

Optimal selection of a portfolio of options under Value-at-Risk constraints: A scenario approach

Citation for published version (APA):

Schyns, M., Crama, Y., & Hübner, G. (2010). Optimal selection of a portfolio of options under Value-at-Risk constraints: A scenario approach. *Annals of Operations Research*, 181(1), 683-708.
<https://doi.org/10.1007/s10479-009-0636-y>

Document status and date:

Published: 01/12/2010

DOI:

[10.1007/s10479-009-0636-y](https://doi.org/10.1007/s10479-009-0636-y)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Optimal selection of a portfolio of options under Value-at-Risk constraints: a scenario approach

M. Schyns · Y. Crama · G. Hübner

Published online: 20 October 2010
© Springer Science+Business Media, LLC 2010

Abstract This paper introduces a multiperiod model for the optimal selection of a financial portfolio of options linked to a single index. The objective of the model is to maximize the expected return of the portfolio under constraints limiting its Value-at-Risk. We rely on scenarios to represent future security prices. The model contains several interesting features, like the consideration of transaction costs, bid-ask spreads, arbitrage-free option pricing, and the possibility to rebalance the portfolio with options introduced at the start of each period. The resulting mixed integer programming model is applied to realistic test instances involving options on the S&P500 index. In spite of the large size and of the numerical difficulty of this model, near-optimal solutions can be computed by a standard branch-and-cut solver or by a specialized heuristic. The structure and the financial features of the selected portfolios are also investigated.

1 Introduction

1.1 General statement of the problem

The main purpose of this work is to develop a multiperiod model for the selection of an optimal portfolio of options subject to Value-at-Risk constraints, to demonstrate the computational feasibility of the approach, and to investigate some of the properties of the resulting portfolios.

M. Schyns · Y. Crama (✉) · G. Hübner
HEC Management School, University of Liège, Rue Louvrex 14 (N1), 4000 Liège, Belgium
e-mail: Yves.Crama@ulg.ac.be

M. Schyns
e-mail: M.Schyns@ulg.ac.be

G. Hübner
School of Business and Economics, Maastricht University, Tongersestraat 53, 6211 LM Maastricht,
The Netherlands
e-mail: G.Hubner@ulg.ac.be

G. Hübner
Gambit Financial Solutions, Rue Forgeur 17, 4000 Liège, Belgium

In general terms, we are interested in the following situation. A fund manager considers investing a total budget B into a portfolio of options written over a single underlying asset (for instance, options over the S&P index) or in the riskless asset. He must make an initial investment decision at the beginning of the horizon, and is allowed to rebalance the portfolio at the end of a first intermediate period. The problem he faces is to select the quantities to be invested in each option so as to optimize the expected value of the portfolio at the end of the horizon, while satisfying predetermined Value-at-Risk constraints at the end of each subperiod.

Although this optimization problem is rather easily stated, modelling it in a rigorous and meaningful way turns out to be quite challenging. Before we discuss the formal details of the model, let us first justify our interest in the problem.

1.2 Motivation

Value-at-Risk (VaR) and its variants, like Conditional Value-at-Risk (CVaR), are very popular concepts for measuring the risk associated with a portfolio of securities. Extensive research has been performed on this subject, with much of the emphasis being placed on the CVaR model (see e.g. Krokmal et al. 2002). The CVaR is attractive for two main reasons: it displays nice financial properties—it is a coherent measure of risk as defined by Artzner et al. (1999), and it is easier to compute than the VaR, which uses a quantile of the portfolio returns distribution. But since the adoption of the Comprehensive Basel II Accord by the Bank for International Settlements (BIS 2006), the role of the VaR has become more central for regulatory purposes. Financial institutions of BIS member countries have to comply strictly with VaR requirements regarding credit risk, operational risk and market risk management. Thus, limitations in VaR exposures typically represent binding constraints for active portfolio management; see e.g. Esch et al. (2005).

The extant literature on portfolio management with VaR requirements has mostly focused on *estimating* the VaR associated with a *predefined* portfolio. However, this approach does not explicitly address the central problem of asset managers, whose objective is to account directly for VaR in the *selection* of optimal complex portfolios.

We restrict here our attention to portfolios of options linked to a single financial asset. Considering only one underlying asset allows us to focus on the core of the problem, namely, portfolio selection under a VaR constraint. Even with only one underlying asset and several linked options, the problem remains complex since the VaR cannot be computed a priori without knowing the probability distribution of the portfolio returns, since the distribution of the portfolio returns cannot be specified before selecting the optimal quantities of its components (and before specifying their respective probability distributions), and since the optimal quantities must be set so as to satisfy the VaR requirement for the whole portfolio. This problem is therefore far more complex than computing the VaR associated with a *predefined* portfolio.

1.3 Contributions

The inherent complexity of the model and of the optimization process justifies our use of a methodology based on *multinomial tree of scenarios*. More precisely, the formulation of the model relies on a collection of scenarios which provide a representative sample of values for the returns of the underlying asset. Since the values of the options are indirectly determined by the scenarios, this approach allows us to capture completely the random nature of the problem.

In order to model the distribution of the underlying asset returns and to generate the associated scenarios, we do not restrict ourselves to standard parametric statistical distributions (like the normal or the Student distribution), but we propose to sample from more general distributions of returns. This approach has the advantage of providing great flexibility in the construction of the model, allowing in particular to integrate different types of distributions and various realistic constraints. We illustrate this with the computation of a distribution of returns implied by the prices of the options available on the market at the time of the investment.

All in all, our work appears to be unique in its simultaneous consideration of multiperiod scenarios, of options, and of a broad range of realistic financial constraints including the VaR measure. In the literature, these three features have been mostly developed in isolation, and so cannot provide reasonable insights into realistic portfolio selection problems. In particular, we pay special attention to the calibration of the parameters defining the tree of scenarios, an issue which turns out to be quite tricky, but extremely relevant when dealing with options.

Our first goal is to show that this approach is valid and tractable in practice. But, thanks to our integrative framework, we can also go a step further, and investigate the composition and the behavior of an optimal portfolio involving options under VaR requirements. We first study the shape of the payoff function and the set of options required to reach the investor's goal. We provide comparisons with the CVaR measure. We also analyze the impact on the VaR of the introduction of a risk-free asset in the options portfolio. This allows us to measure a risk premium induced by the use of options alone. These results are illustrated on a case-study using a set of put and call options written on the S&P500 index.

The results displayed in this paper indicate that, in line with conventional wisdom, the optimal portfolio strategy actually activates a limited set of options, and shares the qualitative features of a protective put. These findings, based on a rigorous optimization approach, contribute to reconciling complex numerical methods with the actual practice of experienced portfolio managers.

The paper is organized as follows. The next section presents the framework of the portfolio management problem with an emphasis on its main financial features and on trees of scenarios. Section 3 gives a complete description of the optimization model to be solved, which is formulated here as a mixed integer linear programming problem, and Sect. 4 discusses the difficulties that arise when instantiating the set of scenarios. In Sect. 5, we report on various numerical experiments with our optimization model and we analyze the results, both in terms of computational efficiency and from the point of view of their financial interpretation. Finally, the last section draws the main conclusions of our work and presents some perspectives for future research.

2 Framework

Before presenting a mathematical formulation of the optimization model to be solved, we first discuss its various components.

2.1 Risk measure

The Value-at-Risk (VaR) of a portfolio at level P is defined as the maximal loss of the portfolio value with probability P , over a specific horizon:

$$\text{VaR} = \min\{V \mid \text{Prob}[\text{loss} \leq V] \geq P\}.$$

If the probability distribution function of losses is continuous and strictly increasing, then the VaR can be expressed more simply by the condition:

$$\text{Prob}[loss \leq \text{VaR}] = P.$$

This risk measure, initially proposed by Edgeworth (1888), has become popular when introduced by JP Morgan in RiskMetrics (1996). It has subsequently been proposed in the “1996 Amendment to the Basel I Capital Accord”, included in the Comprehensive Basel II Accord and is fully applicable nowadays; see e.g. Esch et al. (2005).

In addition to Value-at-Risk requirements, risk management systems often impose “stop-loss” procedures in order to limit the extent of the losses incurred on an individual position. This is equivalent to setting a guaranteed amount at the end of the investment horizon.¹ The same requirement is met for insured portfolios with the Option-Based Portfolio Insurance (OBPI) technique or the Constant Proportion Portfolio Insurance technique (CPPI) proposed by Black and Jones (1987), as the portfolio becomes entirely invested in risk-free securities when the loss incurred in the risky part reaches a given level (see Bertrand and Prigent 2005 for a discussion).

2.2 Market

Many professionally managed equity portfolios involve optional securities. They allow traders or fund managers to shape the future payoff of their portfolios, for instance to ensure a floor in the terminal payoff, which is easily achieved thanks to the intrinsic properties of options. We assume that the manager’s main goal is to maximize the expected value of the portfolio at the end of the investment horizon. To track real market conditions as closely as possible, we take into account the bid-ask spreads and the transaction costs.

We have restricted this first investigation to the case of options linked to a same financial index. Models with several securities and their corresponding derivatives could also be handled by a similar approach but they require more parameters (see Schyns et al. 2009 for an exploratory study). With one underlying index, clear conclusions are easier to draw.

Our model is specially intended for large investments and this allows us to formulate some simplifying hypotheses. In particular, if the initial budget is large, then the fixed commission cost is small with respect to the total invested amount and it can be neglected in the model. Also, since the amount invested in each option can be assumed to be large, we can consider only one proportional tax rate, namely, the rate which applies to the largest trading amounts.

2.3 Multiperiod model

Due to various reasons (transaction costs, market opening times, etc.), it is usually impractical to trade continuously. A discrete multiperiod model is therefore a suitable compromise. In this paper, we restrict our attention to a two-period model. The first aim of the two-period

¹Note that the guarantee constraint is a special case of the Value-at-Risk constraint where the probability P is set to 100%. We could actually impose several VaR constraints in our model. The methodology would be similar for each of them, even if some simplifications could take place (indeed, several VaR constraints would imply overlapping lower bounds). Alternatively, the manager could be imposed a limit on the Conditional Value-at-Risk (a.k.a. Expected Shortfall) beyond the VaR level. This requirement limits the extent of the expected losses when a disaster occurs. We discuss the implications of this requirement in Sect. 5.

approach is to model the degree of freedom given to the portfolio manager when he is offered the opportunity to dynamically rebalance his positions. We define the optimization model so as to provide a global optimal solution over the two periods.

In practice, the model could be used to optimize an investment over a fixed horizon of two periods, such as the position taken by a front-office equity trader or, with a longer time frame, by many closed-end funds. But it could also apply to an investor with a moving horizon, like an open-end mutual fund. Such an investor would certainly not rebalance the portfolio at the end of the first period according to the optimal solution found initially. Rather, a roll-over strategy would be adopted and a new optimization would be carried out at the end of the first period, taking into account all available information as well as a new investment horizon.

Finally, we impose the requirement that the portfolio be self-financing throughout the lifetime of the investment: there is no additional cash-flow at the end of the first period. This assumption naturally derives from the fact that we are considering a fixed horizon and that the portfolio is closed at inception of the strategy. The budget constraint is a natural requirement for this type of strategy. Extensions are however straightforward.

2.4 Multiperiod trees of scenarios

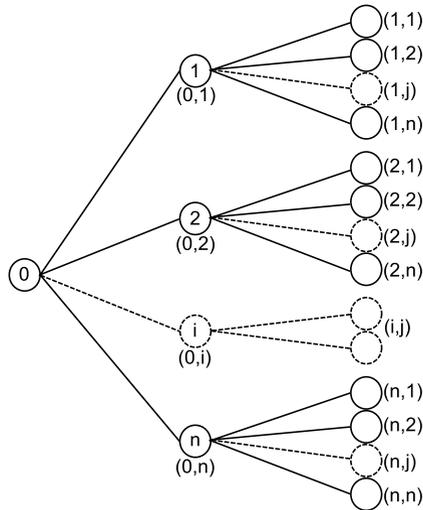
Trees of scenarios constitute a generic, relatively simple approach to represent future states of the world in stochastic optimization problems (see e.g. Birge and Louveaux 1999 or Prekopa 1995 for a broad introduction to stochastic programming). In finance, such trees have been used in numerous models, both in computational and in theoretical frameworks, as in Dembo (1991), Dybvig (1988, 1988), Gülpinar et al. (2004), Kallio and Ziemba (2007), King (2002), Kouwenberg (2001), Larsen et al. (2002), Mulvey (1994), Muzzioli and Torricelli (2005), Naik (1995), Rockafellar and Uryasev (2000), Rubinstein (1994, 1998), etc.

Grinold (1999) ponders some of the relative advantages and drawbacks of scenario-based approaches vs. mean-variance approaches when dealing with portfolio optimization problems. In Grinold's view, scenario-based models are mostly useful—and even indispensable—for portfolio management problems involving options or assets with alternative distributions. They deal with the entire distribution of outcomes and thereby allow for a broad variety of objectives. However, they also have to respond to several major challenges. In particular, setting up a tree of scenarios requires the solution of several complex numerical problems, including the specification of the entire distribution of all assets.

In this paper, we consider scenario trees where each node (either internal or terminal) represents a possible state of the world at a particular date; i.e., each node is characterized by the sequence of asset values along the unique path from the root of the tree to the node under consideration. The intrinsic quality of the tree depends on the process used to instantiate each node and on the number of nodes. Without loss of generality, we assume that all scenarios ending at a same date are equiprobable; they constitute a *representative sample of asset values* at that date, which can be obtained by sampling from the probability distribution function of the asset values (see Sect. 4.1). These values can be used, in turn, to price the options linked to the underlying. (It would be straightforward to extend the model to the case where the scenarios are not equiprobable.)

As explained in Sect. 2.3, we focus on the generation of two-period trees (see Fig. 1). Thus, we consider a dynamic roll-over strategy of investment, whereby the investor can rebalance his portfolio at some intermediate time instant before the end of the horizon. The first period is associated with the initial period of investment and each subtree of the second period corresponds to a portfolio rebalanced by the investor according to the state of the world at the end of the first period.

Fig. 1 Two-period tree of scenarios



2.5 Options

All options considered here are European options with a one-period maturity, i.e., they are created at the beginning of a subperiod, and they expire at the end of the same subperiod.² For the first period, the initial characteristics of traded options are directly observable on the market. For the second period, the set of available options is not initially known, as it depends on the fluctuations of the underlying index during the first period. A major advantage of multinomial trees of scenarios is precisely their ability to represent such dependent outcomes. A realistic set of options can therefore be introduced separately for each subtree of the second period by setting strike prices around the index value observed at the root node of the subtree, according to usual market conventions. We return to this issue in Sect. 4.2.

2.6 Constraints

Note finally that the same guarantee and VaR constraints will be imposed for each subtree of the second period; thus, whatever scenario materializes during the first period, the insurance required by the investor must be satisfied in the corresponding subtree of the second period. We also fix a guarantee constraint to be satisfied at the end of the first period. The level of this constraint can be different and is typically lower than for the second period. This allows the portfolio exposures to remain tightly under control throughout the investment horizon.³

3 Model

Informally, the manager's problem is to select a portfolio with maximum expected value at the end of the horizon, under the following constraints:

²Extensions to other maturities could be included in the model, but they complicate its formulation as well as the interpretation of the results.

³A VaR constraint could also be imposed at the end of the first period, but is disregarded here in order to simplify the model, and because the guarantee constraint is deemed sufficiently protective.

- Budget: the initial cost of settlement does not exceed the available budget. Moreover, at the beginning of the second period, the investor cannot invest more than the value of the current portfolio.
- Guarantee: the value of the portfolio at the end of each period cannot be less than a predefined fraction of the initial budget, under no circumstances.
- Value-at-Risk: with a predefined probability P , the final value of the portfolio cannot be less than a predefined fraction of the initial budget.

In this section, we turn to a mathematical formulation of this portfolio optimization problem.

3.1 Notations

We consider a two-period tree as represented in Fig. 1. We denote by $i = 0, 1, \dots, n$ the non-terminal nodes of the tree, with root-node 0 corresponding to the initial state. For $i = 0, 1, \dots, n$, each node i is the root of a subtree (also indexed by i) consisting of n branches indexed by $j = 1, \dots, n$, and of n (immediate) successor nodes of i associated with n distinct (partial) scenarios: each such scenario is defined by the sequence of asset values along the unique path from node 0 to the end-node of the corresponding branch. Alternatively, it is sometime handy to denote a node (including a terminal node) by a pair (i, j) , indicating that it is the j -th successor of node i in the subtree rooted at i ($i = 0, 1, \dots, n; j = 1, \dots, n$). The period under consideration in subtree i is denoted by t , where $t = 1$ in the subtree $i = 0$, and $t = 2$ for all other subtrees.

Remember that we consider a single underlying asset. At each node i , O_i options are available on this asset ($i = 0, 1, \dots, n$). For $i = 0, 1, \dots, n, j = 1, \dots, n, k = 1, \dots, O_i$, and $t = 1, 2$, the parameters of the model are

- B : the available initial budget;
- $S_{i,j}$: the price of the underlying asset at node (i, j) ;
- $S_{i,0}$: the price of the underlying asset at the root of subtree i (note that $S_{i,0} = S_{0,i}$ when $i \geq 1$);
- $K_{i,k}$: the strike price of option k for all immediate successor nodes of node i ;
- $p_{i,k}^a$: the ask price of option k for all immediate successor nodes of node i ;
- $p_{i,k}^b$: the bid price of option k for all immediate successor nodes of node i ;
- T : the taxation rate;
- r_t : the risk-free rate for period t ;
- ℓ_t : the length of period t ;
- θ_t : the percentage of the initial budget to be preserved at the end of period t in order to respect the guarantee constraint;
- λ : the percentage of the initial budget to be preserved with probability P at the end of period 2 in order to respect the VaR constraint.

Some of the main decision variables are

- $x_{i,k}^a$: the (positive) quantity of option k bought at node i (long position);
- $x_{i,k}^b$: the (positive) quantity of option k sold at node i (short position);
- z_i : the amount invested (positive amount) or borrowed (negative amount) in the risk-free asset at node i ;
- $\alpha_{i,j}$: a $\{0, 1\}$ variable which takes value 1 only if the value of the portfolio exceeds the VaR threshold λB in the j -th successor of node i .

We assume without loss of generality that the portfolio is fully recomposed at the beginning of each subperiod, that is, at the root node i of each subtree, taking into account the securities available at the corresponding date. Thus, the variables $x_{i,k}^a, x_{i,k}^b$ and z_i completely describe the portfolio selected at node i (they are not incrementally defined with respect to the previous period).

For simplicity of notations, we will also assume from now on that the set of options under consideration at each node i is ordered by type (calls first, puts last) and by increasing strike price.

More parameters and variables will be introduced as they are needed, in particular for a definition of the option pricing model in Sect. 4.2.

3.2 Portfolio value

The portfolio value can be easily determined at each node of the tree if we know the quantities invested in each security. Indeed, the portfolio value at a node is essentially the sum of the securities values weighted by the invested quantities. By construction of the tree, the value of each security at each node is known (see Sect. 4 for details). This general scheme just has to be slightly adapted to integrate the transactions costs, i.e., a given percentage of the option values, and the bid, ask or maturity prices.

At maturity, the value of an option is given by:

$$vopt(S, K) = \begin{cases} S - K & \text{for a call option with } S > K, \\ K - S & \text{for a put option with } K > S, \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

where S is the underlying asset price and K is the option strike price.

A distinct portfolio is to be selected in each subtree $i = 0, 1, \dots, n$. The cost incurred in order to compose a portfolio at node i is given by:

$$V_i^0 = \sum_{k=1}^{O_i} [p_{i,k}^a x_{i,k}^a (1 + T) - p_{i,k}^b x_{i,k}^b (1 - T)] + z_i. \tag{2}$$

Assuming that the price of the underlying asset is S at the end of the corresponding time period t , then, when the investor closes his positions, the final value of the portfolio composed at node i is:

$$V_i(S) = \sum_{k=1}^{O_i} vopt(S, K_{i,k}) [x_{i,k}^a (1 - T) - x_{i,k}^b (1 + T)] + z_i (1 + r_t)^{\ell_t}. \tag{3}$$

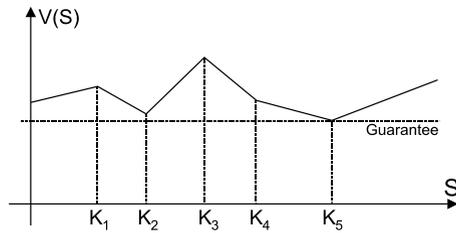
The expected portfolio value is obtained as the sum of the portfolio values at all terminal nodes, multiplied by the probability of the associated scenarios (in our case, this simply amounts to dividing by n^2 , since we consider this number of equiprobable scenarios).

3.3 Risk and investment limits

3.3.1 Guarantee constraints

Recall that θ_1 and θ_2 respectively denote the percentage of the initial budget to be preserved at the end of period 1 and of period 2. In the scenario tree framework, these conditions naturally translate into constraints which impose the minimum guarantee level at every relevant

Fig. 2 Portfolio value as a piecewise linear function of the index value



leaf of the tree, namely:

$$V_0(S_{0,j}) \geq \theta_1 B, \quad j = 1..n; \tag{4}$$

$$V_i(S_{i,j}) \geq \theta_2 B, \quad i = 1..n; \quad j = 1..n. \tag{5}$$

However, an alternative model (initially proposed by Dert and Oldenkamp (2000) and slightly adapted below) is more effective, in that it allows to impose the guarantee constraints for all possible values of the underlying, and not only for the values associated with the scenarios, as in (4)–(5). We now briefly describe this model (see Dert and Oldenkamp 2000 for details).

Since each option is at maturity at the end of one of the subperiods, its payoff function is piecewise linear with only one breakpoint at the corresponding strike price. This implies that the portfolio value, which is a weighted sum of the payoffs of its individual components, also is a piecewise linear function of the underlying index value. Thus, if the guarantee constraint is satisfied at every breakpoint, i.e., at every strike price, then it remains satisfied in the interval between every pair of strike prices; see Fig. 2 for an illustration of this observation. This yields the constraints

$$V_0(K_{0,k}) \geq \theta_1 B \quad k = 1..O_0, \tag{6}$$

$$V_i(K_{i,k}) \geq \theta_2 B \quad i = 1..n; \quad k = 1..O_i. \tag{7}$$

In order to enforce the guarantee constraint for all possible nonnegative values of the index, we must also pay special attention to the extreme values of the index, i.e., when the index value falls under the smallest strike price or rises above the largest one. The minimal possible value of the index is zero (in fact, the index can be viewed as an option with a null strike price). Thus, by adding a virtual “zero” breakpoint, we take care of this first extreme value with the constraints:

$$V_0(0) \geq \theta_1 B, \tag{8}$$

$$V_i(0) \geq \theta_2 B \quad i = 1..n. \tag{9}$$

For the values of the index larger than the largest strike price, since the payoff function is linear, only two cases can occur: either the payoff function is nondecreasing, or it is decreasing. In the first case, the guarantee constraint is always satisfied if it is imposed at the largest strike price. In the second case, the constraint will always be eventually violated, for large enough values of the index. To avoid the latter case, we simply impose that the first derivative of the payoff function must be nonnegative beyond the largest strike price. Fortunately, the first derivative is easy to obtain from (3): it is just a linear combination of

the payoffs of the O_i^c calls (puts disappear from the expression since they are out of the money for large index values). This leads to the constraints:

$$\sum_{k=1}^{O_i^c} [x_{i,k}^a(1 - T) - x_{i,k}^b(1 + T)] \geq 0 \quad i = 0..n. \tag{10}$$

In the formulation (6)–(10), the minimum guarantee level is imposed for *all* nonnegative values of the index, and not only for the values associated with the scenarios (no “leak” can appear between successive values). Another advantage of the formulation (6)–(10) is that the number of constraints to be applied is generally smaller than with the formulation (4)–(5), since the number of options is usually smaller than the number of scenarios.

3.3.2 Value-at-Risk constraint

The VaR constraint is defined as the minimum payoff, say λB , to be achieved at the end of period 2 with probability P . A tree of scenarios is very convenient to model this constraint: if we consider that the set of leaves represents all possible outcomes, then the VaR constraint is satisfied if and only if, for each second-period subtree, the value of the portfolio is greater than or equal to λB in at least $(P \times n)$ scenarios, where n is the number of scenarios in each subtree.

This approach leads to a formulation involving binary variables which express that, for a given leaf, either the threshold λB is reached, or not. More precisely, let us introduce n^2 decision variables $\alpha_{i,j}$ with values in $\{0, 1\}$. Then the VaR requirements can be expressed by the constraints:

$$V_i(S_{i,j}) \geq \alpha_{i,j} \lambda B + (1 - \alpha_{i,j}) \theta_2 B, \quad i = 1..n; \quad j = 1..n; \tag{11}$$

$$\sum_{j=1}^n \alpha_{i,j} \geq Pn, \quad i = 1..n. \tag{12}$$

Due to constraint (11), $\alpha_{i,j}$ can be set to 1 only if the value of the portfolio exceeds the VaR threshold λB at node (i, j) (j -th branch of subtree i); and due to constraint (12), this must happen for at least Pn scenarios. On the other hand, when $\alpha_{i,j} = 0$, constraint (11) is implied by the guarantee constraints as discussed above.

Of course, this formulation suffers from the same weakness as the formulation (4)–(5) of the guarantee constraint (namely, it does not take all possible values of the index into account), and it requires a large number of scenarios to provide an accurate representation of the VaR constraint. But unfortunately, an alternative extended formulation similar to (6)–(10) cannot be obtained for the VaR constraints: the main difficulty being here that the VaR requirement λB must be achieved for a subset of Pn scenarios, but there is no way to know *a priori* the subset of scenarios which will yield this required value.

Note that Dert and Oldenkamp (2000) have suggested a solution to this problem. Using a theorem of Dybvig (1988), they propose to sort *a priori* the scenarios by increasing portfolio values. (In order to do so, in their framework, they only need to know the state-prices associated with the scenarios, but not the portfolio composition.) Then, it is easy to see that the VaR requirement must be fulfilled only at the P -fraction of the leaves associated with the highest values of the portfolio, and this can be expressed in the optimization model.

However, this approach is only valid when three strong assumptions are satisfied: market completeness, market perfectness, normality of returns, and accuracy of the state-prices.

None of these hypotheses are satisfied in our model since we usually consider more scenarios than securities (leading to an incomplete market), transaction costs, departure from normality of returns, and a bid-ask spread for option prices. In particular, because the market is incomplete, the state-prices are not univocally defined by the option pricing process, and since the distribution of the index returns may depart from normality, the state prices are not necessary monotonic with scenario values. We discuss these two issues and our modelling choices in detail in Sects. 4.1 and 4.2.

3.4 The main optimization model

The core of the VaR optimization model can now be expressed by the mixed integer programming (MIP) model **VaROpt**. In this model, ER denotes the expected value of the portfolio at the end of the second period. The initial portfolio cost V_i^0 at node i is given by the expression (2), and the value $V_i(S)$ of the portfolio composed at node i is given by (3) at the end of the subsequent period if the price of the underlying is S at this time. The other parameters, variables and constraints have been defined above.

VaROpt model

$$\max ER = \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n V_i(S_{i,j})$$

subject to

(budget constraints:)

$$\begin{aligned} V_0^0 &\leq B \\ V_i^0 &\leq V_0(S_{0,i}) \quad i = 1..n \end{aligned}$$

(guarantee constraints at the end of period 1:)

$$\begin{aligned} V_0(K_{0,k}) &\geq \theta_1 B \quad k = 1..O_0 \\ V_0(0) &\geq \theta_1 B \\ \sum_{k=1}^{O_c} [x_{0,k}^a(1 - T) - x_{0,k}^b(1 + T)] &\geq 0 \end{aligned}$$

(guarantee constraints at the end of period 2:)

$$\begin{aligned} V_i(K_{i,k}) &\geq \theta_2 B \quad i = 1..n; k = 1..O_i \\ V_i(0) &\geq \theta_2 B \quad i = 1..n \\ \sum_{k=1}^{O_c} [x_{i,k}^a(1 - T) - x_{i,k}^b(1 + T)] &\geq 0 \quad i = 1..n \end{aligned}$$

(VaR constraints at the end of period 2:)

$$V_i(S_{i,j}) \geq \alpha_{i,j} \lambda B + (1 - \alpha_{i,j}) \theta_2 B \quad i = 1..n; j = 1..n$$

(VaR confidence level:)

$$\sum_{j=1}^n \alpha_{i,j} \geq P n \quad i = 1..n$$

$$x_{i,k}^a, x_{i,k}^b \in \mathbb{R}^+ \quad i = 0..n; k = 1..O_i$$

$$z_i \in \mathbb{R} \quad i = 0..n$$

$$\alpha_{i,j} \in \{0, 1\} \quad i = 1..n; j = 1..n$$

3.5 Related work

The **VaROpt** model is a special type of probabilistically constrained stochastic programming model (see Birge and Louveaux 1999; Prekopa 1995). Several related variants of portfolio optimization models subject to VaR (or related) constraints have been investigated in the literature. For instance, Dert and Oldenkamp (2000) propose a single-period model, and they discuss the shape of the optimal portfolio payoff under guarantee or VaR constraints. Daníelsson et al. (2001) establish that a simple variant of our model is already hard to solve (formally, it is “NP-hard”). They also discuss the structure of the optimal portfolio. Gaivoronski and Pflug (2004) consider the problem of minimizing VaR subject to a constraint on the expected portfolio value. They note that

“VaR optimization is inherently more difficult than, for example, variance optimization and efficient solution algorithms for this problem are lacking.”

Larsen et al. (2002) also mention that:

“Because of the multi-extremum structure of the VaR function, it is difficult to optimize it with standardly available methods.”

For the same reason, Gilli et al. (2006) propose a local search (“threshold accepting”) heuristic to minimize VaR.

Yiu (2004) considers a dynamic model where the objective is to maximize the expected value of the portfolio under a VaR constraint which is imposed continuously over time.

Several authors concentrate on models involving Conditional Value-at-Risk (CVaR) constraints. Recall that CVaR is the expected return associated with the fraction of portfolio values smaller than VaR, which implies that CVaR is smaller than or equal to VaR. For portfolios with a normal return/loss distribution, these two measures can be shown to be equivalent (Rockafellar and Uryasev 2000). Since CVaR is in many respects a “well-behaved” risk measure (subadditive, convex), it is tempting to rely on its relation with VaR to handle models involving VaR constraints or objective functions; see e.g. Krokmal et al. (2002), Larsen et al. (2002), Rockafellar and Uryasev (2000), etc., for such approaches. As noted by Gaivoronski and Pflug (2004), however, VaR and CVaR may behave differently in some realistic cases:

“We argue that for investor whose risk preferences are expressed in terms of Value-at-Risk it is important to consider this measure directly because other risk measures like variance, or even seemingly related CVaR, may represent a poor substitute.”

(See also Sect. 5 for a discussion of CVaR in relation to VaR.)

As already mentioned above, and when compared with these related papers, our work appears to be unique in its simultaneous consideration of multiperiod scenarios, of options, of a broad range of realistic financial parameters, and in the emphasis that it places on an adequate calibration of the tree of scenarios (see the next section).

4 Instantiation of the tree of scenarios

In this section, we discuss two specific difficulties that arise when the **VaROpt** model described in Sect. 3 is to be instantiated, namely: the generation of the scenarios, and the determination of the option prices. These questions must be answered before we can solve any instance of the **VaROpt** model.

4.1 Distribution of asset prices

In order to generate the tree of scenarios, a first requirement is to model the probability distribution of the values of the underlying asset over the two periods under consideration. This important question is not exclusively linked to the VaR problem under consideration, but constitutes a broad topic in itself. In order to handle it, we rely here on a characterization of implied probability density functions derived from observed option prices, together with stratified sampling. Note however that other methods could be used as well.

The implied distribution approach is based on papers by Breeden and Litzenberger (1978) and Shimko (1993); its use in connection with scenario trees is discussed by Schyns et al. (2005). In this approach, an analytical expression of the pdf is computed based on the information contained in the current option prices. This allows us to consider nonparametric distributions of returns (well beyond classical, but frequently inadequate, normal or generalized Student distributions) and to preserve consistency with the observed market prices, without relying on long, outdated time series of prices.

Practically, a pdf representing the distribution of the underlying index price is computed for each subtree of scenarios. For the first period, all the information required to generate the pdf is extracted from options available on the market at the time of the investment and with a maturity corresponding to the horizon of the subtree. For the second period, since options cannot be observed ahead of time, we make a simplifying assumption of stationarity and use the most recent set of options available on the market with the required maturity; therefore the same pdf is used for each subtree of the second period.

When the pdfs have been computed, we use stratified sampling in order to obtain a sample of n equiprobable states for each subtree. Stratified sampling preserves information relative to the distributions in the parameterization of the tree. Thanks to this modelling choice, complex continuous problems can be faithfully represented by a relatively small number of nodes, which makes them computationally tractable.⁴

4.2 Option pricing

Once the tree has been instantiated with the prices of the underlying asset, we can turn to the computation of the option prices. In order to be coherent with the **VaROpt** portfolio model and to obtain meaningful results, the option pricing model must be carefully selected. We now discuss this important point in detail.

4.2.1 Pitfalls of naive pricing procedures

The option values ($vopt$) at the leaves of each subtree can be easily determined from (1) since the options are at maturity and the index value at each node is known. The main calibration issue for an arbitrary option k , therefore, consists in determining its ask and bid prices (p_{ik}^a and p_{ik}^b) at the root of the tree i where this option appears. If $i = 0$, then option k is available at the start of the planning horizon and its price can be observed on the market. If $1 \leq i \leq n$, on the other hand, then the option will only appear at the beginning of the second period; therefore, its price is not initially observable and should be estimated, e.g., by the Black and Scholes formula. However, it is important to realize that such straightforward pricing procedures, when applied carelessly in a multiperiod scenario model, could lead to

⁴See Schyns et al. (2005) for a complete description of the procedure for a one-period tree.

undesirable arbitrage opportunities for (at least) two different types of reasons which we now proceed to discuss.

First, “real” arbitrage opportunities arising from the current market prices should be detected and handled appropriately in a preprocessing phase. Note that such opportunities may actually exist on the market and that the optimization model **VaROpt** can be seen as a tool to detect them: indeed, when arbitrage opportunities exist in the model, the optimal expected portfolio return is unbounded (infinite). But, in agreement with financial theory, such opportunities should disappear quickly when arbitrageurs manage to take advantage of them. In contrast with this dynamic situation, the static optimization model **VaROpt** cannot account for any adjustment of option prices resulting from such actions. Given that the final goal is to build a mid or long term portfolio using the **VaROpt** model (a goal which does not preclude the consideration of arbitrage opportunities for independent investment decisions), the investor must either discard any mispriced options from the optimization model (thus reflecting the fact that such options are usually illiquid), or accept minor price adjustments derived from the no-arbitrage conditions (such as those which may ensue from the simultaneous action of arbitrageurs).

Second, some “artificial” arbitrage opportunities could occur due to the calibration of the model. In particular, the tree of scenarios only provides a discrete, incomplete representation of the real market and of the continuous distribution of the index prices, and some of the parameters can only be approximately estimated. Here again, this may translate in undesirable arbitrage opportunities and in the prediction of infinite expected returns.

The theoretical framework for no-arbitrage pricing in scenario-based models is to be found in papers by Naik (1995), King (2002), Kallio and Ziemba (2007). A delicate issue, however, is that in incomplete and imperfect markets such as those considered here, the no-arbitrage conditions do not uniquely determine arbitrage-free option prices. This point is stressed, in particular, by King (2002), or by Kallio and Ziemba (2007); these authors provide various characterizations of the range of admissible prices, but they propose few guidelines as to how prices should actually be fixed in a numerical model.

Similar difficulties have been explicitly identified and discussed by several authors, e.g., Klaassen (1997, 1998) (who speaks of “spurious” profits), Kouwenberg and Vorst (1998) (who refer to artificial occurrences of “money machines”), Gulpinar et al. (2004), and Schyns (2001). In order to resolve them, to identify unsuitable option prices, and to adjust the prices of the remaining options as required, we rely on a pricing procedure adapted from Schyns et al. (2005) and based on no-arbitrage conditions; see also Klaassen (1997, 1998), Kouwenberg and Vorst (1998), Schyns (2001) for related proposals.

4.2.2 Arbitrage-free pricing

In order to set up the **VaROpt** model, we must estimate the bid and ask prices $p_{i,k}^b$ and $p_{i,k}^a$ to be used in (2), for every period i and every option $k = 1, \dots, O_i$ considered in period i .

In line with the approach adopted by Kallio and Ziemba (2007), the no-arbitrage conditions for model **VaROpt** should express that the portfolios associated with the scenario tree cannot generate nonnegative payoffs at every node of the tree, and at least one strictly positive payoff in some node. Using duality theorems as in Kallio and Ziemba (2007), this can be turned into equivalent conditions on the bid and ask prices $p_{i,k}^b$ and $p_{i,k}^a$ (which, as noted earlier, do not uniquely determine the prices).

In order to implement the approach, therefore, we develop a pricing process which consists of three successive stages: for $i = 0, 1, \dots, n$ and $k = 1, \dots, O_i$,

(a) determine tentative bid and ask prices $p_{i,k}^b$ and $p_{i,k}^a$;

- (b) check for the existence of arbitrage-free prices $\tilde{p}_{i,k}$ which lie as close as possible to (and if possible within) the corresponding range $[p_{i,k}^b, p_{i,k}^a]$;
- (c) remove the options whose price $\tilde{p}_{i,k}$ does not fall in the spread range $[p_{i,k}^b, p_{i,k}^a]$.

We now successively discuss the three stages in more detail.

Stage (a): Tentative prices. At the root of the tree (when $i = 0$), $p_{i,k}^b$ and $p_{i,k}^a$ are simply taken to be the observed bid and ask market prices.

At the beginning of the second period (for $i = 1, \dots, n$), as already mentioned, the prices are not observable. Therefore, we first compute the Black-Scholes (henceforth BS) price of each option k based on the volatility of the pdf associated with the corresponding subtree i . This BS price is denoted $target_{i,k}$. Then, a bid price $p_{i,k}^b = target_{i,k} - \delta_{i,k}$ and an ask price $p_{i,k}^a = target_{i,k} + \delta_{i,k}$ are determined. The bid-ask spread ($2\delta_{i,k}$) is determined as a function of the target price (based on a quadratic regression of the observed spreads at the first period with respect to the corresponding mean option prices).

This procedure, albeit somewhat *ad hoc*, yields realistic target prices and bid-ask spreads,⁵ and appears to provide an adequate description of the range of option prices arising in the second period of the model.

Stage (b): Arbitrage conditions. Rather than excluding arbitrage opportunities which may arise from the cash flow streams associated with the *complete* tree of scenarios, we choose for simplicity to impose non-arbitrage conditions on *each* subtree $i = 0, 1, \dots, n$. This slightly stronger condition allows us to reduce the sizes of the pricing subproblems to be solved, as explained hereunder.

More precisely, for each subtree $i = 0, 1, \dots, n$ and each option $k = 1, \dots, O_i$, we want to make sure that the spread range $[p_{i,k}^b, p_{i,k}^a]$ contains a price $\tilde{p}_{i,k}$ such that the collection of prices $\tilde{p}_{i,k}$ satisfy the following no-arbitrage conditions for each i :

$$\begin{pmatrix} 1 \\ \tilde{p}_{i,1} \\ \vdots \\ \tilde{p}_{i,O_i} \end{pmatrix} = \begin{pmatrix} (1+r) & \dots & (1+r) \\ p_{i,1,1} & \dots & p_{i,1,n} \\ \vdots & \vdots & \vdots \\ p_{i,O_i,1} & \dots & p_{i,O_i,n} \end{pmatrix} \begin{pmatrix} \tilde{\psi}_{i,1} \\ \tilde{\psi}_{i,2} \\ \tilde{\psi}_{i,3} \\ \vdots \\ \tilde{\psi}_{i,n} \end{pmatrix}, \tag{13}$$

$$\tilde{p}_{i,k} \geq 0 \quad (k = 1, \dots, O_i), \quad \tilde{\psi}_{i,j} > 0 \quad (j = 1, \dots, n), \tag{14}$$

where O_i is the number of options in the subtree, r is the risk-free rate, $\tilde{p}_{i,k}$ is the initial price of option k , $p_{i,k,j} = vopt(S_{i,j}, K_k)$ is the final value of option k in scenario j , and $\tilde{\psi}_{i,j}$ is the Arrow-Debreu state-price attached to scenario j . The quantities $\tilde{p}_{i,k}$ and $\tilde{\psi}_{i,j}$ are unknown and should be determined by solving the system (13)–(14) subject to the additional conditions

$$p_{i,k}^b \leq \tilde{p}_{i,k} \leq p_{i,k}^a \quad \text{for } k = 1, \dots, O_i. \tag{15}$$

⁵As pointed out by a reviewer, arbitrage-free bid and ask prices can be computed by LP models similar to (13)–(14), and expressing the minimal, resp. maximal, capital needed to sub-replicate, resp. super-replicate each option. In our numerical experiment, however, the gap between these theoretically-derived bounds exceeds by orders of magnitude the real spreads observed on financial market, which emerge as a result of the interaction between buyers, sellers and market makers. In this respect, our approach delivers much more realistic estimates.

Note that the existence of solutions to (13)–(14) guarantees that the system of prices $\tilde{p}_{i,k}$, $k = 1, \dots, O_i$ does not induce arbitrage opportunities in the subtree i . Moreover, no new arbitrage opportunities can arise if the prices $\tilde{p}_{i,k}$ are increased by a positive amount to reflect the buyer’s prices $p_{i,k}^a$, or are decreased by a positive amount to reflect the seller’s prices $p_{i,k}^b$. This intuitive observation is formally expressed by Proposition 1 in Kallio and Ziemba (2007).

Thus, in conclusion, we can check for the absence of arbitrage opportunities by solving the system of linear inequalities (13)–(15). If this system has a solution, then the bid and ask prices $p_{i,k}^b, p_{i,k}^a$ can be used in the portfolio selection model **VaROpt**: since both models involve the same set of scenarios, no arbitrage can arise in **VaROpt** under this system of prices. (Note that the prices $\tilde{p}_{i,k}$ themselves are not explicitly used in the portfolio selection model: we are only interested in asserting their existence, not their exact value.) If the system (13)–(15) has no solution, however, we would like to be able to reject a (hopefully) small subset of options in order to restore the no-arbitrage conditions.

For this purpose, rather than solving (13)–(15), we solve instead the following model, where the variable $\tilde{e}_{i,k}^a$ (resp., $\tilde{e}_{i,k}^b$) measures the amount by which the option price $\tilde{p}_{i,k}$ deviates from the target ask price $p_{i,k}^a$ (resp., from the target bid price $p_{i,k}^b$):

Pricing model (subtree i)

$$\begin{aligned}
 \min \quad & \sum_{k=1}^{O_i} (\tilde{e}_{i,k}^b + \tilde{e}_{i,k}^a) \\
 \text{subject to} \quad & \text{arbitrage equations (13)} \\
 & \tilde{e}_{i,k}^b \geq p_{i,k}^b - \tilde{p}_{i,k} & (k = 1, \dots, O_i) \\
 & \tilde{e}_{i,k}^a \geq \tilde{p}_{i,k} - p_{i,k}^a & (k = 1, \dots, O_i) \\
 & \tilde{e}_{i,k}^a \geq 0, \tilde{e}_{i,k}^b \geq 0 & (k = 1, \dots, O_i) \\
 & \tilde{p}_{i,k} \geq 0 & (k = 1, \dots, O_i) \\
 & \tilde{\psi}_{i,j} > 0 & (j = 1, \dots, n).
 \end{aligned} \tag{16}$$

Stage (c): Removal of options. At the end of the previous stage, after having solved model (16), any option whose arbitrage-free price $\tilde{p}_{i,k}$ does not fall within the bid-ask range $[p_{i,k}^b, p_{i,k}^a]$ is rejected from our portfolio optimization model. Alternatively, it might also be possible to adjust the bid-ask spread of these options around the prices $\tilde{p}_{i,k}$. The first choice, however, seems preferable since it allows us to stick to marked-to-market prices. In practice, in our numerical experiments, very few options are rejected (see Sect. 5).

5 Case study

5.1 Experimental setting

The theoretical model presented above has been tested on a set of real-world option data. The objectives of this experiment are twofold. From a computational perspective, we want to investigate how fast an optimal, or nearly optimal, portfolio of options under VaR constraints can be computed. From a financial perspective, we want to examine the nature and features of portfolios constructed under VaR and guarantee constraints.

We consider a portfolio of options on the S&P500 index. The risk-free asset is included in order to allow more flexibility in market risk management. Conversely, inclusion of the index

itself is not allowed in the portfolio, as it would be redundant with simple combinations of long calls and short puts. In the context of this study, the utility of the index is essentially to model the future and the option prices.

We consider a one year investment horizon, with the opportunity to rebalance the investment after two months. Note that the first period is shorter than the second one, which is justified in the framework of a roll-over strategy: the investor is able to rebalance his portfolio in the short term, but a longer horizon is also taken into account in the optimization process. At the end of the first period, a new problem can be formulated, based on a different horizon and on updated data.

To construct the implied pdfs, the S&P500 price and the corresponding option prices have been collected on the CBOE on an arbitrary date, namely February 16, 2006. The S&P500 price at this date was 1289.38 USD. The distribution of returns for the first period is based on all the available options with a remaining maturity of two months. For the second period, options with a maturity of ten months are used. Each subtree contains 80 equiprobable branches, that is, 6400 possible states of the world are under consideration at the end of the investment horizon.

The risk-free rate for each period is set according to the T-bill rates with corresponding maturities, respectively at 0.91% and 3.81% for each period. Implicit dividend yields are extracted from the application of the put-call parity on option prices and are coherent with observed values. A simple tax rate of 0.17% is applied on option transactions.

For the first period, 67 different options (calls and puts) are observed on the market, and their bid and ask prices are used to instantiate the option pricing model (16). For the second period, a distinct set of options is created in each subtree (note that these options cannot be initially observed). Namely, 15 virtual calls and 15 virtual puts are created at each node $i = 1, 2, \dots, n$, with strike prices spaced by 25 USD around the index value corresponding to node i . For each such option, bid and ask prices are computed as explained in Sect. 4.2.2.

We impose a guarantee on the final portfolio value, which must be larger than or equal to $\theta_2 = 80\%$ of the initial budget. To reduce the exposure, a second guarantee constraint is set at $\theta_1 = 50\%$ at the end of the first period.

The VaR constraint must be satisfied at the end of the horizon with a (typical) confidence level $P = 95\%$. Several VaR thresholds are also considered, ranging from $\lambda = 80\%$ up to $\lambda = 110\%$ of the initial budget.⁶

5.2 Optimization

The global problem involves 2467 options, which are modelled by a tree containing 6400 terminal nodes (or scenarios) and 81 internal nodes. Arbitrage-free prices are obtained by optimizing the option pricing subproblem (16) associated with each subtree. For these subproblems, the CPLEX optimization software (version 10, running on a 2 GHz notebook computer) provides the optimal solutions in a fraction of a second. Only one (extreme) put must be rejected for the first period, and no option is rejected for the second period (the BS prices are almost arbitrage-free for the second period).

As expected, the portfolio optimization **VaROpt** model is much harder to solve. It is a mixed integer programming (MIP) problem. Branch-and-cut algorithms are typically used to solve MIP problems, see e.g. Caprara and Fischetti (1997). At every iteration of a branch-and-cut algorithm, lower and upper bounds on the optimal value of the problem are computed. The lower bound is the value of the best feasible solution computed so far. The upper

⁶Remember that no VaR constraint is imposed at the end of the first period.

Table 1 Optimization results

VaR (λ)	104%	100%	96%	92%	88%	84%	80%
Expected return	6.86%	17.30%	27.73%	38.20%	48.59%	59.01%	69.44%
Final gap	7.14%	6.28%	5.45%	4.19%	3.02%	1.80%	0.00%
Semi-stdev	3.06%	17.30%	31.67%	46.05%	60.46%	74.90%	89.34%
CVaR	102.70%	98.81%	94.94%	90.71%	87.42%	83.83%	80%
Omega	3.65	3.65	3.64	3.65	3.63	3.63	3.62

bound provides a (pessimistic) estimate of the gap between the value of this solution and the optimal value of the problem. This gap decreases in the course of the iterations, as better feasible solutions are found and/or more accurate upper bounds are computed.

Unfortunately, branch-and-cut algorithms are often time-consuming and, after a number of iterations, the residual gap usually decreases slowly, so that the process may have to be terminated before the optimal solution has been found or confirmed.

The standard branch-and-cut algorithm of CPLEX has been used in our experiments; its computation time has been limited to one hour. The results of this optimization process are presented in Table 1. The second line of the table displays the expected return associated with the best portfolio obtained within the one hour limit. The portfolio obtained is feasible, but it is not certified to be optimal. The gap gives an indication of the *maximum* additional return that could be obtained if the optimization process was able to terminate within the time limit.⁷

With its confidence level being set at $P = 95\%$ (a rather high value), the maximal attainable value for the Value-at-Risk should be close to $(1 + r)$, where r is the risk-free rate. The annual risk-free rate is equal to $r = 4.75\%$ and indeed no feasible solution exists for VaR levels higher than $\lambda = 105\%$, corresponding to an increase of 5% of the initial budget. We also notice that, the more constrained the problem, the larger the gap (this is quite typical when solving MIP problems). When the VaR is set at $\lambda = 80\%$, the problem is extremely easy because the guarantee constraint dominates the VaR constraint and the problem becomes a continuous linear programming problem. But when the VaR constraint is strengthened, more computing time is required to reduce the gap, and one hour is not enough to close it completely.

The portfolio semi-standard deviation (*ssd*) is another classical measure of risk and, since it shares some properties with the VaR, we have also displayed it in Table 1. It is classically defined as the standard deviation of the portfolio values where all values larger than the expected value are replaced by the expected value, i.e., with the same notations as in the **VaROpt** model:

$$ssd = \sqrt{\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n [\min\{0, V_i(S_{i,j}) - ER\}]^2}$$

Observe that the volatility (as measured by *ssd*) increases with the exposure to risk (as measured by λ) and with the increase of the expected return. The semi-standard deviation values are of the same order of magnitude as the expected returns.

⁷Note however that a positive gap *does not* mean that the available portfolio is necessarily suboptimal: with more computing time, CPLEX may be able to decrease the upper bound, and hence the gap.

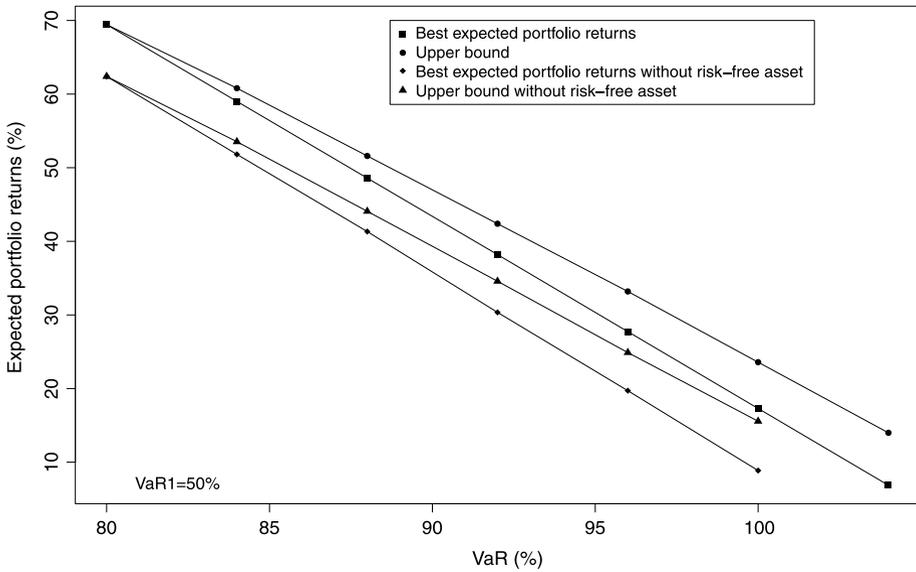


Fig. 3 Expected portfolio returns vs. VaR

Also, recall that the CVaR (Conditional Value-at-Risk) is the expected return associated with the fraction of portfolio values smaller than VaR. In the scenario tree model, we approximate it by the expression:

$$CVaR = \frac{1}{[(1 - P)n^2]} \sum_{k=1}^{[(1-P)n^2]} V_{(k)}$$

where $\{V_{(1)}, \dots, V_{(n^2)}\}$ is the set of final portfolio values $V_i(S_{i,j})$ sorted by increasing order.

Two remarks should be made. First, by definition, CVaR is always smaller than VaR. Secondly, observe that we compute CVaR at the confidence level P on the whole set of final leaves although the VaR constraint was separately specified for each subtree of the second period. This is not incoherent since the local VaR constraints imply exactly the same VaR level at the global level. The advantage is that only one CVaR value, instead of n ones, needs to be computed to perform the comparison.

CVaR and VaR are expressed as percentages of the initial budget. In our experiments, CVaR follows the same trend as VaR and is always close to it; see Table 1. (This is in agreement with the conclusions reported for instance by Rockafellar and Uryasev (2000), although Gaivoronski and Pflug (2004) make the opposite observation in their case-studies.)

Except for a VaR of 104%, which is not economically reasonable, all portfolios' expected returns seem very high with respect to the risk exposure, even after accounting for the remaining optimization gap. It turns out that the expected return is nearly a linear function of the VaR level. This is best seen in Fig. 3, where the best available expected return is represented as a function of the VaR, together with the upper bound on the optimal expected return.

Finally, Table 1 displays the Omega ratio proposed by Keating and Shadwick (2002), a performance measure for actively managed portfolios that suits particularly well the inclusion of options. The Omega ratio is defined as the quotient of the truncated mean return

above a reservation value divided by the truncated mean return below the same reservation value. In other terms, it can be interpreted as the ratio of a call option to a put option on the actively managed portfolio (see Kaplan and Knowles 2004). We took the risk-free rate as the reservation rate. The ratio indicates a very stable value of Omega around 15:4 (3.62 to 3.65). The fact that the portfolio achieves similar option-adjusted performance levels, irrespective of the VaR level, reinforces our observation of an almost-linear pattern of portfolio returns. This provides some indication that higher levels of return are not obtained at the expense of a dramatic increase in risk (measured through the expected downward deviation from the benchmark), but rather through a proportional risk exposure.

5.3 Features of the optimal portfolios

In this section, we intend to discuss some of the most remarkable features of the (near) optimal portfolios obtained when solving the VaR model. We should note, however, that this discussion is not easy to express synthetically, since every feasible solution of the VaR model gives rise to several portfolios: indeed, for every possible state of the world at the end of the first period, the investor has the opportunity to rebalance his portfolio in a different way. Therefore, there are as many final portfolios as there are scenarios in the first period, and this considerably complicates the discussion.

Also, for the sake of conciseness, we limit ourselves to the results obtained when the VaR is fixed at $\lambda = 100\%$ (which is an attractive value for the investor), but most of the conclusions remain valid for other levels of VaR as well.

The portfolio payoff solely depends on the investment made in each subtree of the second period and on the final index price. With 80 possible portfolios for the second period, it is not reasonable to draw a figure for each case. In Fig. 4, we focus on the second-period portfolios linked to two special scenarios, namely, on those portfolios associated with the largest decrease during the first period, or with a stationary index value during the first period. The figure displays the value of each of these two portfolios at the end of period 2, as a function of the index value at that time.

In order to plot faithfully these portfolio values, we have again exploited the properties that were used to model the guarantee constraint; i.e., the portfolio values were only computed at the option strike prices and linearly interpolated between these prices. The strike prices increase from 800 USD up to 1625 USD, which is coherent with the final index values in the tree ranging from 570 USD up to 1750 USD (with an initial price of 1289.38 USD). For the sake of clarity and since the payoff for all portfolios is flat for large index values, the figure was truncated at 1350 USD.

We also draw the “range” of the payoffs generated by the subtrees of the second period: that is, for each strike price defined at the end of the second period, independently of the subtrees, we compute the value of the 80 portfolios and we display the smallest and largest values obtained. This gives rise to two curves, where each curve is therefore typically the combination of portfolio values associated with different subtrees. The curve of smallest values provides a guarantee on the final portfolio value: whatever happens at the end of the first period, the portfolio value never dives below this level at the end of the investment horizon. The largest value indicates the highest return that can possibly be reached at the end of the second period. Note also that the segment of line between two strike prices is only an approximation of the real extreme possible portfolio returns. In such intervals, the approximation is pessimistic for the lowest values and optimistic for the largest values.

Let us first consider the case of the second-period portfolio linked to the first-period scenario corresponding to the smallest index value. As seen in Fig. 4, the shape of the payoff

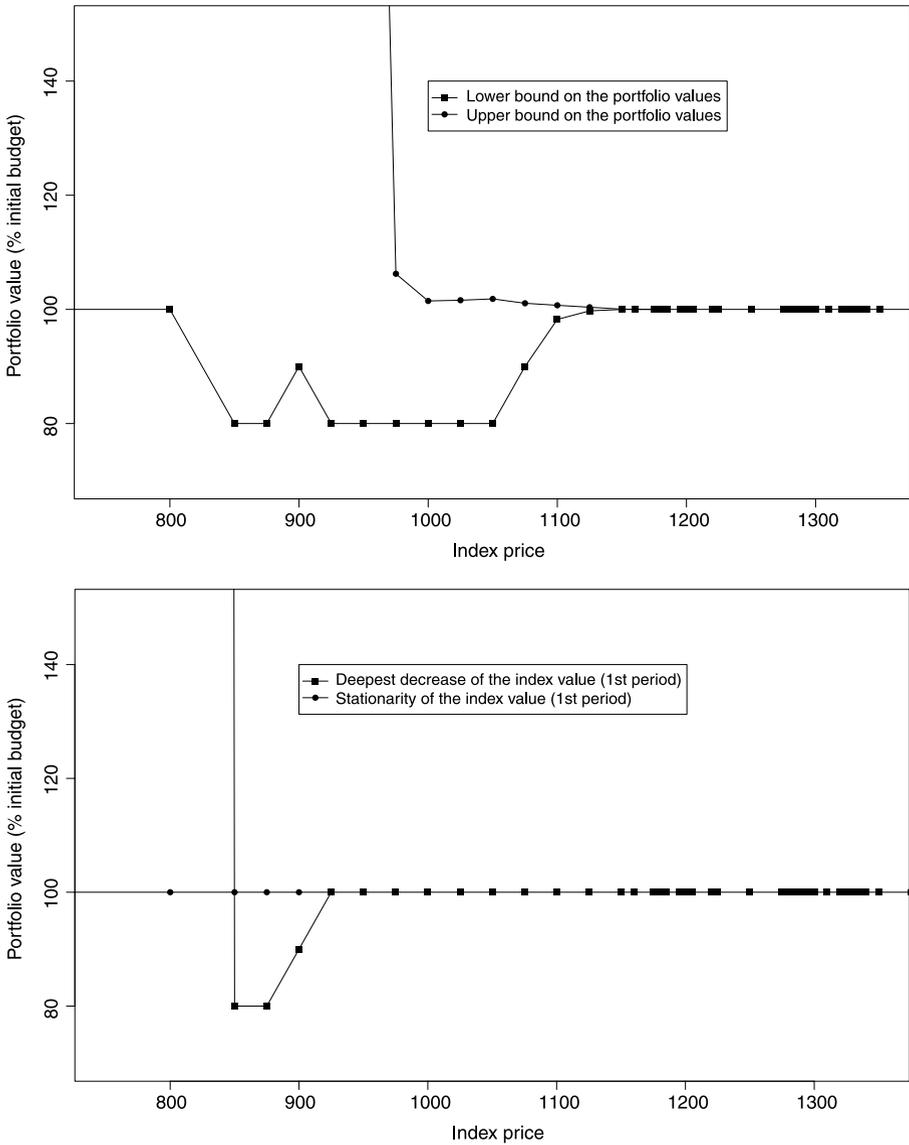


Fig. 4 Portfolio value vs. index price

function is very regular. Except for very small deviations, only two return values dominate the graph, namely the values corresponding to the guarantee level and to the VaR level. Between these two values, the payoff evolves in a linear fashion. For extreme index prices, we observe a flat portfolio payoff function, corresponding to the VaR level on one side, and a linearly decreasing return down to the guarantee level on the other side. This kind of shape is well known in finance, and is a generic feature of VaR models; see, e.g., Dert and Oldenkamp (2000) or Liang et al. (2005).

When we look at final portfolios corresponding to slightly larger index values at the end of the first period, the same behavior is observed. The only differences are essentially limited to slight shifts of the two breakpoints.

Such simple shapes can be obtained using very few options. Three options with strike prices associated to the guarantee and VaR levels appear to suffice.

This behavior can be explained by the high expected rate of return and very large skewness of deeply out-of-the-money option payoffs. In order to maximize the expected portfolio return, the optimal strategy thus involves trading these kinds of options (with a long bias), while the proportion invested in the riskless asset is large enough to ensure the non-violation of the VaR and guarantee constraints.

When we next consider the final portfolios associated with larger values of the index at the end of the first period, and especially for the case of the index stationarity plotted in Fig. 4, we observe that the payoff function is nearly flat around the VaR level. This shape is the most frequently encountered in our numerical experiments (75% of the largest index values). Such a shape can be and is obtained by investing only in the risk-free asset and therefore by avoiding the options. In other words, the investor tries to maximize his investment for small values of the index price and invests only in these cases in options.

Analyzing the composition of the portfolios confirms the previous observations. Table 2 gives the total number of long and short positions in puts and calls for each VaR level during the second period, i.e., the total number over the 80 portfolios of the second period. In a given portfolio, we can observe up to two long positions on puts, one short put, one long call and one short call.

Table 2 also displays the percentage of the initial budget invested in the risk-free asset. The constraints for the first period are weak and therefore the investments essentially aim at satisfying the final constraints. The risk-free amounts initially invested correspond approximately to the minimal amounts required to reach the VaR level at the horizon of investment. This means that the optimal solution consists in locking the VaR level, and then using the remaining budget to maximize the return (for small index values).

More precisely, the initial portfolio always consists of a large investment in the risk-free asset and of a long position in a put (strike price 1175 USD). With the increase of the VaR level, the position in the put is replaced by a larger investment in the risk-free asset. This leads to an interesting conclusion. Indeed, if we assume a roll-over strategy, then the investor will never construct the portfolios proposed for the second period, but he will only invest in the initial, very simple portfolio. At the end of the first period, he will use again a similar model to re-optimize his investment. This suggests that, under the VaR constraint, when the horizon of investment is repeatedly rolled over, the strategy of holding only one position in a single option at any time probably approaches the best one.

Table 2 Portfolio composition

VaR	104%	100%	96%	92%	88%	84%	80%
long puts	29	23	22	47	34	23	3
short puts	20	18	29	35	21	18	0
long calls	4	5	4	14	4	5	0
short calls	4	4	4	14	5	4	0
risk-free	99.29%	95.47%	91.65%	87.83%	84.01%	80.19%	76.37%

5.4 Discussion

The pattern depicted in Fig. 4 is rather atypical and may appear to contradict results by Dert and Oldenkamp (2000), who rely on a theorem of Dybvig (1988) to establish that, under certain conditions, higher portfolio values are associated with higher index values in the optimal portfolio. Actually, such bullish patterns occur frequently when we apply our model on various data sets, and they can be expected to arise in practice as well. Therefore, we now provide some additional comments about the pattern of returns observed in Fig. 4.

First let us recall that Dert and Oldenkamp (2000) make two main assumptions in their paper: market completeness and normality of the index returns. They also state:

“(...) the assumption of lognormally distributed index returns, a positive risk premium, and the objective of maximizing expected return imply the optimality of portfolios with monotonic payoff functions. However, practitioners may want to make different assumptions on the return distribution of the underlying value, which do not necessarily imply optimality of monotonic payoff functions.”

More precisely, we observe that, depending on the distribution of the index, the state-price densities are not always decreasing functions of the index prices—which Dybvig’s theorem does not rule out—and, therefore, while not violating Dybvig’s theorem, the optimal portfolio values should not always be ordered like the index prices. Thus, payoffs corresponding to bearish, butterfly or strangle patterns can also appear when considering alternative distributions that imply various orderings of the state prices relative to the index payoffs.

In this paper, we do not assume a normal distribution of the index prices, but an implied distribution based on the initial option prices. This pdf departs significantly from normality; in particular, it features a large skewness.

A second explanation for the pattern displayed in Fig. 4 is that we consider an incomplete market with frictions. Therefore, Dybvig’s theorem is not applicable in this context, which may also explain the occurrence of alternative patterns for the optimal portfolios.

For the above reasons, the precise relation between portfolio values and index values is not easily stated, and model **VaROpt** cannot be reduced to a simple continuous linear programming model as was the case under the assumptions made by Dert and Oldenkamp (2000).

Note however that for all distributions we considered, some of the conclusions remain similar: the optimal portfolio only involves a few options (puts or calls, depending on the pattern), and the shape of the portfolio payoff is strikingly simple (one level associated with the guarantee constraint and another one associated with the VaR constraint). Dert and Oldenkamp (2000) drew similar conclusions.

Based on these observations, it appears possible to develop a heuristic approach to directly compute a “near-optimal” portfolio. This is explained in the next section.

5.5 Heuristic approach

Note that the complexity of the portfolio optimization model **VaROpt** is due to the presence of the binary variables $\alpha_{i,j}$, which identify those scenarios where the VaR requirement λB is satisfied. If the values of these binary variables could somehow be initially “guessed”, then the optimization problem would boil down to a simple continuous linear programming problem, rather than a difficult MIP. Dert and Oldenkamp (2000) rely on this observation, since the highest portfolio values are associated with the highest index values in their model

Table 3 Heuristic approach

VaR	104%	100%	96%	92%	88%	84%	80%
Heuristic return	6.70%	17.16%	27.62%	38.07%	48.53%	58.99%	69.44%
Branch-and-cut return	6.86%	17.30%	27.73%	38.20%	48.59%	59.01%	69.44%

(see above). A related observation is also heuristically exploited by Gaivoronski and Pflug (1999).

In general, however, it seems difficult to correctly guess the values of the variables $\alpha_{i,j}$. From our previous experimental results, it appears that (for the data set at hand), the investor tries to optimize his investment for the lowest values of the index (this is analogous to a bearish investment strategy). Therefore, assuming at the outset that the portfolio exceeds the VaR for the 95% of the scenarios corresponding to the lowest index prices, accordingly setting the corresponding $\alpha_{i,j}$ -variables to 1, and optimizing the remaining model seems intuitively appealing.

As shown in Table 3, this simple and very fast heuristic approach already yields excellent results (for the sake of comparison, the results delivered by the branch-and-cut algorithm are also reproduced in Table 3). These observations suggest that a near-optimal solution could probably be found analytically, without a costly enumeration process.

Note that the above heuristic is clearly inspired by the results obtained for this specific data set. We can assume, however, that when performing more experiments, other patterns could be observed for the portfolio payoff function: investment in options when the index price is high (analogous to a bullish strategy), or when the index price is either stable or volatile (analogous to a volatility, or to a butterfly strategy), etc. Thus, we could develop a family of heuristics by successively considering several such schemes. The heuristic solutions obtained in this way would quickly provide a sample of feasible solutions (which could also be used, for instance, e.g., as alternative starting points of a branch-and-cut process). A complete investigation of these ideas is left for further research.

6 Conclusions

This paper presents a scenario tree model for a complex computational problem, namely the selection of an optimal portfolio involving options, subject to Value-at-Risk (VaR) constraints. We have shown that formulating this problem as a mathematical optimization problem requires some special care, especially for the generation of scenarios and for the determination of associated option prices which rule out arbitrage opportunities. The resulting mixed integer programming problem is very large, but can be solved (approximately) by CPLEX or by ad hoc heuristics.

Beyond the methodological discussion, an application to real data could presumably produce very complex investment strategies. However, at least two of our conclusions appear to be strikingly simple. First, the optimal solutions obtained for different VaR levels feature an almost linear relation between risk (measured by VaR) and return. This would be a standard result in an asset pricing model framework, provided that the VaR is a measure of risk shared by all investors. Yet, this cannot be the case since Value-at-Risk is not a coherent measure of risk. Second, the optimal trading strategy can be roughly summarized by a combination of a risk-free investment with long positions in a set of out-of-the-money puts. Thus, even though this solution emerges from a complex optimization program, it can still be interpreted in a simple manner.

A strength of our model is its flexibility. Indeed, the generic model already integrates several realistic financial constraints, and the methodology provides many degrees of freedom in order to consistently instantiate additional real market features or requirements. Work in progress includes the consideration of options characterized by different horizons of investment over different periods, as well as the simultaneous consideration of several securities and associated derivatives. In order to handle these complex problems, it may prove necessary to resort to efficient heuristics, since exact methods tend to become prohibitively time consuming for the resulting models.

Acknowledgements This work was initiated as part of a research project jointly conducted with the late Antoon Kolen. Antoon's untimely death prevented him from seeing the end of the project. We wish nevertheless to acknowledge his initial impetus to our work. We also thank several reviewers for their valuable comments and suggestions on earlier drafts of our paper. Financial support from Deloitte Luxembourg is gratefully acknowledged.

References

- Artzner, P., Delbaen, F., Eber, J.-M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9, 203–228.
- Bertrand, P., & Prigent, J.-L. (2005). Portfolio insurance strategies: OBPI versus CPPI. *Finance*, 26(1), 5–32.
- Birge, J. R., & Louveaux, F. (1999). *Introduction to stochastic programming* (2nd Ed.). *Springer series in operations research*. New York: Springer.
- BIS (2006). Bank for International Settlements. International convergence of capital measurement and capital standards—A revised framework. Comprehensive version. Available at BIS: <http://www.bis.org/publ/bcbs107.htm>.
- Black, F., & Jones, R. (1987). Simplifying portfolio insurance. *Journal of Portfolio Management*, 14, 48–51.
- Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51, 621–651.
- Caprara, A., & Fischetti, M. (1997). Branch-and-cut algorithms. In M. Dell'Amico, F. Maffioli, & S. Martello (Eds.), *Annotated bibliographies in combinatorial optimization* (pp. 45–64). New York: Wiley.
- Daniëlsson, J., Jørgensen, B. N., de Vries, C. G., & Yang, X. (2001). *Optimal portfolio allocation under a probabilistic risk constraint and the incentives for financial innovation* (Tinbergen Institute Discussion Paper TI 2001-069/2).
- Dembo, R. S. (1991). Scenario optimization. *Annals of Operations Research*, 30, 63–80.
- Dert, C., & Oldenkamp, B. (2000). Optimal guaranteed return portfolios and the casino effect. *Operations Research*, 48, 768–775.
- Dybvig, P. H. (1988). Inefficient dynamic portfolio strategies or how to throw away a million dollars in the stock market. *The Review of Financial Studies*, 1, 67–88.
- Dybvig, P. H. (1988). Distributional analysis of portfolio choice. *Journal of Business*, 61, 369–393.
- Edgeworth, F. Y. (1888). The mathematical theory of banking. *Journal of the Royal Statistical Society*, 51(1), 113–127.
- Esch, L., Kieffer, R., & Lopez, Th. (2005). *Asset and risk management: risk oriented finance*. Chichester: Wiley.
- Gaivoronski, A. A., & Pflug, G. (1999). Finding optimal portfolios with constraints on Value-at-Risk. In *Proceedings III Stockholm seminar on risk behavior and risk management*, Stockholm.
- Gaivoronski, A. A., & Pflug, G. (2004). Value-at-Risk in portfolio optimization: Properties and computational approach. *Journal of Risk*, 7, 1–31. <http://ssrn.com/abstract=702841>.
- Gilli, M., Kellezi, E., & Hysi, H. (2006). *A data-driven optimization heuristic for downside risk minimization* (Working paper). Department of Econometrics, University of Geneva.
- Grinold, R. C. (1999). Mean-variance and scenario-based approaches to portfolio selection. *The Journal of Portfolio Management* (Winter), 10–22.
- Gülpinar, N., Rustem, B., & Settergren, R. (2004). Simulation and optimization approaches to scenario tree generation. *Journal of Economic Dynamics and Control*, 28, 1291–1315.
- Kallio, M., & Ziemba, W. T. (2007). Using Tucker's theorem of the alternative to simplify, review and expand discrete arbitrage theory. *Journal of Banking and Finance*, 31, 2281–2302.
- Kaplan, P. D., & Knowles, J. A. (2004). Kappa: A generalized downside risk-adjusted performance measure. *Journal of Performance Measurement*, 8, 42–54.

- Keating, C., & Shadwick, W. F. (2002). A universal performance measure. *Journal of Performance Measurement*, 6, 59–64.
- King, A. J. (2002). Duality and martingales: A stochastic programming perspective on contingent claims. *Mathematical Programming, Series B*, 91, 543–562.
- Klaassen, P. (1997). Discretized reality and spurious profits in stochastic programming models for asset liability management. *European Journal of Operational Research*, 101, 374–392.
- Klaassen, P. (1998). Financial asset-pricing theory and stochastic programming models for asset-liability management: A synthesis. *Management Science*, 44, 31–48.
- Kouwenberg, R. (2001). Scenario generation and stochastic programming models for asset liability management. *European Journal of Operational Research*, 134, 279–293.
- Kouwenberg, R., & Vorst, T. (1998). *Dynamic portfolio insurance: A stochastic programming approach* (Erasmus Center for Financial Research Report, Working paper 9909).
- Krokhmal, P., Palmquist, J., & Uryasev, S. (2002). Portfolio optimization with conditional Value-at-Risk objective and constraints. *Journal of Risk*, 4(2), 43–68.
- Larsen, N., Mausser, H., & Uryasev, S. (2002). Algorithms for optimization of Value-at-Risk. In P. Pardalos & V. K. Tsitsiringos (Eds.), *Applied optimization series. Financial engineering, e-commerce and supply chain* (pp. 129–157). Boston: Kluwer Academic.
- Liang, J., Zhang, F., & Li, D. (2005). *Optioned portfolio selection: models and analysis* (Technical Report). Department of Systems Engineering & Engineering Management, The Chinese University of Hong Kong, SEEM2005-01.
- Mulvey, J. M. (1994). Financial planning via multi-stage stochastic programs. In J. R. Birge & K. G. Murty (Eds.), *Mathematical programming: State of the Art 1994* (pp. 151–171). Ann Arbor: University of Michigan Press.
- Muzzioli, S., Torricelli, C. (2005). The pricing of options on an interval binomial tree: an application to the DAX-index option market. *European Journal of Operational Research*, 163, 192–200.
- Naik, V. (1995). Finite state securities market models and arbitrage. In R. A. Jarrow, V. Maksimovic, & W. T. Ziemba (Eds.), *Handbooks in operations research and management science: Vol. 9. Finance* (pp. 31–64). Amsterdam: Elsevier.
- Prekopa, A. (1995). *Stochastic programming*. Dordrecht: Kluwer Academic.
- RiskMetrics (1996). *Technical document* (4th Ed.). New York: Morgan Guaranty Trust Company. <http://www.riskmetrics.com/research>.
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional Value-at-Risk. *Journal of Risk*, 2, 21–41.
- Rubinstein, M. (1994). Implied binomial trees. *The Journal of Finance*, XLIX(3), 771–818.
- Rubinstein, M. (1998). Edgeworth binomial trees. *The Journal of Derivatives*, 5(3), 20–27.
- Schyns, M. (2001). *Modelling financial data and portfolio optimization problems*. Doctoral Thesis, University of Liège.
- Schyns, M., Crama, Y., & Hübner, G. (2005). *Grafting information in scenario trees: Application to option prices* (Working Paper). HEC Management School, University of Liège. Revised, November 2005. Available at SSRN: <http://ssrn.com/abstract=418520>.
- Schyns, M., Hübner, G., & Crama, Y. (2009). Alternative to the mean-variance asset allocation analysis: A scenario methodology for portfolio selection. In G.N. Gregoriou (Ed.), *Stock market volatility* (pp. 231–254). London: Chapman Hall-CRC/Taylor and Francis.
- Shimko, D. (1993). Bounds of probability. *Risk*, 6(4), 33–37.
- Yiu, K. F. C. (2004). Optimal portfolios under a value-at-risk constraint. *Journal of Economic Dynamics and Control*, 28, 1317–1334.