2003): “Statistical arbitrage and securities prices,” *Review of Financial Studies*, 16(3), 875–919.
- BREEDEN, D. T., AND R. H. LITZENBERGER (1978): “Prices of state-contingent claims implicit in option prices,” *Journal of business*, pp. 621–651.
- BROCKETT, P. L., AND Y. KAHANE (1992): “Risk, return, skewness and preference,” *Management Science*, 38(6), 851–866.

- BROCKETT, P. L., M. WANG, AND C. YANG (2005): “Weather derivatives and weather risk management,” *Risk Management and Insurance Review*, 8(1), 127–140.
- BROIHANNE, M.-H., M. MERLI, AND P. ROGER (2008): “Solving some financial puzzles with prospect theory and mental accounting: A survey,” *Revue d’économie politique*, 118(4), 475–512.
- BROWN, D. P., AND J. C. JACKWERTH (2004): *The pricing kernel puzzle: Reconciling index option data and economic theory*. chapter of Contemporary Studies in Economic and Financial Analysis edited by Thornton and Aronson.
- BROWNE, S. (1999): “Beating a moving target: Optimal portfolio strategies for outperforming a stochastic benchmark,” *Finance and Stochastics*, 3(3), 275–294.
- BÜHLMANN, H. (1980): “An economic premium principle,” *Astin Bulletin*, 11(01), 52–60.
- BUSS, A., AND G. VILKOV (2012): “Measuring Equity Risk with Option-implied Correlations,” *Review of Financial Studies*, 25(10), 3113–3140.
- CARLIER, G., AND R.-A. DANA (2011): “Optimal demand for contingent claims when agents have law-invariant utilities,” *Mathematical Finance*, 21(2), 169–201.
- CARMONA, R. (2008): *Indifference Pricing: Theory and Applications*, Princeton Series in Financial Engineering. Princeton University Press.
- CARR, P., AND A. CHOU (1997): “Breaking barriers,” *Risk*, 10(9), 139–145.
- CHABI-YO, F., R. GARCIA, AND E. RENAULT (2008): “State dependence can explain the risk aversion puzzle,” *Review of Financial Studies*, 21(2), 973–1011.
- CHEN, J. S. (2015): “Essays on cost-efficiency,” Ph.D. thesis, Vrije Universiteit Brussel.
- CHEN, Z., K. VETZAL, AND P. A. FORSYTH (2008): “The effect of modelling parameters on the value of GMWB guarantees,” *Insurance: Mathematics and Economics*, 43(1), 165–173.
- CHERUBINI, U., E. LUCIANO, AND W. VECCHIATO (2004): *Copula methods in finance*. John Wiley & Sons.
- CHIU, W. H. (2010): “Skewness preference, risk taking and expected utility maximisation,” *Geneva Risk & Insurance Review*, 35(2), 108–129.
- COCHRANE, J. H. (2005): *Asset pricing*, vol. 1. Princeton university press Princeton, NJ.
- COHEN, M. D. (1995): “Risk-aversion concepts in expected-and non-expected-utility models,” *Geneva Papers on Risk & Insurance Theory*, 20(1), 73–91.
- COLEMAN, T. F., Y. LI, AND M.-C. PATRON (2006): “Hedging guarantees in variable annuities under both equity and interest rate risks,” *Insurance: Mathematics and Economics*, 38(2), 215–228.
- CONT, R., AND R. DEGUEST (2013): “Equity correlations implied by index options: Estimation and model uncertainty analysis,” *Mathematical Finance*, 23(3), 496–530.
- COX, J., AND C. HUANG (1989): “Optimum Consumption and and Portfolio Policies when Asset Prices Follow a Diffusion Process,” *Journal of Economic Theory*, 49(1), 33–83.

- COX, J., AND H. LELAND (1982): “On Dynamic Investment Strategies,” *Proceedings of the seminar on the Analysis of Security Prices, Center for Research in Security Prices, University of Chicago*, 26(2).
- COX, J. C., J. E. INGERSOLL JR, AND S. A. ROSS (1985): “An intertemporal general equilibrium model of asset prices,” *Econometrica: Journal of the Econometric Society*, pp. 363–384.
- COX, J. C., AND H. E. LELAND (2000): “On dynamic investment strategies,” *Journal of Economic Dynamics and Control*, 24(11), 1859–1880.
- CUI, X., J. GAO, X. LI, AND D. LI (2014): “Optimal multi-period mean–variance policy under no-shorting constraint,” *European Journal of Operational Research*, 234(2), 459–468.
- CUI, Z. (2013): “Martingale property and pricing for time-homogeneous diffusion models in finance,” Ph.D. thesis, University of Waterloo.
- DAI, M., Y. KUEN KWOK, AND J. ZONG (2008): “Guaranteed minimum withdrawal benefit in variable annuities,” *Mathematical Finance*, 18(4), 595–611.
- DENUIT, M. M., AND L. EECKHOUDT (2010): “A general index of absolute risk attitude,” *Management Science*, 56(4), 712–715.
- DHAENE, J., M. DENUIT, M. J. GOOVAERTS, R. KAAS, AND D. VYNCKE (2002): “The concept of comonotonicity in actuarial science and finance: theory,” *Insurance: Mathematics and Economics*, 31(1), 3–33.
- DHAENE, J., S. VANDUFFEL, M. GOOVAERTS, R. KAAS, Q. TANG, AND D. VYNCKE (2006): “Risk measures and comonotonicity: a review,” *Stochastic models*, 22(4), 573–606.
- DICKSON, D. C., M. HARDY, M. R. HARDY, AND H. R. WATERS (2013): *Actuarial mathematics for life contingent risks*. Cambridge University Press.
- DRIESSEN, J., P. MAENHOUT, AND G. VILKOV (2012): “Option-Implied Correlations and the Price of Correlation Risk, Working paper,” *Working paper*.
- DRIESSEN, J., P. J. MAENHOUT, AND G. VILKOV (2009): “The price of correlation risk: Evidence from equity options,” *The Journal of Finance*, 64(3), 1377–1406.
- DUFFIE, D. (2010): *Dynamic asset pricing theory*. Princeton University Press.
- DUMAS, B., AND E. LUCIANO (1991): “An exact solution to a dynamic portfolio choice problem under transactions costs,” *The Journal of Finance*, 46(2), 577–595.
- DYBVIG, P. (1988a): “Distributional Analysis of Portfolio Choice,” *Journal of Business*, 61(3), 369–393.
- (1988b): “Inefficient Dynamic Portfolio Strategies or How to Throw Away a Million Dollars in the Stock Market,” *Review of Financial Studies*, 1(1), 67–88.
- DYBVIG, P., AND L. ROGERS (1997): “Recovery of Preferences from Observed Wealth in a Single Realization,” *Review of Financial Studies*, 10, 151–174.
- DYBVIG, P., AND Y. WANG (2012): “Increases in Risk Aversion and the Distribution of Portfolio Payoffs,” *Journal of Economic Theory*, 147(3), 1222–1246.

- EMBRECHTS, P., G. PUCETTI, AND L. RÜSCHENDORF (2013): “Model uncertainty and VaR aggregation,” *Journal of Banking & Finance*, 37(8), 2750–2764.
- FABOZZI, F. J., P. N. KOLM, D. PACHAMANOVA, AND S. M. FOCARDI (2007): *Robust portfolio optimization and management*. John Wiley & Sons.
- FRÉCHET, M. (1960): “Les tableaux dont les marges sont données,” *Trabajos de estadística*, 11(1), 3–18.
- FRIEDMAN, M., AND L. SAVAGE (1948): “The utility analysis of choices involving risk,” *Journal of Political Economy*, 56(4), 279–304.
- GAILLARDETZ, P., AND J. Y. LAKHMIRI (2011): “A New Premium Principle for Equity-Indexed Annuities,” *Journal of risk & insurance*, 78(1), 245–265.
- GHOSSOUB, M. (2011): “Contracting under heterogeneous beliefs,” Ph.D. thesis, University of Waterloo.
- GOETZMANN, W., J. INGERSOLL, M. SPIEGEL, AND I. WELCH (2007): “Portfolio performance manipulation and manipulation-proof performance measures,” *Review of Financial Studies*, 20(5), 1503–1546.
- GOETZMANN, W., J. INGERSOLL, M. I. SPIEGEL, AND I. WELCH (2002): “Sharpening Sharpe ratios,” Discussion paper, National bureau of economic research.
- GOLDFARB, D., AND G. IYENGAR (2003): “Robust portfolio selection problems,” *Mathematics of Operations Research*, 28(1), 1–38.
- GOLDSTEIN, D., E. JOHNSON, AND W. SHARPE (2008): “A new way to measure consumer risk preferences,” *Working Paper*.
- GOLDSTEIN, D. G., E. J. JOHNSON, AND W. F. SHARPE (2008): “Choosing outcomes versus choosing products: Consumer-focused retirement investment advice,” *Journal of Consumer Research*, 35(3), 440–456.
- GRAF, S., A. KLING, AND J. RUSS (2011): “Risk analysis and valuation of life insurance contracts: Combining actuarial and financial approaches,” *Insurance: Mathematics and Economics*, 49(1), 115–125.
- GREEN, R. C., AND S. SRIVASTAVA (1986): “Expected utility maximization and demand behavior,” *Journal of Economic Theory*, 38(2), 313–323.
- GRITH, M., W. KARL HÄRDLE, AND V. KRÄTSCHEMER (2013): “Reference dependent preferences and the EPK puzzle,” Discussion paper, SFB 649 Discussion Paper.
- HARRISON, J. M., AND D. M. KREPS (1979): “Martingales and arbitrage in multiperiod securities markets,” *Journal of Economic theory*, 20(3), 381–408.
- HE, X., AND X. Y. ZHOU (2011): “Portfolio choice via quantiles,” *Mathematical Finance*, 21(2), 203–231.
- HENDERSON, V., AND D. HOBSON (2004): “Utility indifference pricing-an overview,” *Volume on Indifference Pricing*.

- HENS, T., AND C. REICHLIN (2013): “Three solutions to the pricing kernel puzzle,” *Review of Finance*, 17(3), 1065–1098.
- HOBSON, D. (2005): “Bounds for the utility-indifference prices of non-traded assets in incomplete markets,” *Decisions in Economics & Finance*, 28(1), 33–52.
- HODGES, S. D., AND A. NEUBERGER (1989): “Optimal replication of contingent claims under transaction costs,” *Review of futures markets*, 8(2), 222–239.
- HOEFFDING, W. (1963): “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, 58(301), 13–30.
- HOMM, U., AND C. PIGORSCH (2012): “Beyond the Sharpe ratio: An application of the Aumann–Serrano index to performance measurement,” *Journal of Banking & Finance*, 36(8), 2274–2284.
- IBRAGIMOV, R., D. JAFFEE, AND J. WALDEN (2009): “Diversification disasters,” *Journal of Financial Economics*, 99(2), 333–348.
- JIN, H., AND X. YU ZHOU (2008): “Behavioral portfolio selection in continuous time,” *Mathematical Finance*, 18(3), 385–426.
- JORION, P. (2007): *Value at risk: the new benchmark for managing financial risk*, vol. 3. McGraw-Hill New York.
- JOUINI, E., AND H. KALLAL (2001): “Efficient trading strategies in the presence of market frictions,” *Review of Financial Studies*, 14(2), 343–369.
- JOURDAIN, B., AND M. SBAI (2012): “Coupling index and stocks,” *Quantitative Finance*, 12(5), 805–818.
- KAHANE, Y., AND D. NYE (1975): “A portfolio approach to the property-liability insurance industry,” *Journal of risk & insurance*, pp. 579–598.
- KAHNEMAN, D., AND A. TVERSKY (1979): “Prospect Theory: An Analysis of Decision Under Risk,” *Econometrica*, 47(2), 263–291.
- KAPLAN, P. D., AND J. A. KNOWLES (2004): “Kappa: a generalized downside risk-adjusted performance measure,” *Journal of Performance Measurement*, 8, 42–54.
- KARATZAS, I., J. P. LEHOCZKY, AND S. E. SHREVE (1987): “Optimal portfolio and consumption decisions for a small investor on a finite horizon,” *SIAM journal on control and optimization*, 25(6), 1557–1586.
- KEATING, C., AND W. F. SHADWICK (2002a): “An introduction to omega,” *AIMA Newsletter*.
- (2002b): “A universal performance measure,” *Journal of performance measurement*, 6(3), 59–84.
- KELLY, B., AND H. JIANG (2014): “Tail risk and asset prices,” *Review of Financial Studies*, p. hhu039.
- KEMNA, A., AND A. VORST (1990): “A pricing method for options based on average asset values,” *Journal of Banking & Finance*, 14(1), 113–129.

- KORN, R., AND S. TRAUTMANN (1995): “Continuous-time portfolio optimization under terminal wealth constraints,” *Zeitschrift für Operations Research*, 42(1), 69–92.
- KRAUS, A., AND R. H. LITZENBERGER (1976): “Skewness preference and the valuation of risk assets*,” *Journal of Finance*, 31(4), 1085–1100.
- KUMAR, A. (2009): “Who Gambles in the Stock Market?,” *Journal of Finance*, 64(4), 1889–1933.
- LEUNG, P.-L., H.-Y. NG, AND W.-K. WONG (2012): “An improved estimation to make Markowitzs portfolio optimization theory users friendly and estimation accurate with application on the U.S. stock market investment,” *European Journal of Operational Research*, 222(1), 85–95.
- LEVY, H. (2006): *Stochastic Dominance. Investment Decision Making under Uncertainty*. Springer, 2nd edition.
- (2008): “First degree stochastic dominance violations: decision weights and bounded rationality,” *Economic Journal*, 118, 759–774.
- LIM, A. E. (2004): “Quadratic hedging and mean-variance portfolio selection with random parameters in an incomplete market,” *Mathematics of Operations Research*, 29(1), 132–161.
- (2005): “Mean-variance hedging when there are jumps,” *SIAM Journal on Control and Optimization*, 44(5), 1893–1922.
- LO, A. W. (2008): “Where Do Alphas Come from?: A New Measure of the Value of Active Investment Management,” *Journal Of Investment Management (JOIM)*, Third Quarter.
- LONG, J. (1990): “The numeraire portfolio,” *Journal of Financial Economics*, 26, 29–69.
- LOPES, L. (1987): “Between hope and fear: The psychology of risk.,” *Advances in Experimental Social Psychology*, 20, 255–295.
- LOPES, L., AND G. ODEN (1999): “The role of aspiration level in risky choice: A comparison of Cumulative Prospect Theory and SP/A theory.,” *Journal of Mathematical Psychology*, 43, 286–313.
- LUENBERGER, D. (1997): *Investment Science*. Oxford University Press, New York.
- MACHINA, M. J. (1982): “” Expected Utility” Analysis without the Independence Axiom,” *Econometrica: Journal of the Econometric Society*, pp. 277–323.
- MACKEY, A. (2014): “Fee Structure and Surrender Incentives in Variable Annuities,” Ph.D. thesis, University of Waterloo.
- MACKEY, A., M. AUGUSTYNIAK, C. BERNARD, AND M. R. HARDY (2016): “Risk Management of Policyholder Behavior in Equity-Linked Life Insurance,” *Journal of Risk and Insurance*.
- MANOLIU, M., AND S. TOMPAIDIS (2002): “Energy futures prices: term structure models with Kalman filter estimation,” *Applied mathematical finance*, 9(1), 21–43.
- MARKOWITZ, H. (1952): “Portfolio Selection,” *Journal of Finance*, 7(1), 77–91.
- MARTELLINI, L., AND B. UROŠEVIC (2006): “Static mean-variance analysis with uncertain time horizon,” *Management Science*, 52(6), 955–964.

- MERTON, R. (1971): “Optimum Consumption and Portfolio Rules in a Continuous-Time Model,” *Journal of Economic Theory*, 3, 373–413.
- MERTON, R. C. (1972): “An analytic derivation of the efficient portfolio frontier,” *Journal of Financial and Quantitative Analysis*, 7(04), 1851–1872.
- MILEVSKY, M. A., AND T. S. SALISBURY (2006): “Financial valuation of guaranteed minimum withdrawal benefits,” *Insurance: Mathematics and Economics*, 38(1), 21–38.
- MILTERSEN, K. R., AND E. S. SCHWARTZ (1998): “Pricing of options on commodity futures with stochastic term structures of convenience yields and interest rates,” *Journal of Financial and Quantitative Analysis*, 33(01), 33–59.
- MITTON, T., AND K. VORKINK (2007): “Equilibrium underdiversification and the preference for skewness,” *Review of Financial Studies*, 20(4), 1255–1288.
- MØLLER, T. (2003a): “Indifference pricing of insurance contracts in a product space model,” *Finance and Stochastics*, 7(2), 197–217.
- (2003b): “Indifference pricing of insurance contracts in a product space model: applications,” *Insurance: Mathematics and Economics*, 32(2), 295–315.
- MORGENSTERN, O., AND J. VON NEUMANN (1953): *Theory of games and economic behavior*. Princeton University Press.
- MOSSIN, J. (1968): “Optimal multiperiod portfolio policies,” *Journal of Business*, pp. 215–229.
- MÜLLER, A., AND D. STOYAN (2002): *Comparison methods for stochastic models and risks*, vol. 389. Wiley.
- NELSEN, R. (2006): *An Introduction to Copulas*. Second edition, Springer.
- PELEG, B., AND M. YAARI (1975): “A Price Characterization of Efficient Random Variables,” *Econometrica*, 43(2), 283–292.
- PELSSER, A., AND M. STADJE (2014): “time-consistent and market-consistent evaluations,” *Mathematical Finance*, 24(1), 25–65.
- PELSSER, A., AND T. VORST (1996): “Transaction costs and efficiency of portfolio strategies,” *European Journal of Operational Research*, 91(2), 250–263.
- PHAM, H. (2000): “On quadratic hedging in continuous time,” *Mathematical Methods of Operations Research*, 51(2), 315–339.
- PLATEN, E., AND D. HEATH (2006): *A benchmark approach to quantitative finance*, no. v. 13 in Springer finance. Springer.
- POJARLIEV, M., AND R. M. LEVICH (2008): “Do professional currency managers beat the benchmark?,” *Financial Analysts Journal*, 64(5), 18–32.
- POLISSON, M., AND J. K.-H. QUAH (2013): “Revealed preference tests under risk and uncertainty,” Discussion Papers in Economics 13/24, Department of Economics, University of Leicester.
- POLKOVNICHENKO, V. (2005): “Household Portfolio Diversification: A Case for Rank-Dependent Preferences,” *Review of Financial Studies*, 18(4), 1467–1502.

- PRÉKOPA, A. (1973): “On logarithmic concave measures and functions,” *Acta Scientiarum Mathematicarum*, 34, 335–343.
- PRIGENT, J.-L. (2007): *Portfolio optimization and performance analysis*. CRC Press.
- PUC CETTI, G., AND L. RÜSCHENDORF (2012): “Computation of sharp bounds on the distribution of a function of dependent risks,” *Journal of Computational and Applied Mathematics*, 236(7), 1833–1840.
- QUIGGIN, J. (1993): *Generalized Expected Utility Theory - The Rank-Dependent Model*. Kluwer Academic Publishers.
- REICHLIN, C. (2013): “Utility maximization with a given pricing measure when the utility is not necessarily concave,” *Mathematics and Financial Economics*, 7, 531–556.
- RIEGER, M. (2010): “SP/A and CPT: A reconciliation of two behavioral decision theories,” *Economics Letters*, 108(1), 327–329.
- (2011): “Co-monotonicity of optimal investments under the design of structural financial products,” *Finance and Stochastics*, 15(1), 27–55.
- ROY, A. (1952): “Safety First and the Holding of Assets,” *Econometrica*, 20(30), 431–449.
- RYAN, M. J. (2006): “Risk aversion in RDEU,” *Journal of Mathematical Economics*, 42(6), 675–697.
- SCHWEIZER, M. (1992): “Mean-variance hedging for general claims,” *The annals of applied probability*, pp. 171–179.
- (2010): “Mean–variance hedging,” *Encyclopedia of Quantitative Finance*.
- SHAKED, M., AND J. SHANTIKUMAR (2007): *Stochastic Orders*. Springer, New York.
- SHARPE, W. (2007): “Expected Utility Asset Allocation,” *Financial Analysts Journal*, 63(5), 18–30.
- SHARPE, W. F. (1967): “Portfolio analysis,” *Journal of Financial and Quantitative Analysis*, 2(02), 76–84.
- (1994): “The Sharpe ratio,” *The journal of portfolio management*, 21(1), 49–58.
- SHEFRIN, H. (2008): *A behavioral approach to asset pricing*, no. chapters 26 and 27 in Academic Press Advanced Finance Series, 2nd edition. Elsevier.
- SHEFRIN, H., AND M. STATMAN (2000): “Behavioral Portfolio Theory,” *Journal of Financial and Quantitative Analysis*, 35(2), 127–151.
- SORTINO, F. A., AND L. N. PRICE (1994): “Performance measurement in a downside risk framework,” *the Journal of Investing*, 3(3), 59–64.
- SPIVAK, G., AND J. CVITANIĆ (1999): “Maximizing the probability of a perfect hedge,” *Annals of Applied Probability*, pp. 1303–1328.
- STARMER, C. (2000): “Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk,” *Journal of Economic Literature*, 38(2), 332–382.

- TAKAHASHI, A., AND K. YAMAMOTO (2013): “Generating a target payoff distribution with the cheapest dynamic portfolio: an application to hedge fund replication,” *Quantitative Finance*, 13(10), 1559–1573.
- TANKOV, P. (2011): “Improved Fréchet bounds and model-free pricing of multi-asset options,” *Journal of Applied Probability*, 48(2), 389–403.
- TIAN, Y. (2013): “Ironing out the Wrinkles in Executive Compensation: Linking Incentive Pay to Average Stock Prices,” *Journal of Banking and Finance*, 37, 415–432.
- TÜTÜNCÜ, R. H., AND M. KOENIG (2004): “Robust asset allocation,” *Annals of Operations Research*, 132(1-4), 157–187.
- TVERSKY, A., AND D. KAHNEMAN (1992): “Advances in Prospect Theory: Cumulative Representation of Uncertainty,” *Journal of Risk & Uncertainty*, 5(4), 297–323.
- VANDUFFEL, S., A. AHCAN, L. HENRARD, AND M. MAJ (2012): “An explicit option-based strategy that outperforms dollar cost averaging,” *International Journal of Theoretical and Applied Finance*, 15(02).
- VANDUFFEL, S., A. CHERNIH, M. MAJ, AND W. SCHOUTENS (2009): “A note on the suboptimality of path-dependent payoffs in Lévy markets,” *Applied Mathematical Finance*, 16(4), 315–330.
- VON NEUMANN, J., AND O. MORGENSTERN (1947): *Theory of games and economic behavior*. Princeton University Press.
- WAGNER, W. (2009): “Diversification disasters,” *Diversification at financial institutions and systemic crises*, 19(3), 373–386.
- WANG, B., AND R. WANG (2011): “The complete mixability and convex minimization problems with monotone marginal densities,” *Journal of Multivariate Analysis*, 102(10), 1344–1360.
- (2016): “Joint mixability,” *Mathematics of Operations Research*, forthcoming.
- WANG, J., AND P. FORSYTH (2011): “Continuous time mean variance asset allocation: A time-consistent strategy,” *European Journal of Operational Research*, 209(2), 184–201.
- WANG, S. S., V. R. YOUNG, AND H. H. PANJER (1997): “Axiomatic characterization of insurance prices,” *Insurance: Mathematics and Economics*, 21(2), 173–183.
- WÜTHRICH, M. V., H. BÜHLMANN, AND H. FURRER (2010): *Market-consistent actuarial valuation*. Springer Science & Business Media.
- XIA, J., AND J.-A. YAN (2006): “Markowitz’s portfolio optimization in an incomplete market,” *Mathematical Finance*, 16(1), 203–216.
- YAARI, M. (1987): “The Dual Theory of Choice under Risk,” *Econometrica*, 55(1), 95–115.
- YOUNG, V. R. (2004): “Premium principles,” *Encyclopedia of Actuarial Science*.
- ZHANG, W.-G., X.-L. ZHANG, AND W.-L. XIAO (2009): “Portfolio selection under possibilistic mean–variance utility and a SMO algorithm,” *European Journal of Operational Research*, 197(2), 693–700.