CrossMark

Time-consistent risk-constrained dynamic portfolio optimization with transactional costs and time-dependent returns

Davi Valladão[1] · Thuener Silva[1] · Marcus Poggi[2]

Abstract

Dynamic portfolio optimization has a vast literature exploring different simplifications by virtue of computational tractability of the problem. Previous works provide solution methods considering unrealistic assumptions, such as no transactional costs, small number of assets, specific choices of utility functions and oversimplified price dynamics. Other more realistic strategies use heuristic solution approaches to obtain suitable investment policies. In this work, we propose a time-consistent risk-constrained dynamic portfolio optimization model with transactional costs and Markovian time-dependence. The proposed model is efficiently solved using a Markov chained stochastic dual dynamic programming algorithm. We impose one-period conditional value-at-risk constraints, arguing that it is reasonable to assume that an investor knows how much he is willing to lose in a given period. In contrast to dynamic risk measures as the objective function, our time-consistent model has relatively complete recourse and a straightforward lower bound, considering a maximization problem. We use the proposed model for approximately solving: (i) an illustrative problem with 3 assets and 1 factor with an autoregressive dynamic; (ii) a high-dimensional problem with 100 assets and 5 factors following a discrete Markov chain. In both cases, we empirically show that our approximate solution is near-optimal for the original problem and significantly outperforms selected (heuristic) benchmarks. To the best of our knowledge, this is the first systematic approach for solving realistic time-consistent risk-constrained dynamic asset allocation problems.

Keywords Stochastic programming · Time consistency · Dynamic asset allocation · Stochastic dual dynamic programming · Conditional value-at-risk constraints

✉ Davi Valladão
davimv@puc-rio.br

1   Industrial Engineering Department, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rua Marquês de São Vicente, 225 RDC, Gávea, Rio de Janeiro, RJ CEP 22451-900, Brazil

2   Informatics Department, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Rua Marquês de São Vicente, 225 RDC, Gávea, Rio de Janeiro, RJ CEP 22451-900, Brazil

1 Introduction

Portfolio selection is a complex dynamic decision process under uncertainty that must account for risk aversion, price dynamics, transaction costs and allocation constraints. Mathematically represented as a stochastic dynamic programming problem, asset allocation suffers from the curse of dimensionality, particularly when transaction costs are not negligible. Although the simplifying assumptions of previous works make dynamic asset allocation models more intuitive and computationally tractable, they are unreliable policies in practice.

The dynamic portfolio optimization literature begins with Mossin (1968), Samuelson (1969), and Merton (1969, 1971), where they show that a no transaction cost problem is tractable and efficiently solvable. In a discrete-time framework, Constantinides (1979) address a dynamic portfolio selection with proportional transaction costs, but only for a two-asset problem. Davis and Norman (1990) and Shreve and Soner (1994) can be considered continuous-time extensions of the work of Constantinides (1979), whereas Muthuraman (2007) numerically solved a one-asset problem with proportional transaction costs.

In the literature of Multistage Stochastic Programming (MSP) for Asset-Liability (Ziemba et al. 1998; Kouwenberg 2001; Valladão et al. 2014) and Portfolio management (Bradley and Crane 1972; Gülpınar et al. 2004; Topaloglou et al. 2008) realistic assumptions, such as transactional costs, have been considered due to the flexibility given by scenario-tree based models. However, only a limited number of decision stages ensure computational tractability (Shapiro and Nemirovski 2005; Shapiro 2006).

To circumvent the computational burden, Brown and Smith (2011) used dual bounds proposed by Brown et al. (2010) to assess the solution quality of proposed heuristic (non-optimal) allocation strategies for a multi-asset (up to ten assets) model with transaction costs. Bick et al. (2013) proposed a numerical procedure to obtain heuristic allocation strategies as modifications of the closed-form solution of an idealized problem, and a performance bound to measure its quality. Jin and Zhang (2013) considers a jump-diffusion model with investment constraints by embedding the constrained problem in the appropriate family of unconstrained ones. The works of Brown and Smith (2011), Bick et al. (2013) and Jin and Zhang (2013) reiterates the high computational burden of obtaining the actual optimal allocation strategy and the importance of developing an efficient solution algorithm for this problem. More recently, Mei et al. (2016) uses the specific structure of a multiperiod mean-variance to obtain a tractable solution for a multi-asset problem. To the best of our knowledge, no previous works provides an efficient methodology for a dynamic portfolio optimization with transactional costs and time-dependence returns with a large number of decision stages.

The stochastic dynamic programming literature, on the other hand, deals with considerably more complex and large-scale problems in a discrete-time setting. Decomposition methods with a full tree representation, such as progressive hedging, see Rockafellar and Wets (1991), and L-shaped, see Birge and Louveaux (2011), solve medium-sized problems. Additionally, sampling-based decomposition methods, such as stochastic dynamic dual programming (SDDP) by Pereira and Pinto (1991), abridged nested decomposition (AND) by Donohue and Birge (2006) and convergent cutting-plane and partial-sampling algorithm (CUPPS) by Chen and Powell (1999), successfully solve large-scale problems.

Particularly important to our work, the SDDP algorithm can efficiently solve a large-scale stochastic dynamic problem in the context of the operation planning of hydrothermal power systems (Soares et al. 2017; Brigatto et al. 2017; Street et al. 2017). The SDDP overcomes the curse of dimensionality assuming temporal (stage-wise) independence of the stochastic process. For a maximization problem, stage-wise independence guarantees

a unique concave value function for each time stage regardless of the stochastic process realization. Augmenting the state-space with a past realization of the stochastic process, the SDDP framework also accommodates a linear temporal dependence in the right-hand side of a linear constraint (Shapiro et al. 2013). It is important to note that any time dependence (even a linear one) embedded on the left-hand side of a linear constraint (which is exactly the case for a portfolio problem) would not be tractable due to a non-convex augmented value function. Furthermore, the SDDP framework is also suitable for a discrete-state Markov-dependent stochastic process, as in Mo et al. (2001), Philpott and de Matos (2012) and Bruno et al. (2016), which is particularly helpful when solving a stochastic dynamic asset allocation problem.

Finally, risk aversion is an important issue in a dynamic stochastic portfolio optimization. In the portfolio literature, most works maximize the expectation of some parametric concave utility representing the investor's attitude toward risk. The importance of the utility theory notwithstanding, it is not intuitive to determine a risk-averse parameter or even which utility function that suitably represents the investor's risk aversion. The latest approaches in the stochastic dynamic programming literature introduce risk aversion via time-consistent dynamic risk measures. The objective function is a recursive formulation of a one-period coherent risk measure (Shapiro et al. 2013; Philpott and de Matos 2012), generally the convex combination of expectation and conditional value-at-risk (CV@R) (Rockafellar and Uryasev 2000, 2002). The recursive model ensures time-consistent policies (Shapiro 2009) and has a suitable economic interpretation of a certainty equivalent (Rudloff et al. 2014). In practical applications, however, a decision maker must define the relative weights between expectation and CV@R (Rockafellar and Uryasev 2000, 2002) to represent his risk aversion, which is a non-intuitive user-defined risk-aversion parameter.

In this work, we develop and efficiently solve a time-consistent multi-asset stochastic dynamic asset allocation model motivated by the actual decision process in the financial market. Hedge funds hire managers to propose trading strategies that maximize expected returns, while risk departments impose constraints to strategies with a high level of risk. We focus our developments on risk-constrained models (Fernandes et al. 2016; Valladão et al. 2018), arguing that it is reasonable to assume that an investor knows how much he is willing to lose in a given period. We argue that determining a loss tolerance is much more intuitive than other risk aversion parameterizations in a dynamic setting. Our model assumes a discrete-state Markov process to capture the dynamics of common factors of asset returns and imposes one-period conditional CV@R constraints, ensuring a relatively complete recourse and computationally tractable time-consistent risk-averse model.

Our main contribution is to develop and solve a dynamic stochastic asset allocation model in a discrete-time finite horizon that has the following characteristics:

– Realism: multiple assets (high-dimensional space), transactional costs and time-dependent asset returns;
– Time consistency: expected return maximization with one-period CV@R constraints guarantees that planned decisions are actually going to be implemented in the future;
– Intuitive risk aversion parameterization: a user-defined CV@R limit, intuitively interpreted as the maximum percentage loss allowed in one period;
– Flexibility: straight-forward extensions to incorporate operational restrictions such as holdings and turnover constraints;
– Approximation framework for more complex price dynamics: the approximate solution obtained with a Markov chain surrogate is near-optimal and outperforms selected heuristic benchmarks.

– Computational tractability and solution quality: the solution of the SAA problem obtained via MSDDP is computationally tractable and near-optimal for the original (continuously distributed) problem even considering a high-dimensional state-space.

2 Theoretical background

Let us assume a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$, where $\mathcal{F} = \mathcal{F}_T$ and $\mathcal{F}_0 \subseteq \ldots \subseteq \mathcal{F}_T$, and a generic definition of a time-consistent dynamic stochastic programming model for asset allocation. At a given time (stage) t, an investor wishes to construct a portfolio $\mathbf{u_t} = (u_{0t}, \ldots, u_{Nt})^\top$ determining allocation on a risk-free asset[1] $u_{0,t}$ and on risky ones $u_{i,t}$ for $i \in \mathcal{A} = \{1, \ldots, N\}$, each of which has an uncertain return $r_{i,t}$, with a vector notation $\mathbf{r_t} = (r_{0t}, \ldots, r_{Nt})^\top$. At a given time t, we define the feasible set of the asset allocations $\mathcal{U}_t(\mathbf{u_{-1}})$ as a function of the previous allocation $\mathbf{u_{t-1}}$. We also define the terminal wealth $W_T(\mathbf{u_{T-1}})$ as a function of the last allocation and, for every $t = 1, \ldots, T - 1$, a one-period preference functional $\psi_t : L^\infty(\mathcal{F}_{t+1}) \to L^\infty(\mathcal{F}_t)$. Following Rudloff et al. (2014) and Shapiro (2009), we ensure time-consistent policies by defining the recursive setting

$$\max_{\mathbf{u_1} \in \mathcal{U}_1} \psi_1 \left[\max_{\mathbf{u_2} \in \mathcal{U}_2(\mathbf{u_1})} \psi_2 \left[\cdots \max_{\mathbf{u_{T-1}} \in \mathcal{U}_{T-1}(\mathbf{u_{T-2}})} \psi_{T-1} \left[W_T(\mathbf{u_{T-1}}) \right] \right] \right] \tag{1}$$

where u_t and r_t are $\mathcal{F}_t$–measurable.

This generic formulation accounts for different types of models, including maximization of expected utility, minimization of a time-consistent dynamic risk measure and a risk-neutral objective function with one-period risk constraints. Previous works show two modeling issues that pose a significant trade-off between realism and computational tractability: transaction costs and time dependence of asset returns. For the purpose of gaining intuition and keeping the paper self-contained, we review previous results for simplified models: Sect. 2.1 depicts a myopic portfolio model with no transaction costs, while Sect. 2.2 shows that SDDP efficiently solves a stage-wise independent model with transaction costs. Finally, in Sect. 2.3, we discuss available alternatives to accommodates temporal dependence on the SDDP framework.

2.1 No transaction cost model

Let us define the current wealth $W_t = (1 + \mathbf{r_t})^\top \mathbf{u_{t-1}}$ as the accrued value of previous allocation[2] and use it to reallocate the new portfolio, $\mathbf{1}^\top \mathbf{u_t} = W_t$. Hence, in (1), the feasible set would be $\mathcal{U}_t(\mathbf{u_{t-1}}) = \left\{ \mathbf{u_t} \in \mathbb{R}_+^{N+1} \mid \mathbf{1}^\top \mathbf{u_t} = (1 + \mathbf{r_t})^\top \mathbf{u_{t-1}} \right\}$ and the terminal wealth $W_T(\mathbf{u_{T-1}}) = (1 + \mathbf{r_T})^\top \mathbf{u_{T-1}}$. Now, assuming that short selling is not allowed, let us define the model with the dynamic programming equations for $t = T, \ldots, 0$

$$Q_t(W_t) = \max_{\mathbf{u_t} \geq 0} \left\{ \psi_t \left[Q_{t+1} \left((1 + \mathbf{r_{t+1}})^\top \mathbf{u_t} \right) \right] \mid \mathbf{1}^\top \mathbf{u_t} = W_t \right\} \tag{2}$$

where $Q_T(W_T) = W_T$. Traditional dynamic programming techniques efficiently solve no transaction cost portfolio problems when asset returns follow a low-dimensional factor model. For instance, (2) is computationally tractable if asset returns follow a one-factor model

---

[1] The risk-free asset is indexed by $i = 0$, and it has null returns for all $t = 2, \ldots, T$, i.e., $r_{0,t}(\omega) = 0$, $\forall \omega \in \Omega$.

[2] Where $\mathbf{1}$ is a column vector of ones with appropriate dimension.

$\log(1 + \mathbf{r_{t+1}}) = a_r + b_r f_t + \varepsilon_{t+1}$, where the scalar factor follows a simple autoregressive process $f_{t+1} = a_f + b_f f_t + \eta_{t+1}$, Brown and Smith (2011).

Following Blomvall and Shapiro (2006), if we consider a log-utility $u(w) = \log(w)$ and define $\psi_{T-1}[W] = \mathbb{E}[u(W)|\mathcal{F}_{T-1}]$ and $\psi_t[W] = \mathbb{E}[W|\mathcal{F}_t], \forall 1 \leq t \leq T - 2$, we obtain a myopic problem whose current decision $\mathbf{u_t}$ can be obtained as the solution of

$$\max_{\mathbf{u_t} \geq \mathbf{0}} \left\{ \psi_t \left[(\mathbf{1} + \mathbf{r_{t+1}})^\top \mathbf{u_t} \right] \mid \mathbf{1}^\top \mathbf{u_t} = W_t \right\} \tag{3}$$

a tractable two-stage problem only concerning financial gains of the following period $t + 1$. The additional assumption of stage-wise independent returns also leads to a myopic problem if one chooses a power utility $u(w) = w^\eta/\eta$ with $\eta \leq 1$ and for an alternative objective function based on a time-consistent dynamic risk measure, i.e., defining $\psi_t[W] = -\rho_t[W]$, where ρ_t is a conditional coherent risk measure; see Cheridito et al. (2006).

2.2 Transaction costs and stage-wise independent model

Under transaction costs, dynamic asset allocation is no longer myopic but still computationally tractable for stage-wise independence of asset returns. Although unreliable as an actual trading strategy, one can solve the stage-wise independent model using the stochastic dual dynamic programming (SDDP) algorithm (Kozmík and Morton 2015).

Let us assume stage-wise independence and reformulate the problem with an augmented state space for the value function $Q_t(\mathbf{x_t})$, where $\mathbf{x_t}$ is a vector that represents the amount of money in each asset right before buying and selling decisions at time t.

Note that we could extend our framework for any convex transactional cost function $f(\mathbf{x_t}, \mathbf{u_t})$ as in Brown and Smith (2011). For that purpose, we include an additional constraint to obtain the epigraph formulation

$$Q_t(\mathbf{x_t}) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \quad \psi_t \left[Q_{t+1}(\mathbf{R_{t+1}} \cdot \mathbf{u_t}) \right] \tag{4}$$

$$\text{s.t.} \quad u_{0,t} + \kappa = x_{0,t} \tag{5}$$

$$f(\mathbf{x_t}, \mathbf{u_t}) \leq \kappa \tag{6}$$

$$\mathbf{u_t} \geq 0, \tag{7}$$

which is, by definition, a convex problem. For computational efficiency, one can approximate a convex cost by a piecewise linear function to remain on a linear programming framework (Valladão et al. 2018).

For simplicity, we assume proportional costs. We respectively denote buying and selling (dropping) decision variables by $b_{i,t}, d_{i,t}, \forall i \in \mathcal{A}, t \in \{1, \ldots, T - 1\}$, with the vector notation $\mathbf{b_t} = (b_{1,t}, \ldots, b_{N,t})^\top$ and $\mathbf{d_t} = (d_{1,t}, \ldots, d_{N,t})^\top$. Then, we define the dynamic equations

$$Q_t(\mathbf{x_t}) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \quad \psi_t \left[Q_{t+1}(\mathbf{R_{t+1}} \cdot \mathbf{u_t}) \right] \tag{8}$$

$$\text{s.t.} \quad u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t} \tag{9}$$

$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A} \tag{10}$$

$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0, \tag{11}$$

where $Q_T(\mathbf{x_T}) = \mathbf{1}^\top \mathbf{x_T}$, $\mathbf{R_t} = \text{diag}(\mathbf{1} + \mathbf{r_{t+1}})$, $\mathbf{c} = c \cdot \mathbf{1}$, and c is the transaction cost proportion.
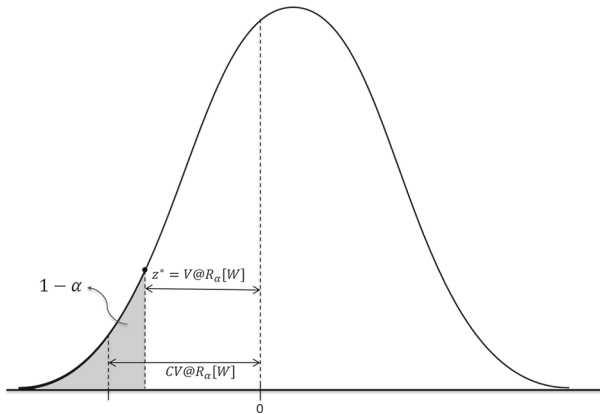
Fig. 1 $V@R_\alpha[W]$ and $CV@R_\alpha[W]$ of an uncertain wealth (W)

As before, the objective function (8) is a risk-adjusted terminal wealth. Constraint (9) ensures that the allocation in the risk-free asset (cash) before $(x_{0,t})$ and after $(u_{0,t})$ buying and selling decisions correctly accounts for the cash balance and transaction costs. Constraint (10) accounts for the reallocation of each risky asset, and (11) are non-negativity of buying and selling decisions making short selling and borrowing prohibition.

Considering stage-wise independence, the model (8–11) is efficiently solved by the standard SDDP algorithm, which iteratively constructs a piecewise linear outer approximation by sampling trial points of $\mathbf{x_t}$ and determining first order approximations, i.e., cutting planes. Given that the function $Q_t(\mathbf{x_t})$ is concave, the piecewise linear outer approximation provides an upper bound to the original problem. In particular, Philpott and de Matos (2012) and Shapiro (2011) propose the time-consistent risk-averse objective

$$\psi_t[Q_{t+1}] = (1-\lambda) \cdot \mathbb{E}[Q_{t+1}|\mathcal{F}_t] - \lambda \cdot CV@R_\alpha[Q_{t+1}|\mathcal{F}_t],$$

where, $Q_{t+1} = Q_{t+1}(\mathbf{R_{t+1}} \cdot \mathbf{u_t})$ and the conditional value-at-risk (CV@R) is defined as

$$CV@R_\alpha[Q_{t+1} \mid \mathcal{F}_t] = \inf_{z \in \mathbb{R}} \left\{ z + \frac{\mathbb{E}\left[\left((-Q_{t+1}) - z \right)^+ \big| \mathcal{F}_t \right]}{1 - \alpha} \right\},$$

with confidence level $\alpha \in (0, 1)$ and $x^+ = \max(x, 0)$. The CV@R concept is illustrated in Fig. 1.

For a risk-neutral formulation, i.e., $\lambda = 0$, a lower bound estimator can be obtained by simulating independent paths of portfolio returns (Shapiro 2011). The recursive risk-averse formulation, i.e., $\lambda > 0$, does not allow for a straightforward lower bound computation, considering a maximization problem. Although Kozmík and Morton (2015) propose a lower bound computation methodology based on importance sampling, most works use ad-hoc stopping criteria for solving such problems.

Notwithstanding Rudloff et al. (2014) intuitive certainty-equivalent interpretation of this recursive objective function, it is not clear how to determine λ to represent the preferences of an investor. Hence, this is a major obstacle for practical implementation.
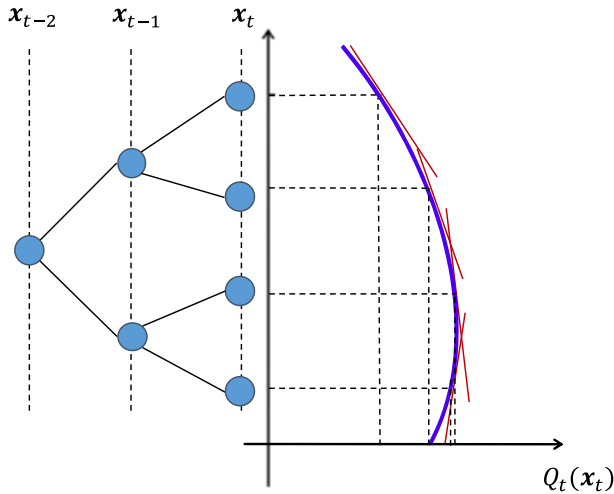
Fig. 2 Illustrative SDDP value function considering stage-wise independence

## 2.3 Stochastic dual dynamic programming and temporal dependence

The SDDP and other sampling-based methods assume stage-wise independence to avoid having multiple value functions per stage. Considering stage-wise independence, SDDP allows for cut sharing and consequently approximates only one value function at each time stage. For a generic time dependence, a decomposition algorithm, such as the L-shaped method, requires a differing value function for each realization of the stochastic process, which significantly increases the computational burden. Figure 2 illustrates an SDDP value function that considers stage-wise independence, while Fig. 3 illustrates differing value functions considering a generic time dependence. For the latter, one would need the full representation of the tree to use Benders-based decomposition algorithms, such as L-shaped, to solve a medium-sized problems. SDDP handles large-scale problems, while other methods based on the full representation of the tree become computationally intractable.

The SDDP framework accommodates some specific types of time dependence, such as a linear autoregressive process embedded in the right-hand-side of linear constraint,[3] and, more generally, for a discrete-state Markov process. For instance, in the hydrothermal model, the state-space includes past rivers' inflows changing the right-hand side accordingly to model a linear time-dependence as in (Infanger and Morton 1996; Pereira and Pinto 1985). For a more generic class of problems, Mo et al. (2001) and Philpott and de Matos (2012) modify the model using the Markov chains preserving convexity, consequently ensuring optimality and computational tractability.

In the dynamic asset allocation problem, the stochastic returns are not right-hand-side coefficients and have a nonlinear dependence (Cont 2001; Morettin and Toloi 2006). For a generic time dependence, full tree representation multistage models as in Valladão et al. (2014) may solve a medium-size dynamic portfolio optimization, but it is not computationally tractable for a large-scale problem including transaction costs, several time stages and multiple assets. In our work, we assume a discrete-state Markov process, where given a state,

---

[3] Note that any time dependent return influences the left-hand side of a balance constraint inducing a non-convex, therefore intractable, augmented value function.

Fig. 3 Illustrative value functions considering generic time dependence
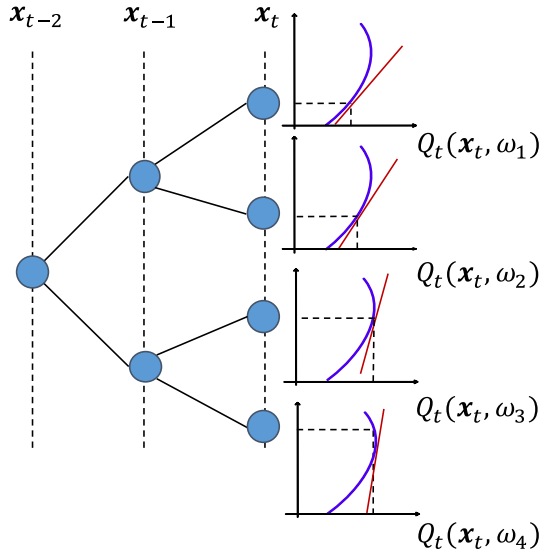
the probability distribution does not depend on past asset returns. In this case, we have a different value function for each time stage and for each Markov state of the system. In the following section, we propose a sequence of risk-constrained models building up complexity and realism with the following assumptions: (i) no transaction costs and stage-wise independence; (ii) transaction costs and stage-wise independence; and (iii) transaction costs and Markov time-dependent asset returns. For the latter, we adapt SDDP based on Mo et al. (2001), Philpott and de Matos (2012) to account for a discrete-state Markov model and solve a realistic dynamic asset allocation problem within reasonable computation time. To the best of our knowledge, no work in the literature has addressed a time-consistent risk-constrained dynamic portfolio optimization with transactional costs and time-dependent returns.

3 Risk-constrained stochastic dynamic asset allocation models

In this section, we propose a class of asset allocation models motivated by the actual decision process in financial markets. Hedge funds hire managers to propose trading strategies that maximize expected returns, while risk departments impose constraints to strategies with a high level of risk. We focus our developments on risk-constrained models, arguing that it is reasonable to assume that an investor knows how much he is willing to lose in a given period.

Let us start by assuming time (stage-wise) independence and consequently obtaining $\mathbb{E}[W|\mathcal{F}_t] = \mathbb{E}[W]$ and $\rho[W|\mathcal{F}_t] = \rho[W]$, for a random gain W. Considering no transaction costs, we define the dynamic programming equations for $t = 0, \ldots, T - 1$

$$Q_t(W_t) = \max_{\mathbf{u_t}} \ \mathbb{E}\left[Q_{t+1}\left((\mathbf{1} + \mathbf{r_{t+1}})^\top \mathbf{u_t}\right)\right] \tag{12}$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}^\top u_t}\right] \le \gamma W_t \tag{13}$$

$$\mathbf{1}^\top \mathbf{u_t} = W_t \tag{14}$$

$$\mathbf{u_t} \ge 0, \tag{15}$$

where $Q_T(W_T) = W_T$. The risk constraint (13) limits percentage portfolio loss by γ, a loss-tolerance determined by the investor. In other words, the risk constraint (13) is equivalent to

$$\rho\left[W_{t+1} - W_t\right] \leq \gamma W_t,$$

where the future wealth is defined as $W_{t+1} = (1 + \mathbf{r_{t+1}})^\top \mathbf{u_t}$ and the current wealth as $W_t = \mathbf{1}^\top \mathbf{u_t}$.

We assume a coherent risk measure $\rho : L^\infty(\mathcal{F}) \rightarrow \mathbb{R}$ with the following properties (Artzner et al. 1999):

- Monotonicity: $\rho(X) \leq \rho(Y)$, for all $X, Y \in L^\infty(\mathcal{F})$, such that $X \geq Y$;
- Translation invariance: $\rho(X + m) = \rho(X) - m$, for all $X \in L^\infty(\mathcal{F})$ and $m \in \mathbb{R}$;
- Positive homogeneity: $\rho(\lambda X) = \lambda \rho(X)$, for all $X \in L^\infty(\mathcal{F})$, $\lambda \geq 0$;
- Subadditivity: $\rho(X + Y) \leq \rho(X) + \rho(Y)$, for all $X, Y \in L^\infty(\mathcal{F})$.

It is important to note that, in general, feasibility is not guaranteed for risk-constrained models. In particular for the SDDP framework, feasibility cuts generate numerical instability and Lagrangean relaxations depend heavily on appropriate choice of multipliers. For a straightforward application of the standard SDDP, risk-constrained models must have relative complete recourse, i.e., there must exist a feasible solution for any possible attainable states of the system.

Given some mild assumptions, the proposed risk-constrained model without transaction costs (12–15) has relatively complete recourse , i.e., we ensure feasibility for any attainable (non-negative) current wealth W_t. Assuming $P(r_t < -1) = 0$, a feasible policy ensures $W_t \geq 0$, since we include short selling and borrowing prohibition. Indeed, a full allocation in the risk-free asset ensures feasibility regarding the risk constraint for a non-negative γ.

Proposition 1 *If* $\mathbb{P}(\{\omega \in \Omega \mid \mathbf{r_{t+1}}(\omega) < -\mathbf{1}\}) = 0$ *and* $\mathbb{P}(\{\omega \in \Omega \mid r_{0,t+1}(\omega) = 0\}) = 1, \forall t \in \{0, \ldots, T-1\}$*, then* (12–15) *has relatively complete recourse for* $\gamma \geq 0$.

Proof of Proposition 1 Assuming $\mathbb{P}(\{\omega \in \Omega \mid \mathbf{r_{t+1}}(\omega) < -\mathbf{1}\}) = 0$ and the short selling and borrowing prohibition $\mathbf{u_t} \geq 0$, then $\mathbb{P}(\{\omega \in \Omega \mid W_{t+1}(\omega) \geq 0\}) = \mathbb{P}(\{\omega \in \Omega \mid (1 + \mathbf{r_{t+1}}(\omega))^\top \mathbf{u_t} \geq 0\}) = 1$. Given a feasible (non-negative) current wealth, $W_t \geq 0$, this model has relatively complete recourse for any $\gamma \geq 0$ since a risk-free allocation, $\mathbf{u_t^{rf}} = (W_t, 0, \ldots, 0)^\top$, is always feasible. Indeed, for $\gamma \geq 0$, $\rho_t\left[\mathbf{r_{t+1}}^\top \mathbf{u_t^{rf}}\right] = 0 \leq \gamma W_t$ given that $\mathbb{P}(\{\omega \in \Omega \mid r_{0,t}(\omega) = 0\}) = 1$. □

Additionally, (12–15) has a myopic solution because the first-stage decision also solves its one-period counterpart problem.

Proposition 2 *The problem* (12–15) *is myopic, i.e., it has the same solution of*

$$\begin{aligned}
\mathcal{Q}_t(W_t) = \max_{\mathbf{u_t}} \quad & \mathbb{E}\left[(1 + \mathbf{r_{t+1}})^\top \mathbf{u_t}\right] \\
s.t. \quad & \rho\left[\mathbf{r_{t+1}}^\top \mathbf{u_t}\right] \leq \gamma W_t \\
& \mathbf{1}^\top \mathbf{u_t} = W_t \\
& \mathbf{u_t} \geq 0,
\end{aligned} \tag{16}$$

Proof Proof of Proposition 2. By definition, $Q_T(W_T) = \mathcal{Q}_T(W_T) = W_T$, and consequently, $Q_T(W_T) = W_T \cdot \mathcal{Q}_T(1)$. Employing the inductive hypothesis $Q_{t+1}(W_{t+1}) = W_{t+1} \prod_{\tau=t+1}^{T} \mathcal{Q}_\tau(1)$ in (12–15), we would have

$$Q_t(W_t) = \left(\prod_{\tau=t+1}^{T} \mathcal{Q}_\tau(1) \right) \cdot \max_{\mathbf{u_t}} \; \mathbb{E}\left[\left((1 + \mathbf{r_{t+1}})^\top \mathbf{u_t} \right) \right]$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}^\top u_t} \right] \leq \gamma W_t$$

$$\mathbf{1}^\top \mathbf{u_t} = W_t$$

$$\mathbf{u_t} \geq 0.$$

Note that if the inductive hypothesis is true, the solution of the 2-stage problem (16) is also the solution of the original problem (12–15).

It is straightforward that inductive hypothesis holds for $W_t = 0$, since $Q_t(0) = 0 \times \prod_{\tau=t}^{T} \mathcal{Q}_\tau(1) = 0$. To prove that the inductive hypothesis holds for $W_t > 0$, we divide all constraints by the positive constant W_t and redefine decision variables as percentage allocation $\tilde{\mathbf{u}}_\mathbf{t} = \mathbf{u_t}/W_t$ to obtain the equivalent problem

$$Q_t(W_t) = W_t \left(\prod_{\tau=t+1}^{T} \mathcal{Q}_\tau(1) \right) \cdot \max_{\tilde{\mathbf{u}}_\mathbf{t}} \; \mathbb{E}\left[(1 + \mathbf{r_{t+1}})^\top \tilde{\mathbf{u}}_\mathbf{t} \right]$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}^\top} \tilde{\mathbf{u}}_\mathbf{t} \right] \leq \gamma$$

$$\mathbf{1}^\top \tilde{\mathbf{u}}_\mathbf{t} = 1$$

$$\tilde{\mathbf{u}}_\mathbf{t} \geq 0.$$

Then, we have that $Q_t(W_t) = W_t \prod_{\tau=t}^{T} \mathcal{Q}_\tau(1)$.   □

Now, let us take one step toward a less unrealistic asset allocation model. For simplicity, we assume proportional transaction costs, even though we could consider a more general convex transaction cost function. As in Sect. 2.2, we define the state space as a vector with the amount of money allocated in each asset right before buying and selling decisions at time t. We redefine the risk-constrained problem with transaction costs and stage-wise independent returns for $t = 0, \ldots, T - 1$,

$$Q_t(\mathbf{x_t}) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \; \mathbb{E}\left[Q_{t+1}\left(\mathbf{R_{t+1}} \cdot \mathbf{u_t} \right) \right] \tag{17}$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}^\top u_t} \right] + \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t}) \leq \gamma \left(\mathbf{1}^\top \mathbf{x_t} \right) \tag{18}$$

$$u_{0,t} + (1 + \mathbf{c})^\top \mathbf{b_t} - (1 - \mathbf{c})^\top \mathbf{d_t} = x_{0,t} \tag{19}$$

$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A} \tag{20}$$

$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0, \tag{21}$$

where $Q_T(\mathbf{x_T}) = \mathbf{1}^\top \mathbf{x_T}$, $\mathbf{R_t} = \text{diag}(1 + \mathbf{r_{t+1}})$ and $\mathbf{c} = c \cdot \mathbf{1}$.

To economically motivate risk constraint (18), we start with the same interpretation used in the no-transaction cost model, i.e., the portfolio percentage loss is limited as by γ,

$$\rho\left[W_{t+1} - W_t \right] \leq \gamma W_t.$$

Defining the current wealth as $W_t = \mathbf{1}^\top \mathbf{x_t}$ and the future wealth as $W_{t+1} = (1 + \mathbf{r_{t+1}})^\top \mathbf{u_t}$, the risk constraint is equivalent to

$$\rho\left[\mathbf{r_{t+1}^\top u_t} - (\mathbf{1}^\top \mathbf{x_t} - \mathbf{1}^\top \mathbf{u_t}) \right] \leq \gamma \left(\mathbf{1}^\top \mathbf{x_t} \right).$$

By summing up constraints (19) and (20), we obtain $\mathbf{1}^\top \mathbf{x_t} - \mathbf{1}^\top \mathbf{u_t} = \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t})$ and, consequently,

$$\rho\left[\mathbf{r}_{\mathbf{t+1}}^\top \mathbf{u_t} - \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t})\right] \leq \gamma \left(\mathbf{1}^\top \mathbf{x_t}\right).$$

Finally, we use the translation invariance property to obtain the final expression as in (18).

As in the previous subsection, feasibility is an important issue for an efficient SDDP implementation. That said, (17–21) has relative complete recourse when the maximum percentage loss is at least the transaction cost rate.

Proposition 3 *If* $\mathbb{P}(\{\omega \in \Omega \mid \mathbf{r}_{t+1}(\omega) < -\mathbf{1}\}) = 0$ *and* $\mathbb{P}(\{\omega \in \Omega \mid r_{0,t+1}(\omega) = 0\}) = 1, \forall t \in \{0, \ldots, T-1\}$, *problem* (17–21) *has relatively complete recourse for* $\gamma \geq c$.

Proof Proposition 3 Assuming $\mathbb{P}(\{\omega \in \Omega \mid \mathbf{r}_{t+1}(\omega) < -\mathbf{1}\}) = 0$ and the short selling and borrowing prohibition $\mathbf{u_t} \geq 0$, then $\mathbb{P}(\{\omega \in \Omega \mid \mathbf{R}_{t+1}(\omega) \cdot \mathbf{u_t} \geq \mathbf{0}\}) = 1$, given that $\mathbf{R_t} = \mathrm{diag}(\mathbf{1} + \mathbf{r_{t+1}})$. First, we show that there exists $\mathbf{b_t}, \mathbf{d_t}$ such that $\mathbf{1}^\top (\mathbf{b_t} + \mathbf{d_t}) \leq \mathbf{1}^\top \mathbf{x_t}$ and consequently $\mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t}) \leq \mathbf{c}^\top \mathbf{x_t}$. If all risky assets are sold, i.e., $\mathbf{b_t} = \mathbf{0}$ and $d_{i,t} = x_{i,t}, \forall i \in \mathcal{A}$. Then,

$$\mathbf{1}^\top (\mathbf{b_t} + \mathbf{d_t}) = \sum_{i \in \mathcal{A}} x_{i,t} \leq \sum_{i \in \mathcal{A} \cup \{0\}} x_{i,t} = \mathbf{1}^\top \mathbf{x_t}$$

Given that we sell all risky assets, $\rho\left[\mathbf{r}_{\mathbf{t+1}}^\top \mathbf{u_t}\right] = 0$. Then, if $\gamma \geq c$,

$$\rho\left[\mathbf{r}_{\mathbf{t+1}}^\top \mathbf{u_t}\right] + \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t}) = \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t}) \leq c \left(\mathbf{1}^\top \mathbf{x_t}\right) \leq \gamma \left(\mathbf{1}^\top \mathbf{x_t}\right). \qquad \square$$

Finally, our last step toward a realistic model incorporates price dynamics by assuming that asset returns follow a discrete-state Markov model. Let us assume a Markov asset return with discrete state space $\mathcal{K}$, where the multivariate probability distribution of $\mathbf{r_t}$ given the Markov state K_t at time t does not depend on past returns, i.e., $r_{t-1}, \ldots, r_1$. The Markov state K_t could be interpreted as the financial market situation. For instance, in Fig. 4, we have three Markov states (bull, neutral or bear market) with three differing probability distributions of asset returns.

For notation simplicity, we denote the state transition probability as

$$P_{k|j} = \frac{\mathbb{P}\left(\{\omega \in \Omega | K_t(\omega) = j, K_{t+1}(\omega) = k\}\right)}{\mathbb{P}(\{\omega \in \Omega | K_t(\omega) = j\})},$$

and, for a given time $t = 0, \ldots, T-1$ and $j \in \mathcal{K}$, we define a different value function for a given state $K_t = j$ via dynamic programming equations

$$
\begin{aligned}
Q_t^j(\mathbf{x_t}) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \quad & \sum_{k \in \mathcal{K}} \mathbb{E}\left[Q_{t+1}^k\left(\mathbf{R_{t+1}} \cdot \mathbf{u_t}\right) \mid K_{t+1} = k\right] P_{k|j} \\
\text{s.t.} \quad & \rho\left[\mathbf{r}_{\mathbf{t+1}}^\top \mathbf{u_t} \mid K_t = j\right] + \mathbf{c}^\top (\mathbf{b_t} + \mathbf{d_t}) \leq \gamma \left(\mathbf{1}^\top \mathbf{x_t}\right) \\
& u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t} \\
& u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A} \\
& \mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0,
\end{aligned}
\tag{22}
$$

where $Q_T^k(\mathbf{x_T}) = \mathbf{1}^\top \mathbf{x_T}$, $\mathbf{R_t} = \mathrm{diag}(\mathbf{1} + \mathbf{r_{t+1}})$, and $\mathbf{c} = c \cdot \mathbf{1}$.

The objective function is a weighted average of $t+1$ value functions, as illustrated in Fig. 5. As a compromise between stage-wise independence, illustrated in Fig. 2, and a generic
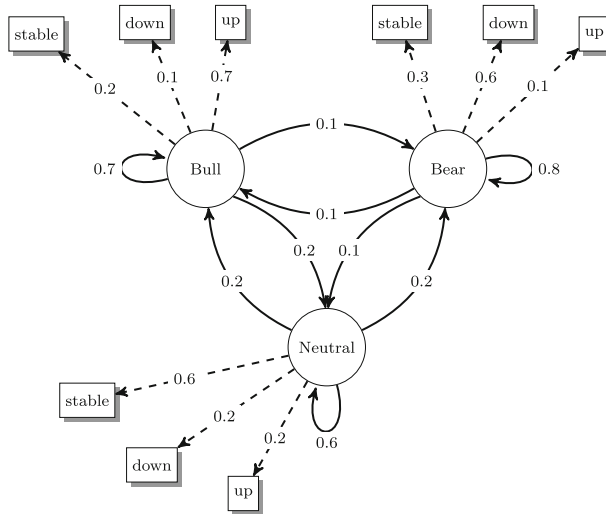
Fig. 4 Illustrative Markov asset return model

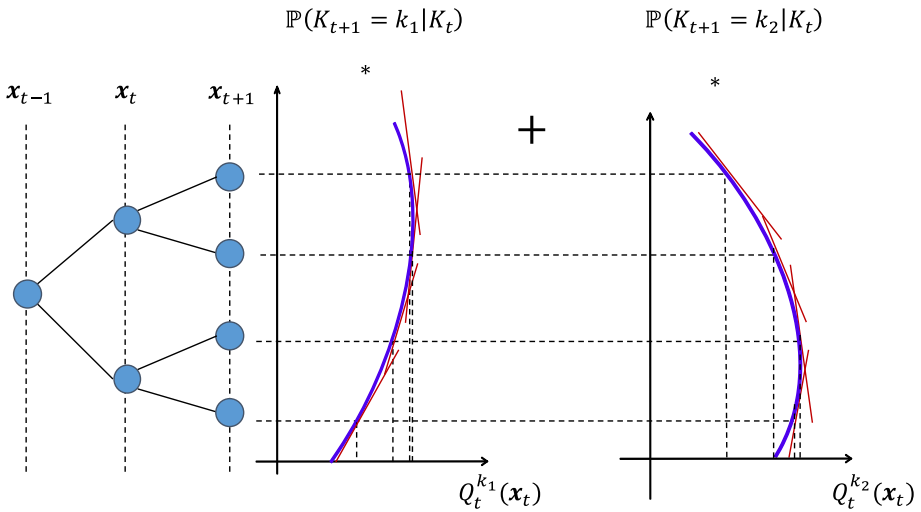$$\mathbb{P}(K_{t+1} = k_1 | K_t) \qquad\qquad \mathbb{P}(K_{t+1} = k_2 | K_t)$$

Fig. 5 Illustrative weighted average of value functions for Markovian asset return model

time-dependent model, illustrated in Fig. 3, the Markov model remains computationally tractable for a reasonably small number of Markov states and introduces the effect of price dynamics in the optimal asset allocation policy.

For a practical and efficient implementation of the SDDP solution algorithm, let us develop a problem equivalent to (22) that explicitly considers intermediate earnings in the objective function. This modeling choice mitigates the effects of a poor approximation of the future value function in early iterations of the SDDP solution algorithm. Given a sequence of decisions, we recursively represent the terminal wealth $W_T = \mathbf{1}^\top \mathbf{x_T}$ at time $t = 0$ as $\mathbf{1}^\top \mathbf{x_T} = \mathbf{1}^\top \mathbf{x_0} + \sum_{\tau=t}^{T-1} \left(-\mathbf{c}^\top (\mathbf{b}_\tau + \mathbf{d}_\tau) + \mathbf{r}_{\tau+1}^\top \mathbf{u}_\tau \right)$. Considering intermediate earnings, we define the value function

$$V_t^j(\mathbf{x_t}) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} -\mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) + \sum_{k \in \mathcal{K}} \mathbb{E}\left[\mathbf{r_{t+1}}^\top \mathbf{u_t} + V_{t+1}^k(\mathbf{R_{t+1}} \cdot \mathbf{u_t}) \mid K_{t+1} = k\right] P_{k|j}$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}}^\top \mathbf{u_t} \mid K_t = j\right] + \mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) \leq \gamma\left(\mathbf{1}^\top \mathbf{x_t}\right)$$

$$u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t}$$

$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A}$$

$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0. \tag{23}$$

Note that we neglect $\mathbf{1}^\top \mathbf{x_0}$ in the objective function since it is a constant at time $t = 0$. Additionally, we assume ρ to be the CV@R with confidence level α conditioned to the current Markov state $K_t = j$,

$$\rho\left[W \mid K_t = j\right] = \inf_{z \in \mathbb{R}}\left\{z + \frac{\mathbb{E}\left[((-W) - z)^+ \mid K_t = j\right]}{1 - \alpha}\right\},$$

which is a coherent risk measure with a suitable economic interpretation easily represented in a linear stochastic programming problem.

For computational tractability, let us assume a discrete number of asset return scenarios and associate probabilities for each given Markov state. Given $K_{t+1} = k$, we denote for each scenario $s \in \mathcal{S}_k$, the asset return vector $\mathbf{r_{t+1}}(s) = (r_{0,t+1}(s), \ldots, r_{N,t+1}(s))^\top$ and the matrix $\mathbf{R_{t+1}}(s) = \text{diag}(\mathbf{1} + \mathbf{r_{t+1}}(s))$. Moreover, we denote

$$p_{s|k} = \frac{\mathbb{P}(\{\omega \in \Omega \mid \mathbf{r_{t+1}}(\omega) = \mathbf{r_{t+1}}(s), \ K_{t+1}(\omega) = k\})}{\mathbb{P}(\{\omega \in \Omega \mid K_{t+1}(\omega) = k\})}$$

as the conditional return probability and define for all $s \in \mathcal{S}_k$ auxiliary decision variables y_s to represent CV@R in the deterministic equivalent linear dynamic stochastic programming problem

$$V_t^j(\mathbf{x_t}) = \max_{z, \mathbf{y}, \mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} -\mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) + \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}_k}\left(\mathbf{r_{t+1}}(s)^\top \mathbf{u_t} + V_{t+1}^k(\mathbf{R_{t+1}}(s) \cdot \mathbf{u_t})\right) P_{k|j} \, p_{s|k}$$

$$\text{s.t.} \quad z + (1 - \alpha)^{-1} \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}_k} y_s \, P_{k|j} \, p_{s|k} + \mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) \leq \gamma(\mathbf{1}^\top \mathbf{x_t})$$

$$u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t}$$

$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t} \quad \forall i \in \mathcal{A}$$

$$\mathbf{r_{t+1}}(s)^\top \mathbf{u_t} + z + y_s \geq 0, \quad \forall k \in \mathcal{K}, s \in \mathcal{S}_k$$

$$y_s \geq 0, \quad \forall k \in \mathcal{K}, s \in \mathcal{S}_k$$

$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0 \tag{24}$$

We argue that the dynamic asset allocation model (24) is: (i) realistic since it considers multiple assets, transactional costs and Markov-dependent asset returns; (ii) time consistent since all planned decisions are actually going to be implemented in the future; (iii) intuitive since its risk averse parameter can be interpreted as the maximum percentage loss allowed in one period; (iv) flexible since additional constraints such as holdings and turnover restrictions can be incorporated by a set of linear inequalities; (v) computationally tractable since it can be efficiently solved by the MSDDP algorithm.

4 Adapting MSDDP for the risk constrained model

Most works in the SDDP literature assume a hazard-decision information structure where decisions are made under perfect information for the current period but under uncertainty for subsequent time stages. For instance, the hydrothermal operation planning problem, first presented by Pereira and Pinto (1985), previous works assume that current operation considers known river inflows for the first month but unknown inflows for subsequent ones. This structure allows for SDDP subproblems to be solved separately for each considered scenario. For our proposed risk-constrained portfolio model however, we no longer can explore the hazard-decision information structure since the conditional (one-period) risk constraint couples all scenarios into a unique sub-problem. For this reason, we adapt the MSDDP proposed by Mo et al. (2001); Philpott and de Matos (2012) to a decision-hazard structure where each SDDP subproblem is a single two-stage stochastic programming model.

Let us denote $\mathfrak{V}_t^j(\cdot)$ as the current approximation of the value function $V_t^j(\cdot)$. The approximate function $\mathfrak{V}_t^j(\cdot)$ is defined as the minimum of a set of cuttings planes. For a given $t \in \{T-1, \ldots, 1\}$ in the backward step of SDDP algorithm, see Pereira and Pinto (1985), Shapiro (2011) for details, we solve the approximate SAA problem

$$
\mathfrak{V}_t^j(\mathbf{x_t}) = \max_{z, \mathbf{y}, \mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} -\mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) + \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}_k} \left(\mathbf{r_{t+1}}(s)^\top \mathbf{u_t} + \mathfrak{V}_{t+1}^k(\mathbf{R_{t+1}}(s) \cdot \mathbf{u_t}) \right) P_{k|j} \, p_{s|k}
$$

$$
\text{s.t.} \quad z + (1-\alpha)^{-1} \sum_{k \in \mathcal{K}} \sum_{s \in \mathcal{S}_k} y_s \, P_{k|j} \, p_{s|k} + \mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) \leq \gamma(\mathbf{1}^\top \mathbf{x_t}) \qquad : \eta_t^j
$$

$$
u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t} \qquad\qquad\qquad\qquad : \pi_{0,t}^j
$$

$$
u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t} \quad \forall i \in \mathcal{A} \qquad\qquad\qquad\qquad\qquad : \pi_{i,t}^j
$$

$$
\mathbf{r_{t+1}}(s)^\top \mathbf{u_t} + z + y_s \geq 0, \quad \forall k \in \mathcal{K}, s \in \mathcal{S}_k
$$

$$
y_s \geq 0, \quad \forall k \in \mathcal{K}, s \in \mathcal{S}_k
$$

$$
\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0
$$

to obtain: the dual vector $\pi_t^j = \{\pi_{0,t}^j, \ldots, \pi_{N,t}^j\}$ comprised of optimal Lagrange multipliers for each asset balance constraint; the dual variable η_t^j for the risk constraint and the objective value $\mathfrak{V}_t^j(\mathbf{x_t})$. Then, we define the linear approximation of the future value function $\mathfrak{V}_t^j(\mathbf{x_t})$ as

$$
l_t^j(\mathbf{x_t}) := \mathfrak{V}_t^j(\hat{\mathbf{x}}_t) + (\pi_t^j + \mathbf{1}\gamma \eta_t^j)^\top (\mathbf{x_t} - \hat{x}_t)
$$

and update the current approximation of the value function, i.e.,

$$
\mathfrak{V}_t^j(\cdot) \leftarrow \min \left(\mathfrak{V}_t^j(\cdot), l_t^j(\cdot) \right).
$$

We propose the following algorithm of the adapted MSDDP:

Algorithm 1 Adapted MSDDP for decision hazard information structure

---

Require: K_0, $\mathbf{x}_0$ and $\{\mathfrak{V}_t^k\}_{t=1,k=1}^{T-1,K}$ (Init. future value function)

Initialize: $UB \leftarrow \mathfrak{V}_0^{K_0}$, $LB \leftarrow 0$, $GAP = \frac{UB-LB}{UB}$

while $GAP > \epsilon$ **do**

    Generate a sample path: $\{\mathbf{r}_t, K_t\}_{t=1}^T$

    $\{\hat{\mathbf{x}}_t\}_{t=1}^T \leftarrow$ **Forward Step**

    $\{\mathfrak{V}_t^j\}_{t=1,j=1}^{T,K} \leftarrow$ **Backward Step**

    $UB \leftarrow \mathfrak{V}_0^{K_0} + \mathbf{1}^\top \mathbf{x}_0$

    $LB \leftarrow$ **Lower Bound**

end while

The Forward and Backward Steps are depicted as follows:

Algorithm 2 Forward Step

Require: K_0, $\mathbf{x}_0$, a sample path $\{\mathbf{r}_t, K_t\}_{t=1}^T$ and $\{\mathfrak{V}_t^k\}_{t=1,k=1}^{T-1,K}$

for $t \in \{0, \ldots, T-1\}$ **do**

    $u_t^* \leftarrow$ solution of $\mathfrak{V}_t^{K_t}$

    $\mathbf{R_{t+1}} \leftarrow \text{diag}(\mathbf{r_{t+1}})$

    $\hat{\mathbf{x}}_{t+1} \leftarrow \mathbf{R_{t+1}} \cdot \mathbf{u_t^*}$

end for

Return $\{\hat{\mathbf{x}}_t\}_{t=1}^T$

Algorithm 3 Backward Step

Require: $\{\hat{\mathbf{x}}_t\}_{t=1}^T$ and $\{\mathfrak{V}_t^k\}_{t=1,k=1}^{T-1,K}$

for $t \in \{T-1, \ldots, 1\}$ **do**

    for $j \in \{1, \ldots, K\}$ **do**

        $l_t^j(\mathbf{x_t}) := \mathfrak{V}_t^j(\hat{\mathbf{x}}_t) + (\pi_t^j + \mathbf{1}\gamma\eta_t^j)^\top(\mathbf{x_t} - \hat{\mathbf{x}}_t)$

        $\mathfrak{V}_t^j(\cdot) \leftarrow \min\left(\mathfrak{V}_t^j(\cdot), l_t^j(\cdot)\right)$

    end for

end for

Return $\{\mathfrak{V}_t^k\}_{t=1,k=1}^{T-1,K}$

The direct approach to compute a statistical lower bound is to simulate the optimal policy to obtain different realizations of the terminal wealth. To reduce variability however, we propose an alternative method for lower bound computation. Essentially we simulate the optimal policy and accumulate the current (one-period) expected return conditioned to the state of the system, see details in Algorithm 4.

Algorithm 4 Lower Bound evaluation

Require: $\mathbf{x_0}$ and $\{\mathfrak{V}_t^k\}_{t=1,k=1}^{T-1,K}$
  for $s \in \{1, \ldots, S_{LB}\}$ **do**
    $W(s) \leftarrow \mathbf{1}^\top \mathbf{x_0}$
    Generate a sample path: $\{\mathbf{r}_t, K_t\}_{t=1}^T$
    for $t \in \{0, \ldots, T-1\}$ **do**
      $\mathbf{u_t^*} \leftarrow$ solution of the problem $\mathfrak{V}_t^{K_t}(\mathbf{x_t})$
      $\mathbf{R_{t+1}} \leftarrow \text{diag}(\mathbf{r_{t+1}})$
      $\mathbf{x_{t+1}} \leftarrow \mathbf{R_{t+1}} \cdot \mathbf{u_t^*}$
      $W(s) \leftarrow W(s) - c^\top(\mathbf{b_t} + \mathbf{d_t}) + \sum_{k \in \mathcal{K}} \sum_{s_b \in \mathcal{S}_k} \left(\mathbf{r_{t+1}}(s_b)^\top \mathbf{u_t^*}\right) P_{k|j=K_t} \, p_{s_b|k}$
    end for
  end for
  $\overline{W} = \frac{1}{S_{LB}} \sum_{s=1}^{S_{LB}} W(s)$
  $\sigma_W^2 = \frac{1}{S_{LB}-1} \sum_{s=1}^{S_{LB}} \left(W(s) - \overline{W}\right)^2$
  $LB = \overline{W} - Z_{90\%} \cdot (\sigma_W / \sqrt{S_{LB}})$

Given that $Z_{90\%}$ denotes the 90% quantile of a standard Normal distribution, gap is constructed using the lower limit of the lower bound. This imposes a probabilistic guarantee of 90% that the gap of the SAA problem is smaller than or equal to 1%.

5 Empirical analysis

In this section, our objective is to compare the performance of the proposed risk-constrained time-consistent dynamic asset allocation model against selected benchmark strategies considering: (i) an illustrative problem with 3 assets and 1 factor with an autoregressive dynamic; (ii) a high-dimensional problem with 100 assets and 5 factors following a discrete Markov chain. We use a discrete state Markov stochastic dual dynamic programming (MSDDP) algorithm to solve problems considering transactional costs and time dependence. We assume factor models for asset returns and argue that our framework is fairly general even when factor dynamics are not theoretically defined by a discrete Markov chain. For low-dimensional time-dependent factors, we can estimate a Markov chain that represents either a given stochastic process or historical data.

For all case studies, we estimate using simulations or historical data a factor model for asset returns while representing factor dynamics by a Gaussian mixture in the hidden Markov model (HMM) scheme.[4] Rydén et al. (1998) show that modeling the financial time series as a Gaussian mixture, as illustrated in Fig. 6, according to the states of an unobserved Markov chain reproduces most of the stylized facts for daily return series demonstrated by Granger and Ding (1994). Thus, HMM is a natural and fairly general methodology for modeling financial time series (Hassan and Nath 2005; Elliott and Siu 2014; Elliott and Van der Hoek 1997; Mamon and Elliott 2007, 2014).

We divide our experiments into two case studies: in the first case study, we approximate a known one-factor dynamic by a Markov chain, solve a 3-asset problem and empirically test its in-sample performance. In the second case study, we consider a unknown 5-factor model (Fama and French 2015, 2016) whose time dependence is extracted from historical data to

---

[4] To fit HMM in the MSDDP framework, we assume, as an approximation, that the most probable Markov state is actually observable.
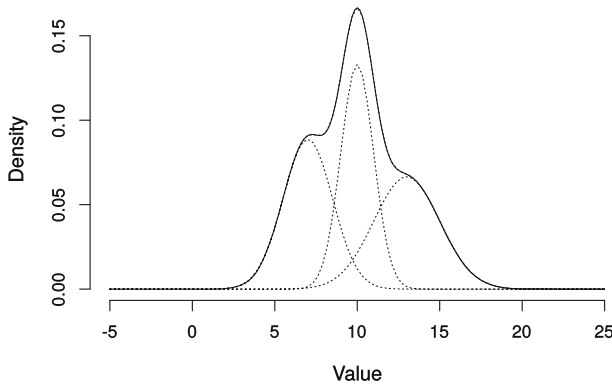
Fig. 6 An illustrative example of a Gaussian mixture model

solve a 100-asset problem ensuring with 90% probability an optimality gap smaller than or equal to 1%.

5.1 Illustrative 3-asset case study

Following Brown and Smith (2011), we consider a problem with a 12-month horizon ($T = 12$), three risky assets where uncertain monthly returns follow a factor model with its single (uncertain) factor $\tilde{f}_t$ as an autoregressive process. Objectively,

$$\log(1 + \mathbf{r_{t+1}}) = \mathbf{a_r} + \mathbf{b_r}\tilde{f}_t + \varepsilon_{t+1} \quad \text{and} \quad \tilde{f}_{t+1} = a_f + b_f\tilde{f}_t + \eta_{t+1}, \quad (25)$$

where (ε_t, η_t) are independent and identically distributed multivariate normal distribution with zero mean and covariance matrix $\Sigma_{\varepsilon\eta}$. With no transactional costs, the stochastic dynamic program

$$
\begin{aligned}
Q_t(W_t, f_t) = \max_{\mathbf{u_t}} \quad & \mathbb{E}\left[Q_{t+1}\left((1 + \mathbf{r_{t+1}})^\top \mathbf{u_t}, \tilde{f}_{t+1}\right)\middle|\tilde{f}_t = f_t\right] \\
\text{s.t.} \quad & \rho\left[\mathbf{r_{t+1}}^\top \mathbf{u_t}\middle|\tilde{f}_t = f_t\right] \le \gamma W_t \\
& \mathbf{1}^\top \mathbf{u_t} = W_t \\
& \mathbf{u_t} \ge 0.
\end{aligned}
$$

can be accurately approximated by its sample average approximation (SAA) counterpart and efficiently solved using approximate dynamic programming techniques due to a low-dimensional state space (Powell 2011).We use positive homogeneity, i.e., $Q_t(W_t, f_t) = W_t \cdot Q_t(1, f_t)$, simulate S realizations of (ε_t, η_t) denoted by $(\varepsilon_t(s), \eta_t(s))$, $\forall s = 1, \ldots, S$ and consider a CV@R with probability level α to solve the deterministic equivalent SAA problem

$$
\begin{aligned}
Q_t(W_t, f_t) = \max_{z,,\mathbf{y},\mathbf{u_t}} \quad & S^{-1}\sum_{s=1}^{S}\left((1 + \mathbf{r_{t+1}}(s))^\top \mathbf{u_t} \cdot Q_{t+1}(1, f_{t+1}(s))\right) \\
\text{s.t.} \quad & z + (S(1 - \alpha))^{-1}\sum_{s=1}^{S} y_s \le \gamma W_t
\end{aligned}
$$

$$\mathbf{r_{t+1}}(s)^{\top}\mathbf{u_t} + z + y_s \geq 0, \quad \forall s = 1, \ldots, S$$
$$y_s \geq 0, \quad \forall s = 1, \ldots, S$$
$$\mathbf{1}^{\top}\mathbf{u_t} = W_t$$
$$\mathbf{u_t} \geq 0, \tag{26}$$

where $\log(1 + \mathbf{r_{t+1}}(s)) = \mathbf{a_r} + \mathbf{b_r} f_t + \varepsilon_{t+1}(s)$ and $f_{t+1}(s) = a_f + b_f f_t + \eta_{t+1}(s)$.

Note that the value function $Q_{t+1}(1, f_{t+1})$ only varies with the factor realization, which is a scalar random variable. For an efficient approximation we go backward in time evaluating Q_t for many values of f_t and we use linear interpolation to approximate Q_t for the remaining values of f_t.

To evaluate solution quality of the SAA policy, we compute a probabilistic optimality for the "true" (continuously distributed) problem as in Kleywegt et al. (2002). For $S = 750$, $\alpha = 90\%$ and $\gamma \in \{0.02, 0.05, 0.08, 0.10, 0.20\}$ we obtain a optimality gap smaller than 1% with a probabilistic guarantee of 90%.

Considering proportional transaction costs, traditional dynamic programming techniques become computationally intractable due to the high-dimensional state space. To efficiently solve the problem with MSDDP, we need to approximate factor dynamics by a Markov chain and formulate it as in (24). In order to estimate the Markov chain we: (i) construct 1000 scenarios paths with 240 periods using Brown and Smith (2011) stochastic return model; (ii) use the simulated scenarios as inputs to the Baum–Welch algorithm (Baum et al. 1970) to construct the Markov Model. For this purpose, we also need to determine the number of Markov states required for an accurate approximation. For the no-transactional-cost case, we know the optimal policy obtained by (26), and we use its optimal expected return as a benchmark to compare with MSDDP policies with an increasing number of Markov states and different values of γ.

The objective function of (26) represents the expected terminal wealth, which is defined as the initial wealth plus the optimal expected return. To assess solution quality, we consider only the expected return part since the initial wealth is a constant. Said that, we assess the return difference (optimal policy minus Markov proxy) using 1000 out-of-sample paths to compute a sample average estimator $\overline{D}$ and perform a pairwise t-test with the null hypothesis $H_0 : \overline{D} \geq \varepsilon_\gamma$, where the error tolerance ε_γ is defined as 2.5% of the optimal expected rate of return for a given γ. Rejecting the null hypothesis represents a strong evidence that the optimality gap is smaller than a given tolerance ε_γ. According to the p values presented in Table 1, we observe that policies with intermediate values of γ have more complex return dynamics and, consequently, require a higher number of Markov states to obtain an accurate approximation. Indeed, extreme values of γ induce single-asset portfolios with much simpler return dynamics. Since Markov proxy policies with $K \geq 3$ rejects the null hypothesis for all risk limits (γ), we present further analysis for the policy with $K = 3$.

For comparison purposes, we devise heuristic investment strategies, using the same parameters as in Brown and Smith (2011) for the stochastic return model, to compare with the MSDDP solution; the parameters are specified in the Appendix.

In our case, all benchmark strategies must impose the same risk constraint of the original model. We devise two benchmark strategies namely, future-blind strategy (FBS) and future-cost-blind strategy (FCBS). On the one hand, the FBS solves

$$H_t(\mathbf{x_t}, f_t) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \quad \mathbb{E}\left[(\mathbf{1} + \mathbf{r_{t+1}})^{\top}\mathbf{u_t} \,\middle|\, \tilde{f}_t = f_t\right]$$
$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}^{\top}}\mathbf{u_t} \,\middle|\, \tilde{f}_t = f_t\right] + \mathbf{c}^{\top}(\mathbf{b_t} + \mathbf{d_t}) \leq \gamma\left(\mathbf{1}^{\top}\mathbf{x_t}\right)$$

Table 1 p Values for the difference of expected return associated with SDP versus K-state Markov SDDP

γ	Markov SDDP policies			
	$K = 1$	$K = 2$	$K = 3$	$K = 4$
0.02	0.253	0.554	0.000	0.000
0.05	0.418	0.632	0.000	0.007
0.08	0.454	0.669	0.000	0.013
0.10	0.059	0.057	0.001	0.006
0.20	0.000	0.000	0.000	0.000

$$u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t}$$
$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A}$$
$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0 \tag{27}$$

optimizing the next period return considering time dependence and transaction costs and ignoring the future value function. On the other hand, the FCBS solves

$$H_t(\mathbf{x_t}, f_t) = \max_{\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t}} \quad \mathbb{E}\left[(\mathbf{1} + \mathbf{r_{t+1}})^\top \mathbf{u_t} \cdot Q_{t+1}\left(1, \tilde{f}_{t+1}\right) \Big| \tilde{f}_t = f_t\right]$$

$$\text{s.t.} \quad \rho\left[\mathbf{r_{t+1}}^\top \mathbf{u_t} \Big| \tilde{f}_t = f_t\right] + \mathbf{c}^\top(\mathbf{b_t} + \mathbf{d_t}) \leq \gamma\left(\mathbf{1}^\top \mathbf{x_t}\right)$$

$$u_{0,t} + (\mathbf{1} + \mathbf{c})^\top \mathbf{b_t} - (\mathbf{1} - \mathbf{c})^\top \mathbf{d_t} = x_{0,t}$$

$$u_{i,t} - b_{i,t} + d_{i,t} = x_{i,t}, \quad \forall i \in \mathcal{A}$$

$$\mathbf{u_t}, \mathbf{b_t}, \mathbf{d_t} \geq 0, \tag{28}$$

considering a the simplified future value function Q_{t+1} defined in (26) that disregards future transaction costs.

To compare the MSDDP solution[5] with these alternative heuristic approaches (FBS and FCBS), we simulate all strategies using 1000 return scenarios generated using (25). We present the results for $\gamma = \{0.05, 0.1, 0.2, 0.3\}$ and transaction cost $c = \{0.005, 0.013, 0.02\}$ in Table 2.

Note that for a small transaction cost rate, the expected return is almost the same for MSDDP and the benchmark strategies. As transactional costs increase, the expected return significantly outperforms selected benchmarks for all loss-tolerance levels γ. For better visualization of the results, we plot risk-return curves (γ versus expected return) for all strategies and differing transactional costs in Fig. 7.

The influence of the transactional costs is explicit in Fig. 7. When it is low, the results are very similar, but for the highest values, the difference between the approaches significantly increases. Thus, for $c = 0.02$, FBS an FCBS allocate everything in the risk-free asset because of the costs negative effect. Note that the heuristic strategies barely invest in risky assets when transactional costs are significantly high, while in the MSDDP, we observe only a small shift in the efficient frontier.

5.2 Realistic 100-asset and 5-factor case study

Now, let us consider a realistic case study with 100 assets with transactional costs and a factor model for asset returns. We use Kenneth French monthly data for 100 portfolios by

---

[5] Here, we assume $K = 3$ Markov states and $S = 750$ for the sample average approximation.

Table 2 Expected return (in 12 months) for different methods, γ and transactional costs

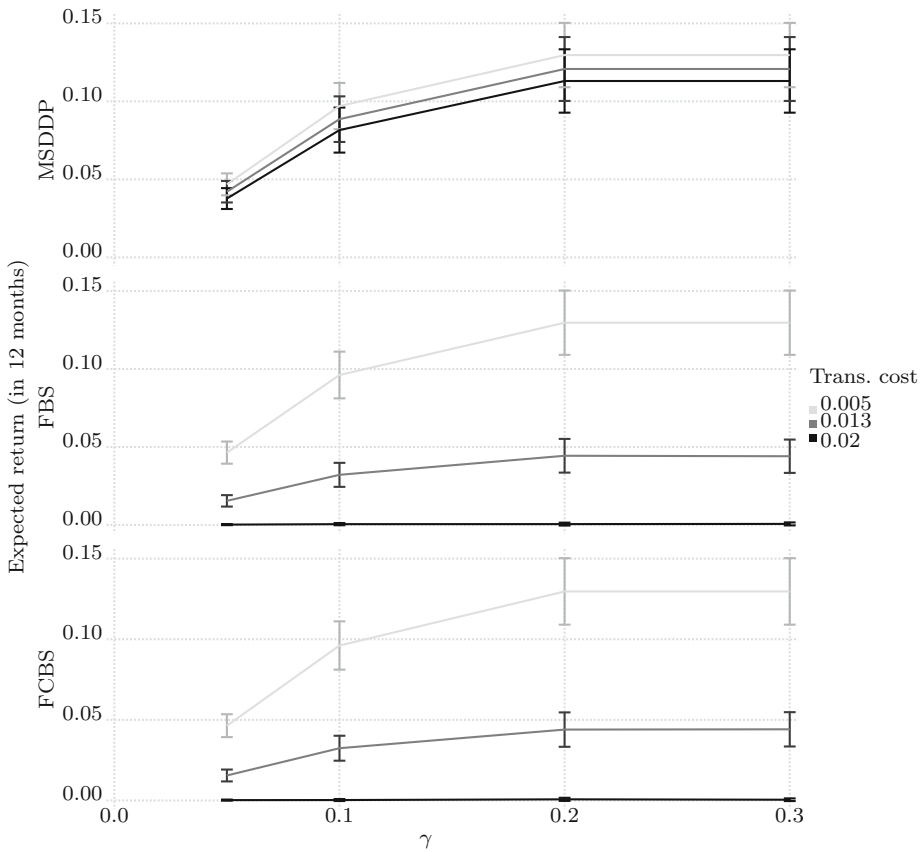| Trans. cost	γ	MSDDP	FBS	FBCS
0.005	0.05	0.0468	0.0464	0.0464
	0.1	0.0970	0.0961	0.0961
	0.2	0.1297	0.1297	0.1297
	0.3	0.1297	0.1297	0.1297
0.013	0.05	0.0421	0.0155	0.0155
	0.1	0.0886	0.0321	0.0324
	0.2	0.1207	0.0444	0.0440
	0.3	0.1207	0.0441	0.0441
0.02	0.05	0.0377	0.0003	0.0002
	0.1	0.0816	0.0006	0.0003
	0.2	0.1131	0.0006	0.0007
	0.3	0.1131	0.0007	0.0005

Fig. 7 Risk return curves for 3-asset and 1-factor case study

book-to-market and operating profitability, which are intersections of 10 portfolios formed on size (market equity) and 10 portfolios portfolios formed on profitability that include most of the NYSE, AMEX, and NASDAQ stocks, and 5 factors as in Fama and French (2015,
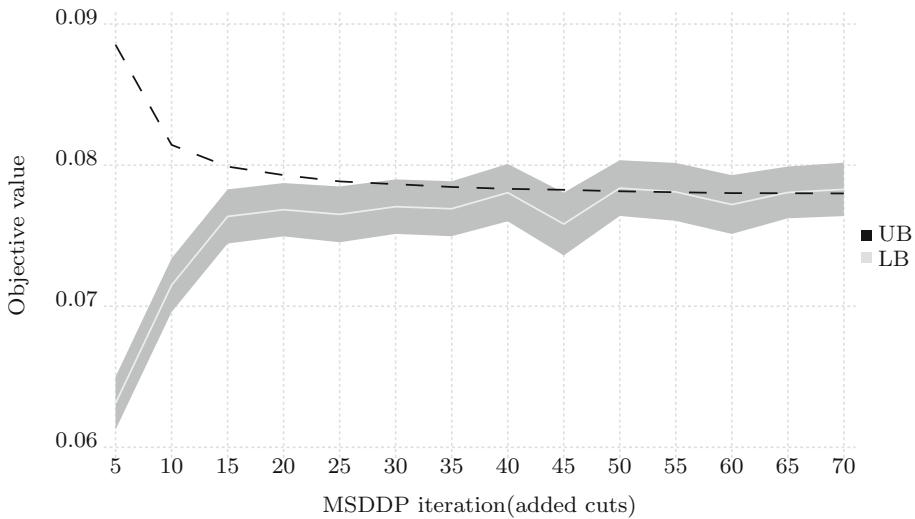
Fig. 8 Convergence of the MSDDP for 100-asset

2016). We estimate a Markov chain with 3 discrete state to represent factor dynamics and assume the estimated stochastic process is the true one; see Appendix for details. We consider a sample average approximation (SAA) of (23) and solve its discrete representation as in (24) for $T = 12$, $S = 750$, $\alpha = 90\%$ and $K = 3$ with a probabilistic guarantee of 90% that the optimality gap is smaller than 0.01 of the upper bound. We present the results of a particular SAA instance for $\gamma = 0.05$ and $c = 0.01$ in Fig. 8 with the upper bound and the confidence interval of the lower bound (y-axis) evolution over MSDDP iterations (x-axis). In the last MSDDP iteration, the upper bound for the expected rate of return is 7.8%, and the difference between the upper bound and lowest value of the lower bound confidence interval is smaller than 1%.

Moreover, we observe in Fig. 9 how the percentage allocation changes as MSDDP converges. Note that in the first iteration, the value function is not considered in the allocation, and consequently, it starts with 100% in the risk-free asset, which is the first asset. As the approximation of the value function improves, the allocation changes, converging to a portfolio with 12 of the 100 assets.

To assess solution quality of the SAA for the "true" (continuously distributed) problem we solve 10 randomly generated instances of the SAA problem and compute a probabilistic optimality gap as in Kleywegt et al. (2002). We solve each instance using a single core of Intel Xeon E5-2680 2.7 GHz with 128 GB of memory. The implementation was performed in Julia (Bezanzon et al. 2012) using JuMP (Lubin and Dunning 2015) and CPLEX(V12.5.1) to solve linear programming problems. The computational time for solving each SAA instance is around 5 hours. Considering a probabilistic guarantee of 90% we obtain a gap smaller than 1.02% relative to expected rate of return. In other Moreover, we perform out-of-sample simulation and compare with the previously described heuristic strategies based on Brown and Smith (2011): (i) future-blind strategy (FBS); (ii) future-cost-blind strategy (FCBS). In Fig. 10, we present the risk-return curves for the MSDDP, FBS and FCBS. The FBS and FCBS curves are virtually the same while the MSDDP provides significant improvements over the previously stablished heuristic approaches.
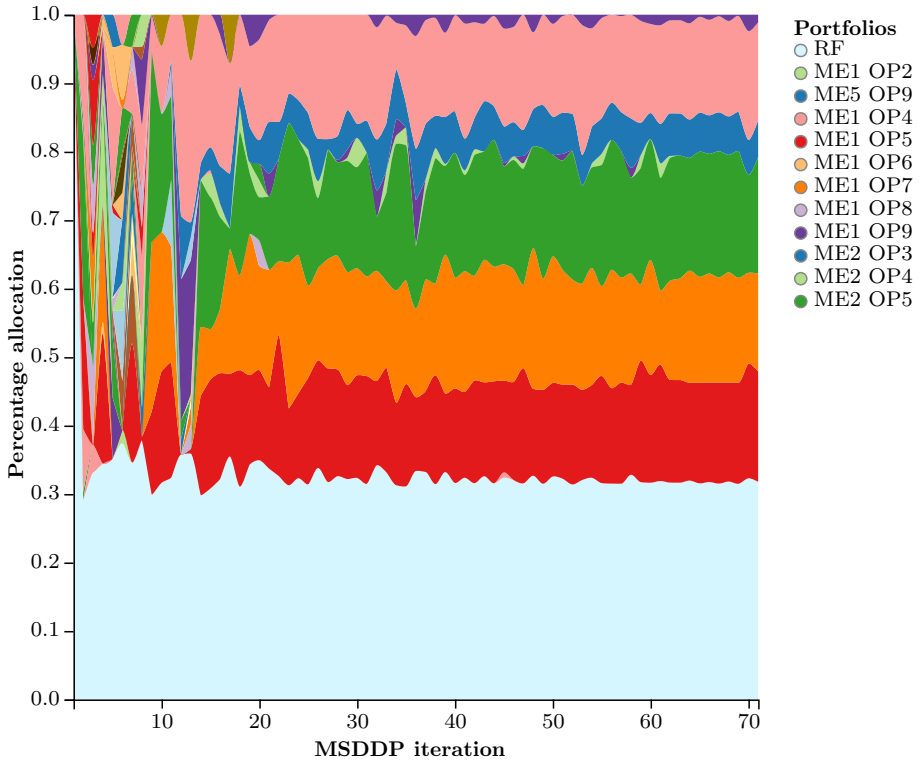
Fig. 9 Percentage allocation by MSDDP iteration for 100-asset
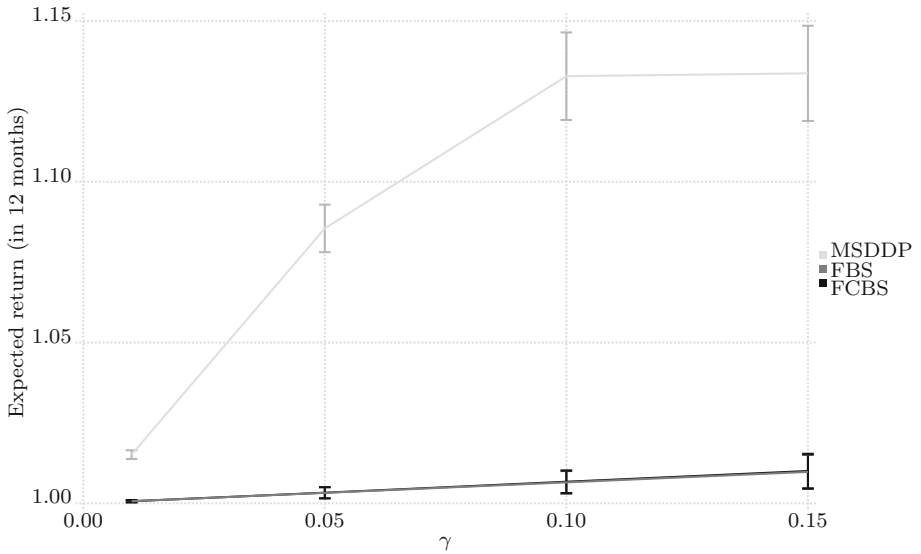

Fig. 10 Risk-return curves for the 100-asset and 5-factor case study

6 Conclusions and future works

In this work, we proposed a realistic dynamic asset allocation model that considers multiple assets, transactional costs and a Markov factor model for asset returns. We introduce risk aversion using an intuitive user-defined parameter to limit the maximum percentage loss in one period. The model maximizes the expected portfolio return with one-period conditional CV@R constraints, guaranteeing a time-consistent optimal policy, i.e., planned decisions are actually going to be implemented. We extend the literature results of myopic policies for our risk-constrained model if we assume no transactional costs and stage-wise (time) independent asset returns. Moreover, we proved a relatively complete recourse model when the maximum percentage loss limit is greater than or equal to the transactional cost rate. We efficiently solve the proposed model using Markov chained stochastic dual dynamic programming (MSDDP) for an illustrative problem with 3 assets and 1 factor with an autoregressive dynamic and for a high-dimensional problem with 100 assets and 5 factors following a discrete Markov chain.

For the 3-asset problem, we obtain a Markov chain surrogate for the true factor dynamic and solve the sample average approximation of the proposed model via MSDDP. Considering no transactional costs, we assess solution quality comparing the optimal policy (obtained via traditional dynamic programming techniques) against our approximate solution. We show the Markov approximate policy to be sufficiently accurate since one cannot reject the hypothesis that its expected return is equal to the optimal expected return. With transaction cost, we empirically show that our approximate solution outperforms selected heuristic benchmarks. We observe that these heuristics are close the optimal when transactional costs are small but perform poorly otherwise. Indeed, we observe that, at a certain level of transactional costs, one-period portfolio optimization does not invest in risky assets, while our approximate policy suggests significant allocation in the risky assets.

To show the computational tractability and scalability of our approach, we solve a sample average approximation of a high-dimensional 100-asset problem with 5 factors following a Markov chain with 3 discrete states. We provide convergence graphs that illustrate a probabilistic guarantee of optimality, i.e., we ensure with 90% probability that the optimality gap is smaller than or equal to 1.02% in terms of expected rate of returns. We also illustrate how first stage percentage asset allocation changes as the MSDDP converges. Indeed, the allocation in the first iterations of the algorithm is close to the static (one-period) counterpart and converges to the optimal solution of the dynamic problem. Finally, we assess solution quality of our approach by solving 10 randomly generated SAA instances to obtain a probabilistic optimality gap for the "true" (continuously distributed) problem. We empirically show a optimality gap for the "true" problem is smaller than 1.02% with a probabilistic guarantee of 90%. We also present superior risk-return curves computed with randomly generated out-of-sample paths for asset returns.

To the best of our knowledge, this is the first systematic approach for time-consistent risk-constrained dynamic portfolio optimization with transactional costs and time-dependent returns. We argue that our framework is flexible enough to approximately represent a relevant set of dynamic portfolio models and operational constraints on holdings and turnovers. In addition, we believe that time-consistent risk-constrained policies are better suited for practical applications since its risk aversion parameterization is direct and intuitive for all investors capable of determining a maximum percentage loss allowed in a given period.

It is important to note that the proposed model uses the Markov model estimates as the true underlying stochastic process. In practice however, investors face estimation errors that compromise significantly out-of-sample portfolio performance. As future research topic, we

Table 3 3-asset regression model

	Large-cap	Mid-cap	Small-cap	Dividend Yield
Mean (a_r, a_z)	0.0053	0.0067	0.0072	0.0000
Regression coeff. (b_r, b_z)	0.0028	0.0049	0.0061	0.9700
Covariance (Σ_{ev})				
Large-cap	0.002894	0.003532	0.00391	−0.000115
Mid-cap		0.004886	0.005712	−0.000144
Small-cap			0.007259	−0.000163
Dividend yield				0.052900

envision a distributionally robust extension of the proposed model that can efficiently handle estimation error seeking superior out-of-sample results.

Details of case studies

Illustrative 3-asset case study

The main goal of this experiment was to compare the SDP alternative approach with MSDDP. Regression coefficients a_r and b_r was the same as Brown and Smith (2011) as shown in Table 3.

and $r_f = 0.00042$.

Realistic 100-asset and 5-factor case study

As it is a factor model the Markov states are only evaluated for the factors. To construct the returns regression it was used vectors a_r and b_r that are obtained using linear regression. The estimates of factor probability distributions conditioned to each Markov state is given by

$$
f_1 \sim N \left[\begin{pmatrix} 0.0232 \\ 0.0073 \\ -0.0135 \\ -0.0145 \\ -0.0024 \end{pmatrix}, \begin{pmatrix} -0.0004 & 0.0039 & & & \\ -0.0015 & -0.0008 & 0.0033 & & \\ -0.0001 & -0.0017 & 0.0016 & 0.0029 & \\ -0.0005 & 0.0004 & 0.0014 & -0.0002 & 0.0014 \end{pmatrix} \right],
$$

$$
f_2 \sim N \left[\begin{pmatrix} 0.0070 \\ 0.0021 \\ 0.0024 \\ 0.0024 \\ 0.0017 \end{pmatrix}, \begin{pmatrix} 0.0013 & & & & \\ 0.0003 & 0.0006 & & & \\ -0.0002 & 0.0000 & 0.0005 & & \\ 0.0000 & -0.0001 & -0.0001 & 0.0002 & \\ -0.0001 & 0.0000 & 0.0003 & -0.0001 & 0.0003 \end{pmatrix} \right],
$$

$$
f_3 \sim N \left[\begin{pmatrix} -0.0375 \\ -0.0013 \\ 0.0195 \\ 0.0125 \\ 0.0174 \end{pmatrix}, \begin{pmatrix} 0.0062 & & & & \\ 0.0024 & 0.0023 & & & \\ 0.0000 & 0.0005 & 0.0022 & & \\ -0.0007 & -0.0001 & 0.0006 & 0.0013 & \\ -0.0009 & -0.0001 & 0.0012 & 0.0007 & 0.0012 \end{pmatrix} \right],
$$

where $\log(1 + \mathbf{r_{t+1}}) = \mathbf{a_r} + \mathbf{b_r} \mathbf{F_t} + \varepsilon_\mathbf{t}$ and $\mathbf{F_t} = \mathbf{f_k}$ if $K_t = k$.

References

Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, *9*(3), 203–228.

Baum, L. E., Petrie, T., Soules, G., & Weiss, N. (1970). A maximization technique occurring in the statistical analysis of probabilistic functions of markov chains. *The Annals of Mathematical Statistics*, *41*(1), 164–171.

Bezanzon, J., Karpinski, S., Shah, V., & Edelman, A. (2012). Julia: A fast dynamic language for technical computing. In: *Lang.NEXT*. http://julialang.org/images/lang.next.pdf.

Bick, B., Kraft, H., & Munk, C. (2013). Solving constrained consumption–investment problems by simulation of artificial market strategies. *Management Science*, *59*(2), 485–503.

Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Boston, MA: Springer.

Blomvall, J., & Shapiro, A. (2006). Solving multistage asset investment problems by the sample average approximation method. *Mathematical Programming*, *108*(2–3), 571–595.

Bradley, S. P., & Crane, D. B. (1972). A dynamic model for bond portfolio management. *Management Science*, *19*(2), 139–151.

Brigatto, A., Street, A., & Valladão, D. M. (2017). Assessing the cost of time-inconsistent operation policies in hydrothermal power systems. *IEEE Transactions on Power Systems*, *32*(6), 4541–4550.

Brown, D. B., & Smith, J. E. (2011). Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science*, *57*(10), 1752–1770.

Brown, D. B., Smith, J. E., & Sun, P. (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research*, *58*(4–part–1), 785–801.

Bruno, S., Ahmed, S., Shapiro, A., & Street, A. (2016). Risk neutral and risk averse approaches to multistage renewable investment planning under uncertainty. *European Journal of Operational Research*, *250*(3), 979–989.

Chen, Z. L., & Powell, W. B. (1999). Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse. *Journal of Optimization Theory and Applications*, *102*(3), 497–524.

Cheridito, P., Delbaen, F., Kupper, M., et al. (2006). Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability*, *11*(3), 57–106.

Constantinides, G. M. (1979). Multiperiod consumption and investment behavior with convex transactions costs. *Management Science*, *25*(11), 1127–1137.

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, *1*(2), 223–236.

Davis, M. H., & Norman, A. R. (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research*, *15*(4), 676–713.

Donohue, C. J., & Birge, J. R. (2006). The abridged nested decomposition method for multistage stochastic linear programs with relatively complete recourse. *Algorithmic Operations Research*, *1*(1), 20–30.

Elliott, R. J., & Siu, T. K. (2014). Strategic asset allocation under a fractional hidden markov model. *Methodology and Computing in Applied Probability*, *16*(3), 609–626.

Elliott, R. J., & Van der Hoek, J. (1997). An application of hidden markov models to asset allocation problems. *Finance and Stochastics*, *1*(3), 229–238.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, *116*(1), 1–22.

Fama, E. F., & French, K. R. (2016). Dissecting anomalies with a five-factor model. *Review of Financial Studies*, *29*(1), 69–103.

Fernandes, B., Street, A., Valladão, D., & Fernandes, C. (2016). An adaptive robust portfolio optimization model with loss constraints based on data-driven polyhedral uncertainty sets. *European Journal of Operational Research*, *255*(3), 961–970.

Granger, C. W., & Ding, Z. (1994). *Stylized facts on the temporal and distributional properties of daily data from speculative markets*. UCSD Department of Economics Discussion Paper (pp. 94–19).

Gülpınar, N., Rustem, B., & Settergren, R. (2004). Simulation and optimization approaches to scenario tree generation. *Journal of Economic Dynamics and Control*, *28*(7), 1291–1315.

Hassan, M. R., & Nath, B. (2005). Stock market forecasting using hidden markov model: a new approach. In *Intelligent systems design and applications, 2005 ISDA'05. Proceedings of 5th international conference on IEEE* (pp. 192–196).

Infanger, G., & Morton, D. P. (1996). Cut sharing for multistage stochastic linear programs with interstage dependency. *Mathematical Programming*, *75*(2), 241–256.

Jin, X., & Zhang, K. (2013). Dynamic optimal portfolio choice in a jump-diffusion model with investment constraints. *Journal of Banking & Finance*, *37*(5), 1733–1746.

Kleywegt, A. J., Shapiro, A., & de Mello, T. H. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, *12*(2), 479–502.

Kouwenberg, R. (2001). Scenario generation and stochastic programming models for asset liability management. *European Journal of Operational Research*, *134*(2), 279–292.

Kozmík, V., & Morton, D. P. (2015). Evaluating policies in risk-averse multi-stage stochastic programming. *Mathematical Programming*, *152*(1–2), 275–300.

Lubin, M., & Dunning, I. (2015). Computing in operations research using julia. *INFORMS Journal on Computing*, *27*(2), 238–248.

Mamon, R. S., & Elliott, R. J. (Eds.). (2007). *Hidden Markov models in finance, international series in operations research & management science* (Vol. 104). Boston, MA: Springer.

Mamon, R. S., & Elliott, R. J. (Eds.). (2014). *Hidden Markov models in finance, international series in operations research & management science* (Vol. 209). Boston, MA: Springer.

Mei, X., DeMiguel, V., & Nogales, F. J. (2016). Multiperiod portfolio optimization with multiple risky assets and general transaction costs. *Journal of Banking & Finance*, *69*, 108–120.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics*, *51*(3), 247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, *3*(4), 373–413.

Mo, B., Gjelsvik, A., & Grundt, A. (2001). Integrated risk management of hydro power scheduling and contract management. *IEEE Transactions on Power Systems*, *16*(2), 216–221.

Morettin, P. A., & Toloi, C. (2006). *Análise de séries temporais*. Sao Paulo: Blucher.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business*, *41*(2), 215–229.

Muthuraman, K. (2007). A computational scheme for optimal investment-consumption with proportional transaction costs. *Journal of Economic Dynamics and Control*, *31*(4), 1132–1159.

Pereira, M., & Pinto, L. (1985). Stochastic optimization of a multireservoir hydroelectric system: A decomposition approach. *Water Resources Research*, *21*(6), 779–792.

Pereira, M. V. F., & Pinto, L. M. V. G. (1991). Multi-stage stochastic optimization applied to energy planning. *Math Program*, *52*(2), 359–375.

Philpott, A. B., & de Matos, V. L. (2012). Dynamic sampling algorithms for multi-stage stochastic programs with risk aversion. *European Journal of Operational Research*, *218*(2), 470–483.

Powell, W. B. (2011). *Approximate dynamic programming: Solving the curses of dimensionality* (2nd ed., Vol. 703). Hoboken: Wiley.

Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, *2*, 21–42.

Rockafellar, R., & Uryasev, S. (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, *26*(7), 1443–1471.

Rockafellar, R. T., & Wets, R. J. B. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, *16*(1), 119–147.

Rudloff, B., Street, A., & Valladão, D. M. (2014). Time consistency and risk averse dynamic decision models: Definition, interpretation and practical consequences. *European Journal of Operational Research*, *234*(3), 743–750.

Rydén, T., Teräsvirta, T., & Åsbrink, S. (1998). Stylized facts of daily return series and the hidden markov model. *Journal of Applied Econometrics*, *13*(3), 217–244.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics*, *51*(3), 239–246.

Shapiro, A. (2006). On complexity of multistage stochastic programs. *Operations Research Letters*, *34*(1), 1–8.

Shapiro, A. (2009). On a time consistency concept in risk averse multistage stochastic programming. *Operations Research Letters*, *37*(3), 143–147.

Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, *209*(1), 63–72.

Shapiro, A., & Nemirovski, A. (2005). *On complexity of stochastic programming problems* (pp. 111–146). Boston, MA: Springer.

Shapiro, A., Tekaya, W., da Costa, J. P., & Soares, M. P. (2013). Risk neutral and risk averse stochastic dual dynamic programming method. *European Journal of Operational Research*, *224*(2), 375–391.

Shreve, S. E., & Soner, H. M. (1994). Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, *4*(3), 609–692.

Soares, M. P., Street, A., & Valladão, D. (2017). On the solution variability reduction of stochastic dual dynamic programming applied to energy planning. *European Journal of Operational Research*, *258*(2), 743–760.

Street, A., Brigatto, A., & Valladão, D. M. (2017). Co-optimization of energy and ancillary services for hydrothermal operation planning under a general security criterion. *IEEE Transactions on Power Systems*, *32*(6), 4914–4923.

Topaloglou, N., Vladimirou, H., & Zenios, S. A. (2008). A dynamic stochastic programming model for international portfolio management. *European Journal of Operational Research*, *185*(3), 1501–1524.

Valladão, D. M., Veiga, Á., & Veiga, G. (2014). A multistage linear stochastic programming model for optimal corporate debt management. *European Journal of Operational Research*, *237*(1), 303–311.

Valladão, D. M., Veiga, Á., & Street, A. (2018). A linear stochastic programming model for optimal leveraged portfolio selection. *Computational Economics*, *51*(4), 1021–1032.

Ziemba, W., Mulvey, J., & For Mathematical Sciences INI. (1998). *Worldwide asset and liability modeling. Includes bibliographies*. Cambridge: Cambridge University Press.