



Multi-period portfolio optimization with linear control policies[☆]

Giuseppe Carlo Calafiore

Dipartimento di Automatica e Informatica, Politecnico di Torino, Italy

ARTICLE INFO

Article history:

Received 28 March 2007
Received in revised form
2 October 2007
Accepted 3 February 2008
Available online 13 May 2008

Keywords:

Risk management
Finance
Dynamic portfolio allocation
Multi-stage decisions
Convex optimization

ABSTRACT

This paper is concerned with multi-period sequential decision problems for financial asset allocation. A model is proposed in which periodic optimal portfolio adjustments are determined with the objective of minimizing a cumulative risk measure over the investment horizon, while satisfying portfolio diversity constraints at each period and achieving or exceeding a desired terminal expected wealth target. The proposed solution approach is based on a specific affine parameterization of the recourse policy, which allows us to obtain a sub-optimal but exact and explicit problem formulation in terms of a convex quadratic program.

In contrast to the mainstream stochastic programming approach to multi-period optimization, which has the drawback of being computationally intractable, the proposed setup leads to optimization problems that can be solved efficiently with currently available convex quadratic programming solvers, enabling the user to effectively attack multi-stage decision problems with many securities and periods.

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

The fundamental goal of portfolio theory is to help the investor allocating money among different financial securities in an “optimal” way. In the classical Markowitz framework, (Markowitz, 1991), the selection is guided by a quantitative criterion that considers a tradeoff between the return of an investment and its associated risk. Specifically, in the Markowitz approach, each asset is described by means of its return over a fixed period of time (e.g. one month), and the vector of asset returns is assumed to be random, with known expectation and covariance matrix. An optimal portfolio of assets is hence selected, for instance, by minimizing the investment risk (as expressed by the portfolio variance, or “volatility”) subject to a given lower bound on the expected return at the end of the period. From the computational side, this classical paradigm results in a quadratic programming problem, which may be efficiently solved numerically on a computer.

However, a drawback of this basic approach is that it is tuned to a single period, and it can therefore provide short-sighted strategies of investment, if applied repeatedly over many subsequent periods. To overcome this issue, one may formulate from the beginning the allocation problem over an horizon composed of multiple periods ($T \geq 1$ periods), with the goal of minimizing the total risk over the investment path, while satisfying

constraints on the portfolio composition and on desired expected return at all the intermediate stages.

Seminal contributions to multi-stage decisions in finance have been given in Merton (1971, 1973), where an approach based on continuous-time dynamic programming is proposed. This paradigm is still in use – see for instance (Ait-Sahalia & Brandt, 2001; Brennan, Schwartz, & Lagnado, 1997; Lynch & Balduzzi, 2000). Note, however, that the dynamic programming approach is impractical for actual numerical implementation, due to the “curse of dimensionality.” Indeed, Brandt (1999) and Brennan et al. (1997) report that incorporating more than a few state variables in the dynamic programming formulation leads to unworkable problem size (and probably for this reason most of the multi-period models encountered in the literature are actually two-periods models with only a few securities). Analytical or reduced complexity solutions can be obtained only in very special cases. For instance, a closed form solution in continuous-time is proposed in Zhou and Li (2000) when no transaction costs are considered and no constraints are imposed on the portfolio composition. On a similar note, a mean-variance discrete-time problem is reduced to a control problem with only one state variable in Infanger (2006), under the hypotheses of no transaction costs, no composition constraints and serially independent returns. It should hence be remarked that the presence of constraints on portfolio composition and/or of transaction costs makes the problem harder from the computational viewpoint.

The mainstream computational model to solve recursive decision problems in the presence of uncertainty is currently provided by multi-stage stochastic programming, see e.g. Birge (1997), Birge and Louveaux (1997), Gulpinar, Rustem, and Settergren (2002), and Ruszczyński and Shapiro (2003) and the many references

[☆] This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Editor Berc Rüstern.

E-mail address: giuseppe.calafiore@polito.it.

therein. However, while stochastic programming provides a conceptually sound framework for posing multi-stage decision problems, from the computational side it results to be impervious to exact and efficient numerical solution, (Shapiro, 2005). The key difficulty in the stochastic programming formulation comes from the fact that the stage decisions are actually conditional decision rules, or “policies” that define which action should be taken in response to past outcomes. To model the conditional nature of the problem in some “tractable” way, a discretization of the decision space is typically introduced by constructing a “scenario tree,” and this scenario tree may grow exponentially if an accurate and representative discretization is needed, see e.g. Ermoliev and Wets (1988). On the other hand, if branching is kept low in the scenario tree, the resulting discretization cannot be guaranteed to be a reliable representation of reality. These computational difficulties are witnessed by the fact that most multi-period problems discussed in the literature deal with few securities over only two or three periods.

In the specific context of financial allocation, a classical stochastic programming method based on Benders decomposition is proposed in Dantzig and Infanger (1993), and techniques for construction of scenario trees are discussed for instance in Gulpinar, Rustem, and Settergren (2004), Pflug (2001) and Ziemba and Mulvey (1998). A possibly improved approach that uses scenario trees together with simulated paths has been recently studied in Hibiki (2006). Also, in Rustem and Gulpinar (2007), the authors extend the multi-period mean-variance model to deal with rival uncertainty scenarios, and propose a worst-case decision approach which uses a min-max technique in synergy with a stochastic optimization algorithm based on scenario trees. Scenario-based stochastic programming models aiming at maximizing expected portfolio value while taking into account cost variability have been recently proposed also in Pinar (2007) and Takriti and Ahmed (2004). A survey with theoretical analysis of multiperiod models based on scenario trees is provided in Steinbach (2001).

In this paper, we propose a different route to multi-stage portfolio allocation, which prescind from the use of decision trees or sample paths, and which leads to explicit convex constrained quadratic programming models that can be solved globally and efficiently.

We achieve this goal by considering recourse actions that are prescribed by policies with fixed structure. In particular, we shall consider recourse actions that are affine functions of the past periods returns, where the coefficients of these functions become the decision variables of the problem. While with this position we lose some generality, since the control policy is now restricted to the affine functions class, we also gain decisive advantages. First, we show that it is possible to express explicitly the expected value and variance of the portfolio at any stage, as a function of the decision variables. Furthermore, the optimization objective and constraints result to be convex quadratic functions of these variables, and therefore the optimal strategy, under the affine recourse hypothesis, can be found *exactly* and *numerically efficiently* by means of standard codes for convex quadratic programming. Second, the optimal recourse parameters returned by the algorithm have a simple and insightful interpretation as nominal actions and market reaction sensitivities, as further discussed in Section 5.

The idea of using affine recourse has been inspired by a similar “adjustable variables” technique recently proposed in Ben-Tal, Goryashko, Guslitzer, and Nemirovski (2004) in the context of robust optimization (i.e. optimization problems where the data is subject to deterministic unknown-but-bounded uncertainties), where the authors impose an affine dependence of variables on the data, so that variables can “tune themselves to varying data”. Adjustable variables have been proposed in the context of model




Fig. 1. Time axis and periods for multi-period allocation.

predictive control in Löfberg (2003), and, more generally, affine recourse is reminiscent of the classical linear feedback laws used for control of dynamical systems. A similar idea is also employed in dynamic programming for approximating the optimal cost-to-go function by fitting a parameterized function approximator, see for instance Bertsekas and Tsitsiklis (1996); Farias and Van Roy (2006).

This paper is organized as follows. Preliminary concepts are given in Section 2. Section 2.1 describes the recursive equations for the portfolio dynamics, and Section 3 presents the basic open-loop mean-variance optimization setup. The key Section 4 introduces the recourse model into the dynamic decision problem, and states a convexity result for affinely parameterized policies, in Lemma 1. An explicit multi-stage optimization model is developed in Section 5, under an additional assumption of independence of the market gains and a specific structure of the recourse rule. Numerical experiments are presented in Section 6, and conclusions are drawn in Section 7. Technical proofs are contained in the Appendix.

Notation. If $x \in \mathbb{R}^{n,1}$ is an n -dimensional vector, then $\text{diag}(x)$ denotes a diagonal matrix having the entries of x on the diagonal. A^\top denotes the transpose of matrix A . The operator \odot denotes the Hadamard (entry-wise) product of conformably sized matrices. For a random vector x taking values in $\mathbb{R}^{n,1}$, we denote with $E\{x\}$ the expected value of x , and with $\text{var}\{x\} \doteq E\{(x - E\{x\})(x - E\{x\})^\top\}$ the covariance matrix of x . We shall denote with an over bar the expectation of random quantities and with a tilde the centered quantities, i.e. the quantities with the expectation subtracted, that is $\bar{x} \doteq E\{x\}$, $\tilde{x} \doteq x - \bar{x}$.

2. Preliminaries

An investment instrument that can be bought and sold is here named an *asset*. In our discussion, an asset can be a single stock, cash, or an index representing the value of a class of securities. It is assumed that no dividends/coupons are paid, or if paid, they are immediately reinvested in the same asset.

Consider a universe of n assets $\{a_1, \dots, a_n\}$ and an investment timeline divided into T periods of equal duration Δ , as depicted in Fig. 1. Δ is a fixed time span, for instance $\Delta = 1$ month, or $\Delta = 1$ year.

We denote with $x_i(k)$ the Euro value of the portion of the investor's total wealth invested in security a_i at time k . The vector having components $x_i(k)$, $i = 1, \dots, n$ is here named the *portfolio*:

$$x(k) \doteq \begin{bmatrix} x_1(k) \\ \vdots \\ x_n(k) \end{bmatrix}.$$

The investor's total wealth at time k is given by

$$w(k) = \sum_{i=1}^n x_i(k) = \mathbf{1}^\top x(k),$$

where $\mathbf{1}$ denotes a vector of ones of suitable dimension. At the end of each period, the investor has the opportunity of adjusting his/her investment, by rebalancing the portfolio composition. Specifically, we denote with $x^+(k)$ the portfolio composition just after the adjustment $u(k)$ occurred at time k :

$$x^+(k) \doteq x(k) + u(k),$$




Fig. 2. Portfolio dynamics.

where $u(k) \doteq [u_1(k) \cdots u_n(k)]^\top$ is the vector of adjustments. A value of $u_i(k) > 0$ indicates that the portfolio content in asset a_i is increased by $u_i(k)$ Euros (by buying this asset), whereas $u_i(k) < 0$ indicates that the portfolio content in asset a_i is decreased by $u_i(k)$ Euros (by selling this asset). Overall, it must hold that

$$\sum_{i=1}^n u_i(k) = c(k),$$

where $c(k)$ is the amount of cash injected (if $c(k) > 0$) or withdrawn (if $c(k) < 0$) from the portfolio at time k . A *self-financing portfolio* is obtained by setting $c(k) = 0$ for all k , which we shall assume henceforth, without loss of generality.

Let $p_i(k)$ denote the market value of a_i at time k . The (simple) return of an investment in security a_i over the period of time $[k-1, k]$ is defined as

$$r_i(k) \doteq \frac{p_i(k) - p_i(k-1)}{p_i(k-1)} = \frac{p_i(k)}{p_i(k-1)} - 1.$$

The one-period *gain* (or total return) of the same investment is

$$g_i(k) \doteq \frac{p_i(k)}{p_i(k-1)} = r_i(k) + 1.$$

If $p(k)$ denotes the collection of prices of the n assets: $p(k) = [p_1(k) \ p_2(k) \ \cdots \ p_n(k)]^\top$, the return vector $r(k)$ and gain vector $g(k)$ are defined accordingly.

In this paper, we shall take a standard stochastic view of the market, and assume that the asset gains $g(k)$, $k = 1, \dots$ form a possibly non-stationary discrete-time stochastic process with finite mean and covariance.

2.1. Portfolio dynamics

Let $x(0)$ be the initial portfolio composition at time $k = 0$. At $k = 0$, we conduct transactions on the market and therefore adjust the portfolio by increasing or decreasing the amount invested in each asset. The adjusted portfolio is $x^+(0) = x(0) + u(0)$. Notice that constraints typically exist on the portfolio holdings and hence on the admissible adjustment vector $u(0)$. These constraints usually involve bounds on exposure in each single security or in clusters of securities and are discussed in more detail in Section 3.1.

Suppose now that the portfolio is held unchanged for the first period of time Δ . At the end of this first period, the portfolio composition is $x(1) = G(1)x^+(0) = G(1)x(0) + G(1)u(0)$, where $G(1) = \text{diag}(g(1))$ is a diagonal matrix of the asset gains over the period from time 0 to time 1. At time $k = 1$, we perform again an adjustment of the portfolio: $x^+(1) = x(1) + u(1)$, and then hold the updated portfolio for another period of duration Δ . At time $k = 2$ the portfolio composition is hence $x(2) = G(2)x^+(1) = G(2)x(1) + G(2)u(1)$. Proceeding in this way for $k = 0, 1, 2, \dots$, we determine the iterative dynamic equations of the portfolio composition at the end of period $(k+1)$:

$$x(k+1) = G(k+1)x(k) + G(k+1)u(k), \quad k = 0, \dots, T-1 \quad (1)$$

as well as the equations for portfolio composition just after the $(k+1)$ -th transaction (see Fig. 2)

$$x^+(k) = x(k) + u(k).$$

Notice that since the asset gains are random, recursion (1) defines a stochastic process $x(k)$, $k = 1, \dots, T$:

$$\begin{aligned} x(k) = & \Phi(1, k)x(0) \\ & + [\Phi(1, k) \quad \Phi(2, k) \quad \cdots \quad \Phi(k-1, k) \quad \Phi(k, k)] \\ & \times \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(k-2) \\ u(k-1) \end{bmatrix}, \end{aligned} \quad (2)$$

where we defined $\Phi(v, k)$, $v \leq k$, as the *compounded gain matrix* from the beginning of period v to the end of period k :

$$\Phi(v, k) \doteq G(k)G(k-1) \cdots G(v), \quad \Phi(k, k) \doteq G(k),$$

and where it holds that $\Phi(v, k+1) = G(k+1)\Phi(v, k)$. From (2) we obtain an expression for the total wealth at time k :

$$w(k) = \mathbf{1}^\top x(k) = \phi^\top(1, k)x(0) + \sum_{j=1}^k \phi^\top(j, k)u(j-1), \quad (3)$$

where

$$\phi^\top(j, k) \doteq \mathbf{1}^\top \Phi(j, k).$$

3. The basic (open-loop) optimization model

The framework in which we pose our basic asset allocation problem is a multi-period version of the classical mean-risk Markowitz model. More precisely, we quantify the total investment risk as a weighted sum of all stage wealth volatilities (variances):

$$J(T) = \sum_{k=1}^T \gamma(k) \text{var}\{w(k)\}, \quad (4)$$

where $\gamma(k) \geq 0$ are given weights. A possible choice is, for instance, to set $\gamma(k) = 0$ for $k = 1, \dots, T-1$ and $\gamma(T) = 1$; in this way only the final-stage risk would be accounted for in the cost index $J(T)$.

We consider optimal allocation strategies that minimize risk $J(T)$, while satisfying a lower-bound constraint for the total return at the final stage, as well as constraints on the expected portfolio composition at each intermediate stage.

The basic elements of the optimization model will first be described under a “naive,” open-loop approach. The open-loop approach takes the viewpoint of a decision maker that at time $k = 0$ wants to compute and freeze the whole sequence of optimal adjustments $u(0), \dots, u(T-1)$. This approach would clearly perform poorly in practice, since only $u(0)$ is the “here-and-now” action to be taken at $k = 0$. For the subsequent adjustment $u(1)$, the decision maker would better “wait and see” what the actual outcome of his/her first-stage decision is, and rethink the course of action accordingly, and so on for $u(2)$, etc. This smarter “recourse” approach and the ensuing optimization model is the core subject of the paper, and it is discussed in detail in the remaining sections. Here, we temporarily take the open-loop viewpoint, in order to illustrate more simply the basic components (minimization objective and constraints) of the optimization model.

3.1. Portfolio constraints

The following standard constraints are explicitly considered in our model. Notice however that generic linear constraints on the expected portfolio composition and on the adjustments at any stage may be accommodated in our setup.

Terminal return constraint. The expected return of the investment at the end of the decision horizon should be greater than a given lower bound Φ_{lb} :

$$E\{w(T)\} \geq \Phi_{lb}w(0).$$

Notice that the same type of constraint could be imposed also at every or at some of the intermediate decision stages.

Budget constraints. Each time the portfolio is adjusted, money value is transferred from one asset to another (cash may be part of the portfolio assets), but the net value remains unchanged, except for possible loss due to transaction costs and/or money withdrawn from or injected to the portfolio. This fact is expressed by the *budget constraints* $\mathbf{1}^\top u(k) = c(u(k))$, $k = 0, 1, \dots$, where $c(u(k))$ denotes all costs associated to the transactions, plus cash withdrawn/injected. Since the focus of this study is on a new technique for multi-stage problems, we shall assume for simplicity that $c(u(k)) = 0$. Under this assumption, the budget constraints become

$$\mathbf{1}^\top u(k) = 0, \quad k = 0, 1, \dots$$

Portfolio content constraints At each period k when the portfolio is rebalanced, we can impose constraints on the minimum and maximum expected exposure in an individual security, i.e.

$$\underline{b}_i(k) \leq E\{x_i^+(k)\} \leq \bar{b}_i(k),$$

$$i = 1, \dots, n; k = 0, 1, \dots, T-1,$$

where $\underline{b}_i(k), \bar{b}_i(k)$ are the given lower and upper bounds on portfolio holding in security a_i at time k , after rebalancing. For instance, if no shortselling is allowed, we impose this constraint using the previous bounds, with $\underline{b}_i(k) = 0, \bar{b}_i(k) = \infty$, i.e.

$$E\{x_i^+(k)\} \geq 0, \quad i = 1, \dots, n; k = 0, 1, \dots, T-1.$$

Portfolio relative diversity constraints We can limit the fraction of total (post-transaction) wealth held in each asset:

$$E\{x_i^+(k)\} \leq v_i(k) \mathbf{1}^\top E\{x^+(k)\},$$

$$i = 1, \dots, n; k = 0, 1, \dots, T-1,$$

where $v_i(k) \in [0, 1]$ is the limit fraction to be invested in security a_i at time k .

More generally, we can group the assets in compartments $\{c_1, \dots, c_\ell\}$ where each compartment $c_j \subseteq \{1, \dots, n\}$ contains the indices relative to a group of assets, and impose limits on the fraction of total (post-transaction) wealth held in each compartment:

$$v_{lb,j}(k) \mathbf{1}^\top E\{x^+(k)\} \leq \sum_{i \in c_j} E\{x_i^+(k)\} \leq v_{ub,j}(k) \mathbf{1}^\top E\{x^+(k)\},$$

$$j = 1, \dots, \ell; k = 0, \dots, T-1,$$

where $v_{lb,j}(k), v_{ub,j}(k) \in [0, 1]$ are respectively the limit lower bound and upper bound fractions to be invested in compartment c_j at time k .

Remark 1 (Constraints on Random Variables). Notice that, since the portfolio components $x_i^+(k)$ are random variables that causally depend on past returns, one cannot a-priori impose “deterministic” constraints on these values. There are indeed essentially three ways to treat these constraints in a multistage stochastic model. If scenario trees are used, then one may ask the constraints to hold for all the considered scenarios (point-wise constraints). Otherwise, the constraints can be imposed to hold with a given sufficiently high probability. The third possibility, which is the one implemented in this paper, is to impose the constraints in expectation, as shown in the previous paragraphs. Each option has its motivations and drawbacks. Point-wise constraints can be used only with scenario trees, and may lead to growth

of the constraints set with resulting increase of computational complexity; moreover, the constraints are still not guaranteed to hold for outcomes that were not included in the scenarios. Probability constraints are a sound way to specify risk of constraint violation, and may be imposed even when knowledge of the distribution is imprecise, by means of Chebychev-type bounds, see for instance (Calafiore & El Ghaoui, 2006). The drawback in this case is that linear constraints are transformed into second-order-cone constraints, which are more complicated to deal with numerically. The expected value approach has the advantage of preserving linearity of the constraints, while imposing them for the average case. While this may seem critical for “hard” constraints, such as the no-short-selling conditions, we argue that in practice the optimization model is typically implemented in a sliding horizon fashion (see Section 5.2), hence only the first decision and thus the first portfolio $x^+(0)$ are actually implemented. Since $x^+(0)$ is a deterministic quantity (it does not depend on past returns) the constraints in expectation are met exactly by the implemented portfolio. In this way, imposing constraints on expectation provides a good tradeoff between model computability and significance, and guarantees that the actual decisions are feasible in reality, whenever the optimal strategy is implemented in a receding horizon fashion. \diamond

3.2. Explicit open-loop formulation

Notice that since (3) is affine in the decision variables $u(k)$, the wealth variance at each stage is a convex quadratic function of these variables (for a simple proof of this statement one may follow the same reasoning as for the proof of Lemma 1 in Section 4.1). Hence, since the portfolio expectations are affine functions of the $u(k)$'s, the overall open-loop optimization problem can be cast as a convex quadratic programming problem. More explicitly, we have:

$$\min_{u(0), \dots, u(T-1)} \sum_{k=1}^T \gamma(k) \text{var}\{w(k)\}$$

subject to:

$$E\{w(T)\} \geq \Phi_{lb}w(0),$$

$$\mathbf{1}^\top u(k) = 0, \quad k = 0, \dots, T-1$$

$$\underline{b}(k) \leq E\{w(k)\} + u(k) \leq \bar{b}(k), \quad k = 0, \dots, T-1$$

$$v_{lb,j}(k) \mathbf{1}^\top E\{x(k)\} \leq \sum_{i \in c_j} (E\{x_i(k)\} + u_i(k)) \leq v_{ub,j}(k) \mathbf{1}^\top E\{x(k)\},$$

$$j = 1, \dots, \ell; k = 0, \dots, T-1,$$

where

$$\begin{aligned} E\{x(k)\} &= \bar{\Phi}(1, k)x(0) \\ &\quad + [\bar{\Phi}(1, k) \quad \bar{\Phi}(2, k) \quad \dots \quad \bar{\Phi}(k-1, k) \quad \bar{\Phi}(k, k)] \\ &\quad \times \begin{bmatrix} u(0) \\ u(1) \\ \vdots \\ u(k-2) \\ u(k-1) \end{bmatrix} \end{aligned}$$

and

$$E\{w(k)\} = \bar{\Phi}^\top(1, k)x(0) + \sum_{j=1}^k \bar{\Phi}^\top(j, k)u(j-1).$$

4. Multi-stage decision model with recourse

The problem described in Section 3.2 relates to an open-loop dynamic setting, in which the whole adjustment sequence $u(0), u(1), \dots, u(T-1)$ is fixed at the decision time $k = 0$. As already discussed, this setup hardly exploits the sequential nature of the decision problem at hand. Indeed, we remark that at time $k = 0$

the whole future sequence of asset gains is uncertain. However, only action $u(0)$ is actually applied to the portfolio at $k = 0$. At the subsequent time $k = 1$, the *actual outcome* of asset returns over the first period is revealed to the decision maker, and therefore his/her next action $u(1)$ should take into account this knowledge. In other words, at the decision time $k = 1$, only the future asset gains relative to periods from 2 onwards are still uncertain. In general, at the decision stage k , the past gains over periods $1, \dots, k$ have been observed, hence are exactly known, while the future gains of periods $k+1, \dots, T$ are still uncertain. Thus, uncertainty is reduced as the decision stage moves forward. Therefore, in a “controlled,” or closed-loop approach, we let the adjustment decisions $u(k)$ depend on the information state of the system at time k .

In this dynamic optimization setting, one seeks for an optimal action as a function of the information state of the system, i.e. $u(k)$ is given by a policy, or rule, π_k that associates an action to a given state of knowledge (recourse strategy). In the problem at hand, the information state is the observed sequence of portfolios $x(1), \dots, x(k)$. However, it can be proved that one can equivalently choose the observed sequence of returns $g(1), \dots, g(k)$ as the information state (see Section II.B of Calafiore and Campi (2005) for a formal proof of this statement). Hence we shall consider, without loss of generality, control policies of the form

$$u(k) = \pi_k(g(1), \dots, g(k)).$$

Notice that in this setting our optimization problem is a functional one, since we need to search over the infinite-dimensional space of policies π_k . An exact solution to this problem is numerically unfeasible. Techniques such as (constrained) dynamic programming, or stochastic optimization try to solve approximations of the problem, and invariably result in NP-hard problem formulations, see, e.g., Birge and Louveaux (1997) and Shapiro (2005).

We here propose a sub-optimal solution approach based on a finite-dimensional affine parameterization of the policies. This approach is discussed in the next section.

4.1. Affinely parameterized policies

Suppose that policy π_k is restricted to belong to a family Π_k of functions of $g(1), \dots, g(k)$ of given structure, that are *affinely parameterized* by a finite-dimensional collection of parameters. Then, it can be easily verified that also the portfolio composition $x(k)$ is an affine function of these decision parameters. Moreover, it can be proved (see Lemma 1) that the expectation of $x(k)$ is also affine in the parameters, and that the variance of the total wealth $w(k) = \mathbf{1}^\top x(k)$ is a convex quadratic function of the parameters. These desirable properties hold in general, for any stochastic and possibly dependent model of the market (i.e. for any stochastic description of the $G(k)$'s), and any given functional dependence of π_k on $g(1), \dots, g(k)$, provided that the parameterization of the policy family Π_k is affine in the parameters. Formally, we write the policy π_k as

$$u(k) = \pi_k(g(1), \dots, g(k)) = \sum_{i=1}^{n_k} f_{i,k}(g(1), \dots, g(k)) v_i(k), \quad (5)$$

where $v_1(k), \dots, v_{n_k}(k)$ are the decision parameters and $f_{i,k}(g(1), \dots, g(k))$, $i = 1, \dots, n_k$ are fixed basis functions. The following lemma holds for this class of policies.

Lemma 1. Consider the portfolio dynamic Eq. (1), with the control policy (5), and let $G(t) = \text{diag}(g(t))$, $t = 1, 2, \dots$ be a generic, possibly dependent, stochastic return process. Then, the portfolio expectation $E\{x(k)\}$ is an affine function of the policy parameters $\{v_1(t), \dots, v_{n_t}(t)\}_{t=0, \dots, k-1}$, and the total wealth variance $\text{var}\{w(k)\}$ is a convex quadratic function of these parameters.

A simple proof of this lemma is reported in the Appendix.

Remark 2. It follows from Lemma 1 that, under the policy class (5), the problem of minimizing $J(T)$ under the constraints described in Section 3.1 is a convex quadratic programming problem in the decision variables $\{v_1(t), \dots, v_{n_t}(t)\}_{t=0, \dots, k-1}$. The same would also be true for any other set of linear constraints on the expected portfolio compositions and on the adjustments.

However, obtaining an exact analytic formulation of this convex program is very hard in practice, since generic multi-dimensional expectations of market parameters need to be computed. Nevertheless, the required expectations can be estimated via sampling, hence an approximation of the problem (obtained by substituting the real expectations with their sample approximations) can be solved in practice with great numerical efficiency via convex quadratic programming. This is an important and practically relevant feature descending from the policy family (5). Notice also that if the information state $x(1), \dots, x(k)$ were used instead of $g(1), \dots, g(k)$ in the policy, the convexity structure would be lost. \diamond

While the approach described above leads to an approximate model that can be solved efficiently in practice, it still entails one level of approximation due to the sampling estimation of expected values of parameters related to market returns. In the next section, we propose a specific linear basis function setup within which it is possible to obtain explicit analytic expressions of the portfolio expectations and variances, and thus to derive an explicit and exact formulation of the multi-stage decision problem (under the assumed policy family) which is solvable efficiently by means of convex quadratic programming.

5. Explicit model with linear recourse

In this section, we develop an explicit closed-loop model for multi-period asset allocation, under additional hypotheses on the stochastic return model and on the policy parameterization.

Assumption 1 (Independent Gains). We shall make the following standard assumptions on the stochastic behavior of the asset gains.

1. $g_i(k_1)$ is statistically independent of $g_j(k_2)$, for all i, j and for all $k_2 \neq k_1$. In other words, the returns over different periods are assumed to be independent.
2. The first two moments of the total return vectors $g(k)$, $k = 1, 2, \dots$, are known.

These assumptions are compatible with the standard “efficient market hypothesis” (EMH), see Fama (1970); Malkiel (1987). \star

The expected gains and covariances over periods are defined as

$$\bar{g}(k) \doteq E\{g(k)\}, \quad k = 1, 2, \dots$$

$$\Sigma(k) \doteq \text{var}\{g(k)\} = M(k) - \bar{g}(k)\bar{g}^\top(k), \quad k = 1, 2, \dots,$$

with

$$M(k) \doteq E\{g(k)g^\top(k)\}, \quad k = 1, 2, \dots$$

In practice, the expected returns $\bar{g}(k)$ and covariances $\Sigma(k)$ are estimated by means of past time series analysis and/or elicited by expert advice. It should, however, be realized that these parameters are difficult to obtain with precision in reality (see, e.g., Section 8.5 of Luenberger (1997)), hence the results of an ensuing optimization model will always be affected by a level of “estimation uncertainty”. This issue is well known in classical single-stage portfolio optimization: for instance, a study of sensitivity of optimal portfolios to changes in the expected returns is performed in Best and Grauer (1991), while more recently some authors investigated the robustness of optimal portfolios to possibly non infinitesimal variations in the parameters of

the model, see Tütüncü and Koenig (2004), El Ghaoui, Oks, and Oustry (2003) and Goldfarb and Iyengar (2003). This kind of analysis is, however, outside the intended scope of this paper. Therefore, for the purpose of the optimization model developed here, expectation and covariances are assumed to be given input data.

Assumption 2 (Linear Recourse Policies). Consider recourse policies (5) with linear basis functions:

$$\pi_k(g(1), \dots, g(k)) = \bar{u}(k) + \sum_{i=1}^k \theta_i(k)(g(i) - \bar{g}(i)), \quad (6)$$

where the parameters are $\bar{u}(k) \in \mathbb{R}^n$ and $\theta_i(k) \in \mathbb{R}^{n,n}$.

In the following developments we assume a special case of policy (6) with a single-stage memory depth. That is, we consider the explicit control rule

$$u(k) = \bar{u}(k) + \theta(k)(g(k) - \bar{g}(k)), \quad k = 0, \dots, T-1. \quad \star \quad (7)$$

Remark 3. The control policy in (7) has the following interpretation: $\bar{u}(k)$ is the “nominal” adjustment action that we would perform at time k , if the market during period k performed as expected. Since the market will never perform exactly as expected, we correct the nominal decision $\bar{u}(k)$ with a term which is proportional to the market deviation from expectation ($g(k) - \bar{g}(k)$). The coefficients of the correction are collected in the market reaction matrix $\theta(k)$. In particular, element $\theta_{ij}(k)$ in row i and column j of matrix $\theta(k)$ represents the sensitivity of the control action in the i -th security, $u_i(k)$, with respect to deviations from expectation of the return of the j -th security. \diamond

By substituting the control law (7) into (1), we obtain the dynamic equations for the closed-loop controlled portfolio, and for its adjusted version:

$$x(k+1) = G(k+1)[x(k) + \bar{u}(k)] + G(k+1)\theta(k)[g(k) - \bar{g}(k)] \quad (8)$$

$$x^+(k) = x(k) + \bar{u}(k) + \theta(k)[g(k) - \bar{g}(k)].$$

From (8) we also easily find the recursion for the expected value of the portfolio and for its centered value

$$\bar{x}(k+1) = \bar{G}(k+1)[\bar{x}(k) + \bar{u}(k)] \quad (9)$$

$$\tilde{x}(k+1) = G(k+1)\tilde{x}(k) + \tilde{G}(k+1)[\bar{x}(k) + \bar{u}(k)] + G(k+1)\theta(k)\tilde{g}(k). \quad (10)$$

The above recursions are initialized with

$$\bar{x}(0) = x(0), \quad \tilde{x}(0) = 0, \quad \tilde{g}(0) = 0, \quad \theta(0) = 0.$$

An explicit recursion can also be derived for the portfolio covariance and for the wealth variance. The following key result holds (see the Appendix for a proof).

Lemma 2 (Closed-Loop Portfolio Expectation and Covariance). Let Assumption 1 be satisfied, and consider the portfolio dynamic Eq. (1), with the control policy (7). Then, the portfolio expectation $E\{x(k)\}$ is an affine function of the policy parameters $\{\bar{u}(t)\}_{t=0,\dots,k-1}$, and the total wealth variance $\text{var}\{w(k)\}$ is a convex quadratic function of $\{\bar{u}(t), \theta(t)\}_{t=0,\dots,k-1}$. Specifically, the portfolio expectation is

$$E\{x(k)\} = \bar{\Phi}(1, k)x(0) + \sum_{j=1}^k \bar{\Phi}(j, k)\bar{u}(j-1) \quad (11)$$

and the portfolio covariance obeys to the recursion

$$\begin{aligned} \Gamma(k+1) &\doteq E\{\tilde{x}(k+1)\tilde{x}^\top(k+1)\} \\ &= \Gamma(k) \odot M(k+1) + \theta(k)\Sigma(k)\theta^\top(k) \odot M(k+1) \\ &\quad + \text{diag}(\bar{x}(k-1) + \bar{u}(k-1))\Sigma(k)\theta^\top(k) \odot M(k+1) \\ &\quad + \theta(k)\Sigma(k)\text{diag}(\bar{x}(k-1) + \bar{u}(k-1)) \odot M(k+1) \\ &\quad + (\bar{x}(k) + \bar{u}(k))(\bar{x}(k) + \bar{u}(k))^\top \odot \Sigma(k+1), \\ &\quad k = 1, 2, \dots, T-1, \end{aligned} \quad (12)$$

where $M(k+1) = \Sigma(k+1) + \bar{g}(k+1)\bar{g}^\top(k+1)$, with the initialization

$$\theta(0) = 0, \quad \Gamma(0) = 0, \quad \Sigma(0) = 0, \quad \bar{x}(0) = x(0);$$

$$\Gamma(1) = (\bar{x}(0) + \bar{u}(0))(\bar{x}(0) + \bar{u}(0))^\top \odot \Sigma(1).$$

For the total wealth variance, it holds that

$$\begin{aligned} \text{var}\{w(k+1)\} &= \mathbf{1}^\top \Gamma(k+1) \mathbf{1} \\ &= \text{Tr}\Gamma(k)M(k+1) + \text{Tr}\Sigma(k)\theta^\top(k)M(k+1)\theta(k) \\ &\quad + 2\text{Tr}\text{diag}(\bar{x}(k-1) + \bar{u}(k-1))\Sigma(k)\theta^\top(k)M(k+1) \\ &\quad + (\bar{x}(k) + \bar{u}(k))^\top \Sigma(k+1)(\bar{x}(k) + \bar{u}(k)). \quad \star \end{aligned} \quad (13)$$

5.1. Explicit closed-loop formulation

Notice that, as stated by Lemma 2, the total wealth covariance $\text{var}\{w(k)\}$ is a convex quadratic function of $\bar{u}(j)$, $\theta(j)$, $j = 0, \dots, k-1$, hence the cost objective $J(T)$, which is a weighted sum of wealth covariances, is jointly convex in these variables. Thus, we can effectively devise a convex quadratic programming problem for solving multi-stage portfolio allocation in closed loop. The explicit form of this program is given next.

$$(\mathcal{P}) : \min_{\bar{u}(0), \dots, \bar{u}(T-1); \theta(1), \dots, \theta(T-1)} \sum_{k=1}^T \gamma(k) \text{var}\{w(k)\} \quad (14)$$

subject to:

$$\mathbf{1}^\top \bar{u}(k) = 0, \quad k = 0, \dots, T-1 \quad (15)$$

$$\mathbf{1}^\top \theta(k) = 0, \quad k = 1, \dots, T-1 \quad (16)$$

$$E\{w(T)\} \geq \Phi_b w(0), \quad (17)$$

$$\bar{b}(k) \leq E\{w(k)\} + \bar{u}(k) \leq \bar{b}(k), \quad k = 0, \dots, T-1 \quad (18)$$

$$\nu_{lb,j}(k) \mathbf{1}^\top E\{x(k)\} \leq \sum_{i \in c_j} (E\{x_i(k)\} + \bar{u}_i(k)) \leq \nu_{ub,j}(k) \mathbf{1}^\top E\{x(k)\}, \quad (19)$$

$$j = 1, \dots, \ell; k = 0, \dots, T-1.$$

The value of $E\{x(k)\}$ and $E\{w(k)\}$ needed in constraints (17)–(19) is obtained recursively from Eq. (9). The value of $\text{var}\{w(k)\}$ needed in the objective (14) is instead obtained from the covariance recursion (12) and using (13).

Remark 4. Notice from Eq. (9) that the expressions for the expected portfolios remain unchanged with respect to the open-loop approach, that is, the introduction of the recourse policy (7) does not change the expected portfolios and the expected wealth. What is indeed affected by the recourse policy is the covariance of the portfolios and, as a consequence, the wealth variance. Specifically, $\Gamma(k)$ depends on the sequence of past nominal adjustments $\bar{u}(j)$, $j = 0, \dots, k-1$ and on the sequence of past market reaction matrices $\theta(j)$, $j = 1, \dots, k-1$. These latter variables give us additional “degrees of freedom” that permit to further reduce the optimization cost, with respect to the open-loop situation. \diamond

5.2. Shrinking-horizon implementation

In a real-world investment endeavor, the proposed multi-stage decision model might be implemented using a shrinking-horizon scheme. That is, the optimization problem is solved repeatedly at each time k , over a shrinking horizon, and only the first computed adjustment is applied at k . More precisely, consider the following shrinking-horizon version of problem (14)–(19): at time k (for $k = 0, 1, \dots, T-1$), we solve problem (\mathcal{P}_k) below:

$$(\mathcal{P}_k) : \min_{\bar{u}(k), \dots, \bar{u}(T-1); \theta(k+1), \dots, \theta(T-1)} \sum_{j=k+1}^T \gamma(j) \text{var}\{w(j)\} \quad (20)$$

subject to:

$$\begin{aligned}
\mathbf{1}^\top \bar{u}(j) &= 0, \quad j = k, \dots, T-1 \\
\mathbf{1}^\top \theta(j) &= 0, \quad j = k+1, \dots, T-1 \\
E_k \{w(T)\} &\geq \Phi_{lb,k} w(k), \\
\bar{b}(j) &\leq E_k \{w(j)\} + \bar{u}(j) \leq \bar{b}(j), \quad j = k, \dots, T-1 \\
\nu_{lb,i}(j) \mathbf{1}^\top E_k \{x(j)\} &\leq \sum_{h \in c_i} (E_k \{x_h(j)\} + \bar{u}_h(j)) \leq \nu_{ub,i}(j) \mathbf{1}^\top E_k \{x(j)\}, \quad (21) \\
i &= 1, \dots, \ell; j = k, \dots, T-1.
\end{aligned}$$

and actuate the computed adjustment $u(k) = \bar{u}(k)$. The adjustment needed at $k+1$ is computed at time $k+1$, by re-solving the one-step-ahead version of the above problem, etc. Notice that the expectations and variances appearing in problem (20) and (21) are denoted by suffix k since they are actually conditioned to the information state at time k , that is they are conditioned upon $g(1), \dots, g(k)$.

The advantage of using a shrinking-horizon implementation is that at each decision stage we can profit from observations of actual market behavior during the preceding period, and use these information, for instance, to feed fresh estimates to the model and to update risk weights and return targets.

The way the model is updated at each stage is largely user-dependent, since the decision-maker might decide to arbitrarily modify his/her initial investment profile, in the presence of a particular market opportunity/situation. In the following section, we propose a simple model update strategy. The sole scope of this strategy is to define a repeatable experimental framework within which to test our decision model consistently.

5.2.1. A model update strategy

For the purpose of the numerical test presented in Section 6.1, we propose the following simple rules to be used to update the problem data (moment estimates, objective weights and target return) in a shrinking-horizon implementation of the multi-stage model.

Updating the moment estimates. At each time k , the optimization problem needs as input data the expected asset returns $\bar{g}(j)$ and covariance matrices $\Sigma(j)$ over the residual forward horizon $j = k+1, \dots, T$. These data can be recomputed (using the user's favorite estimation technique) by taking into account the most recent historical data.

In the numerical experiments of Section 6, we used the following simple technique: at each decision time k , we use one year of daily return data up to k , and assume that the forward expected returns $\bar{g}(k+1), \dots, \bar{g}(T)$ are constant and equal to the current empirical return expectation (i.e. the return computed by averaging one year of past historical returns data). Similarly, the covariances $\Sigma(k+1), \dots, \Sigma(T)$ are set to a constant matrix that represents the current empirical covariance, based on one-year-backwards data. Better techniques might certainly be devised for predicting the forward expected returns and covariances. However, from the point of view of our optimization model, these are input data and their estimation falls out of the intended scope of the discussion.

Updating the risk weights. At $k = 0$ we start with some pre-defined weights $\gamma_0(j) = \gamma(j)$, $j = 1, \dots, T$. At $k = 1$ one of these weights (namely, $\gamma_0(1)$) is dropped and “inherited” by the remaining weights, in proportion to their relevance:

$$\gamma_1(j) = \gamma_0(j) + \gamma_0(1) \frac{\gamma_0(j)}{\sum_{i=2}^T \gamma_0(i)}, \quad j = 2, \dots, T.$$

The same rule is used at the subsequent times $k = 2, 3, \dots$. This means that the weights are recursively computed according to the

rule

$$\gamma_k(j) = \gamma_{k-1}(j) + \gamma_{k-1}(k) \frac{\gamma_{k-1}(j)}{\sum_{i=j}^T \gamma_{k-1}(i)}, \quad j = k+1, \dots, T.$$

Notice that the sum of the weights $\gamma_k(k+1) + \dots + \gamma_k(T)$ remains constant, for all k .

Updating the target return. Let Φ_{lb} be the lower bound that we wish to impose initially on the end-of-horizon investment return, that is at $k = 0$ we set $\Phi_{lb,0} = \Phi_{lb}$ and solve problem (\mathcal{P}_0) . Two cases are possible: either problem (\mathcal{P}_0) is feasible, or not. If it is feasible, we solve it, implement the computed first-stage adjustment $\bar{u}(0)$, and wait one period. If (\mathcal{P}_0) is unfeasible, we need to adjust (lower) our return request. In our experiments, we adjusted the lower bound by first computing the maximum admissible return $\Phi_{ub,0}$ (to compute this, we maximize the return, with no constraint on risk), and then setting $\Phi_{lb,0}$ to a fixed fraction η (say, $\eta = 0.8$) of $\Phi_{ub,0}$, that is $\Phi_{lb,0} = \eta \Phi_{ub,0}$. With this adjustment, problem (\mathcal{P}_0) is certainly feasible, hence we solve it, implement the computed first-stage adjustment $\bar{u}(0)$, and wait one period.

At the next adjustment time (time $k = 1$) we know exactly (deterministically) the current value of total wealth $w(1)$. The actual return experienced by the investment at time $k = 1$ is thus $w(1)/w(0)$, hence if we want to keep up to the initial return target Φ_{lb} , we should try to impose $\Phi_{lb,1} = \frac{w(0)}{w(1)} \Phi_{lb}$. Since once again problem (\mathcal{P}_1) might be unfeasible with this constraint in place, we actually set $\Phi_{lb,1} = \frac{w(0)}{w(1)} \Phi_{lb}$, if feasible; or $\Phi_{lb,1} = \eta \Phi_{ub,1}$ otherwise, where $\Phi_{ub,1}$ is the maximum expected return achievable over the residual period, if no constraint on risk is imposed (this value can be computed in a numerically efficient way by solving a linear program). Proceeding in accordance to these rules, at the generic adjustment time k , we set

$$\Phi_{lb,k} = \begin{cases} \left(\frac{w(0)}{w(1)} \frac{w(1)}{w(2)} \dots \frac{w(k-1)}{w(k)} \right) \Phi_{lb} & \text{if feasible} \\ \eta \Phi_{ub,k} & \text{otherwise.} \end{cases}$$

We stress again that this procedure is reported only for the purpose of describing the numerical example in Section 6.1. This update rule insists in maintaining the return target originally designed for the problem. It may however happen that the actual path of the investment makes this plan unfeasible or not convenient. In these cases, the decision-maker would divert from his/her initial plans, and set the return target in accordance to the new market conditions and opportunities.

6. Numerical experiments

The numerical examples presented in this section have been developed in MATLAB Version 7.2.0.232 (R2006a) under operating system Microsoft Windows XP, on a quad Opteron workstation. The modelling language (parser) used to describe the problem is YALMIP (Löfberg, 2004), with the freely available CLP quadratic programming solver. Historical market data have been obtained from the Yahoo finance server.

6.1. Strategic allocation among sectors

Consider a fund manager facing a strategic problem of allocating investments among three equity sectors and cash. The allocation in these four sectors can be adjusted every $\Delta = 21$ trading days, and the investment horizon is equal to 12 periods of duration Δ (about one trading year). Thus, we have a multi-period allocation problem with $T = 12$ periods and $n = 4$ asset classes: (1) industrial, (2) utilities, (3) commodities, (4) cash.

The manager fixes a target objective of a 8% annual return for this portfolio, and implements blindly the shrinking-horizon




Fig. 3. Relative portfolio composition at the decision stages.

strategy described in Section 5.2. This investment approach is simulated on real data from 30-Dec-2003 to 4-Jan-2007. The following market indices were considered as proxies of the investment sectors: (1) Dow Jones Industrial Index (DJI), (2) Dow Jones Utilities Index (DJU), (3) Dow Jones Commodities Index (DJC), whereas cash (4) is assumed to provide a 3% risk-free annual interest rate.

The following parameters have been selected in the optimization model.

$\Phi_{lb} = 1.08$ annual return lower bound

$\eta = 0.9$ target return damping factor

$c_j = \{j\}, j = 1, \dots, n$ each security is a compartment

$\gamma(k) = 1, k = 1, \dots, T$ risk weights

$\bar{b}_i(k), \bar{b}_i(k), i = 1, \dots, n; k = 0, \dots, T-1$ not specified

$v_{ub,i}(k) = 1, i = 1, \dots, n; k = 0, \dots, T-1$

$v_{lb,i}(k) = 0, i = 1, \dots, n; k = 0, \dots, T-1$

no short-selling is allowed.

The simulation spans three years. Following the shrinking-horizon strategy, twelve instances of problem (\mathcal{P}_k) are solved for each year, and the investment decision computed for each period is actuated and held for $\Delta = 21$ days. Overall, 36 instances of problem (\mathcal{P}_k) have been solved, requiring a total of about 50 s computing time. The optimal portfolio relative weights computed at each decision stage are displayed in Fig. 3.

Fig. 4 shows the time value of one Dollar invested on 30-Dec-2003 in the individual component securities and in the optimal portfolio (PORT). It can be seen from Fig. 4 that the optimal portfolio would have been successful in maintaining a smooth and positive trend. Indeed, computing the average returns and volatilities for the component securities and the optimal portfolio, on a daily basis from 30-Dec-2003 to 4-Jan-2007, the investment opportunities result to be placed on the volatility-return plane as shown in Fig. 5.

Figs. 4 and 5 also show the performance of an open-loop optimal investment strategy, that is the performance of portfolios PORT(OL) computed via a shrinking-horizon implementation of the open-loop (no recourse) model of Section 3.2, feeded by the same parameter estimates that we used for computing the optimal portfolios PORT. It can be noticed that the recourse strategy indeed dominates the no-recourse one. However, since the shrinking-horizon actuation of the investment plans already incorporates a degree of reactivity to market movements, the difference between the open-loop and the recourse strategy may not be so apparent in simulations where market parameters have not been accurately estimated. The ability of the recourse strategy in reducing risk can be better appreciated when efficient frontiers are compared, as it is shown in Fig. 6 in the next example.




Fig. 4. Time value of component securities and of optimal portfolios.




Fig. 5. Risk-return characteristics of component securities and optimal portfolios.

6.2. The multi-period frontier

Problem (14)–(19) can be used effectively to construct a numerical approximation of the (constrained) efficient multi-period frontier. Let for instance $\gamma(k) = 0, k = 1, \dots, T-1$, and $\gamma(T) = 1$. Then, the optimal objective value (14) represents the minimal terminal wealth variance achievable, given the terminal return lower bound Φ_{lb} (and possibly given the other portfolio composition constraints). Solving repeatedly problem (14)–(19) for a range of increasing values of Φ_{lb} , and plotting the corresponding minimal risk level on a risk-return plane, one obtains a plot of a portion of the multi-period risk-return efficient frontier.

As a numerical example, consider a four-period allocation problem involving Equity, Bond and Cash (i.e., $T = 4, n = 3$). Specifically, we assume a horizon of one year, with the following expectations and covariances for the quarterly gains of the considered asset classes:

$$\bar{g}(1) = \begin{bmatrix} 1.0400 \\ 1.0100 \\ 1.0000 \end{bmatrix}, \quad \bar{g}(2) = \begin{bmatrix} 1.0500 \\ 1.0100 \\ 1.0000 \end{bmatrix}, \quad \bar{g}(3) = \begin{bmatrix} 1.0600 \\ 1.0150 \\ 1.0000 \end{bmatrix},$$

$$\bar{g}(4) = \begin{bmatrix} 1.0600 \\ 1.0150 \\ 1.0000 \end{bmatrix};$$

$$\Sigma(k) = (1 + 0.1(k-1)) \begin{bmatrix} 0.02 & -0.0008 & 0.0 \\ -0.0008 & 0.0016 & 0.0 \\ 0.0 & 0.0 & 0.0 \end{bmatrix},$$

$$k = 1, 2, 3, 4.$$

Solving problem (14)–(19), with the only additional no short-selling constraints, and with $x(0) = [0 \ 0 \ 1]^T$, for 40 increasing values of Φ_{lb} in the range $[1.035, 1.10]$, we obtain the portion

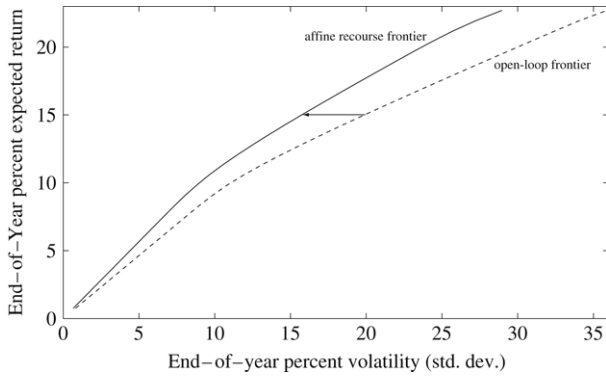


Fig. 6. Multi-period efficient frontiers for open-loop and affine recourse portfolios. The horizontal arrow highlights the reduction in risk, at 15% return level.

of multi-period efficient frontier shown in Fig. 6. This optimal closed-loop frontier can be compared with the efficient frontier one obtains from the open-loop (no recourse) multi-period model of Section 3.2. The target final return being the same, it is apparent from Fig. 6 how the closed-loop optimal strategies considerably reduce risk with respect to the open-loop optimal ones. Notice further that multi-period efficient frontiers are extremely difficult to obtain via the classical stochastic programming approaches, already for small number of periods; see for instance an example of experimental frontier computation via scenario sampling in Section 6 of Pinar (2007). Contrary, we remark that with the proposed method the approximated frontier in this example is obtained on a standard workstation in about two seconds, that is with extreme numerical efficiency.

To examine closer the computed recourse policies, consider for instance a 15% return level in Fig. 6, that is $\phi_{lb} = 1.15$. The linear policy assumption prescribes policies of the form

$$u(k) = \bar{u}(k) + \theta(k)(g(k) - \bar{g}(k)), \quad k = 0, 1, 2, 3; \quad \theta(0) = 0,$$

and, for $\phi_{lb} = 1.15$, problem (14)–(19) achieves the optimal objective value of $J(T) = \text{var}\{w(T)\} = 0.0248$, and returns the following optimal policy parameters:

$$\begin{aligned} \bar{u}(0) &= \begin{bmatrix} 0.6560 \\ 0.3440 \\ -1 \end{bmatrix}, & \bar{u}(1) &= \begin{bmatrix} 0.0285 \\ -0.0285 \\ 0 \end{bmatrix}, \\ \bar{u}(2) &= \begin{bmatrix} -0.1322 \\ 0.1322 \\ 0 \end{bmatrix}, & \bar{u}(3) &= \begin{bmatrix} -0.1788 \\ 0.1788 \\ 0 \end{bmatrix}; \\ \theta(1) &= \begin{bmatrix} -1.5108 & -0.4482 & 0 \\ -3.0000 & -1.9172 & 0 \\ 4.5108 & 2.3654 & 0 \end{bmatrix}, \\ \theta(2) &= \begin{bmatrix} -1.8437 & -0.5083 & 0 \\ -3.9075 & -2.0720 & 0 \\ 5.7512 & 2.5803 & 0 \end{bmatrix}, \\ \theta(3) &= \begin{bmatrix} -1.8783 & -0.9350 & 0 \\ -4.0735 & -3.4671 & 0 \\ 5.9518 & 4.4021 & 0 \end{bmatrix}. \end{aligned}$$

7. Conclusions

In this paper, we examined a multi-period version of the classical mean-variance portfolio optimization problem. In general, this is a hard infinite-dimensional dynamic optimization problem. We showed in Section 4 that if the decision policy uses the past returns as information states and if the policy is restricted to a finite-dimensional affine parametric family, then the multi-period

mean-variance problem can be cast as a finite-dimensional convex quadratic programming problem. The actual numerical solution of the problem in this setup, however, would still require sampling estimation of market expectations that cannot be determined analytically, in general.

In Section 5 we hence made further assumptions on the market behavior (statistic independence of returns among periods) and proposed an explicit linear decision rule in (7). Under these additional hypotheses, we derived an explicit analytic representation of the multi-period mean-variance problem as a constrained quadratic program. The practical effectiveness of this model has been tested numerically in the examples in Section 6.

While the proposed approach is sub-optimal with respect to a fully general search over infinite-dimensional unrestricted policy spaces, it does avoid resorting to coarse scenario approximations, it is fully analytic, and it leads to efficiently computable (polynomial-time) solutions. This latter aspect is particularly important, since it allows us to tackle practical decision problems involving many securities and periods. The proposed model can also be extended to account for proportional transaction costs, which is the subject of current research.

Appendix

Proof of Lemma 1. Substituting (5) into (2) we have

$$\begin{aligned} x(k) &= \Phi(1, k)x(0) \\ &\quad + \sum_{j=1}^k \sum_{i=1}^{n_{j-1}} \Phi(j, k) f_{i,j-1}(g(1), \dots, g(j-1)) v_i(j-1). \end{aligned}$$

Thus, taking expectation,

$$\begin{aligned} E\{x(k)\} &= E\{\Phi(1, k)x(0)\} \\ &\quad + \sum_{j=1}^k \sum_{i=1}^{n_{j-1}} E\{\Phi(j, k) f_{i,j-1}(g(1), \dots, g(j-1))\} v_i(j-1) \end{aligned}$$

we directly verify that $E\{x(k)\}$ is affine in the policy parameters $\{v_1(t), \dots, v_{n_t}(t)\}_{t=0, \dots, k-1}$. Similarly, defining

$$\begin{aligned} \tilde{\phi}^\top(1, k) &\doteq \phi^\top(1, k) - E\{\phi^\top(1, k)\}, \\ \tilde{\phi}_i^\top(j, k) &\doteq \phi^\top(j, k) f_{i,j-1}(g(1), \dots, g(j-1)) \\ &\quad - E\{\phi^\top(j, k) f_{i,j-1}(g(1), \dots, g(j-1))\}. \end{aligned}$$

we write the centered value of the total wealth

$$\begin{aligned} \tilde{w}(k) &= w(k) - E\{w(k)\} \\ &= \tilde{\phi}^\top(1, k)x(0) + \sum_{j=1}^k \sum_{i=1}^{n_{j-1}} \tilde{\phi}_i^\top(j, k) v_i(j-1) \\ &= \psi^\top(k) \begin{bmatrix} 1 \\ v(k-1) \end{bmatrix}, \end{aligned}$$

where

$$\begin{aligned} \psi^\top(k) &\doteq [\tilde{\phi}^\top(1, k)x(0) | \tilde{\phi}_1^\top(1, k) \cdots \tilde{\phi}_{n_0}^\top(1, k) | \tilde{\phi}_1^\top(2, k) \cdots \\ &\quad \tilde{\phi}_{n_1}^\top(2, k) | \cdots \tilde{\phi}_1^\top(k, k) \cdots \tilde{\phi}_{n_{k-1}}^\top(k, k)] \\ V^\top(k-1) &= [v_1^\top(0) \cdots v_{n_0}^\top(0) | v_1^\top(1) \cdots v_{n_1}^\top(1) | \cdots | \\ &\quad v_1^\top(k-1) \cdots v_{n_{k-1}}^\top(k-1)]. \end{aligned}$$

Hence, the wealth variance is

$$\begin{aligned} \text{var}\{\tilde{w}(k)\} &= E\{\tilde{w}^\top(k)\tilde{w}(k)\} \\ &= \begin{bmatrix} 1 \\ v(k-1) \end{bmatrix}^\top E\{\psi(k)\psi^\top(k)\} \begin{bmatrix} 1 \\ v(k-1) \end{bmatrix}. \end{aligned}$$

Since the covariance matrix $E\{\psi(k)\psi^\top(k)\}$ is positive semidefinite, it follows that $\text{var}\{\tilde{w}(k)\}$ is convex and quadratic in $V(k-1)$, which proves the statement. \square

Proof of Lemma 2. Expression (11) is obtained by recursive application of (8). From direct inspection of (11) it follows that $E\{x(k)\}$ is affine in the variables $\tilde{u}(t)$, $t = 0, \dots, k-1$.

Convexity of the wealth variance is established by Lemma 1. To build an explicit recursion for the controlled portfolio covariance we use (10) and we write

$$\begin{aligned} \Gamma(k+1) &\doteq E\{\tilde{x}(k+1)\tilde{x}^\top(k+1)\} \\ &\stackrel{(1)}{=} E\{G(k+1)\tilde{x}(k)\tilde{x}^\top(k)G(k+1)\} \\ &\quad + E\{G(k+1)\tilde{x}(k)[\tilde{x}(k) + \tilde{u}(k)]^\top \tilde{G}(k+1)\} \\ &\quad + E\{G(k+1)\tilde{x}(k)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} + (2)^\top \\ &\quad + E\{\tilde{G}(k+1)[\tilde{x}(k) + \tilde{u}(k)][\tilde{x}(k) + \tilde{u}(k)]^\top \tilde{G}(k+1)\} \\ &\quad + E\{\tilde{G}(k+1)[\tilde{x}(k) + \tilde{u}(k)]\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &\quad + (3)^\top + (6)^\top \\ &\stackrel{(9)}{=} E\{G(k+1)\theta(k)\tilde{g}(k)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\}. \end{aligned}$$

We now evaluate each term in the previous expression, recalling that $\tilde{x}(k)$ is independent of $G(k+1)$, but it depends on $\tilde{g}(k)$.

$$\begin{aligned} (1) \quad &E\{G(k+1)\tilde{x}(k)\tilde{x}^\top(k)G(k+1)\} \\ &= E\{\tilde{x}(k)\tilde{x}^\top(k)\} \odot E\{g(k+1)g^\top(k+1)\} \\ &= \Gamma(k) \odot M(k+1); \\ (2) \quad &E\{G(k+1)\tilde{x}(k)[\tilde{x}(k) + \tilde{u}(k)]^\top \tilde{G}(k+1)\} = 0; \\ (3) \quad &E\{G(k+1)\tilde{x}(k)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &= (3a) + (3b) + (3c) \quad (\text{to be specified later on}); \\ (5) \quad &E\{\tilde{G}(k+1)[\tilde{x}(k) + \tilde{u}(k)][\tilde{x}(k) + \tilde{u}(k)]^\top \tilde{G}(k+1)\} \\ &= E\{[\tilde{x}(k) + \tilde{u}(k)][\tilde{x}(k) + \tilde{u}(k)]^\top\} \odot E\{\tilde{g}(k+1)\tilde{g}^\top(k+1)\} \\ &= [\tilde{x}(k) + \tilde{u}(k)][\tilde{x}(k) + \tilde{u}(k)]^\top \odot \Sigma(k+1); \\ (6) \quad &E\{\tilde{G}(k+1)[\tilde{x}(k) + \tilde{u}(k)]\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &= E\{[\tilde{x}(k) + \tilde{u}(k)]\tilde{g}^\top(k)\theta^\top(k) \odot \tilde{g}(k+1)\tilde{g}^\top(k+1)\} = 0; \\ (9) \quad &E\{G(k+1)\theta(k)\tilde{g}(k)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &= E\{\theta(k)\tilde{g}(k)\tilde{g}^\top(k)\theta^\top(k) \odot g(k+1)g^\top(k+1)\} \\ &= \theta(k)\Sigma(k)\theta^\top(k) \odot M(k+1). \end{aligned}$$

To write item (3) explicitly, we substitute $\tilde{x}(k)$ with its expression (10):

$$\begin{aligned} \tilde{x}(k) &= G(k)\tilde{x}(k-1) + \tilde{G}(k)[\tilde{x}(k-1) + \tilde{u}(k-1)] \\ &\quad + G(k)\theta(k-1)\tilde{g}(k-1) \end{aligned}$$

thus obtaining

$$\begin{aligned} (3) \quad &E\{G(k+1)\tilde{x}(k)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &\stackrel{(3a)}{=} E\{G(k+1)G(k)\tilde{x}(k-1)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \end{aligned}$$

$$\begin{aligned} &\stackrel{(3b)}{+} E\{G(k+1)\tilde{G}(k)[\tilde{x}(k-1) + \tilde{u}(k-1)]\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &\stackrel{(3c)}{+} E\{G(k+1)G(k)\theta(k-1)\tilde{g}(k-1)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\}. \end{aligned}$$

Since $\tilde{x}(k-1)$ is independent of $G(k)$, $\tilde{g}(k)$, and $G(k+1)$, we have that

$$\begin{aligned} (3a) \quad &E\{G(k+1)G(k)\tilde{x}(k-1)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} = 0; \\ (3b) \quad &E\{G(k+1)\tilde{G}(k)[\tilde{x}(k-1) + \tilde{u}(k-1)]\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} \\ &= E\{\tilde{G}(k)[\tilde{x}(k-1) + \tilde{u}(k-1)]\tilde{g}^\top(k)\theta^\top(k)\} \\ &\quad \odot E\{\tilde{g}(k+1)\tilde{g}^\top(k+1)\} \\ &= E\{\text{diag}(\tilde{x}(k-1) + \tilde{u}(k-1))\tilde{g}(k)\tilde{g}^\top(k)\theta^\top(k)\} \odot M(k+1) \\ &= (\text{diag}(\tilde{x}(k-1) + \tilde{u}(k-1))\Sigma(k)\theta^\top(k)) \odot M(k+1); \\ (3c) \quad &E\{G(k+1)G(k)\theta(k-1)\tilde{g}(k-1)\tilde{g}^\top(k)\theta^\top(k)G(k+1)\} = 0. \end{aligned}$$

Putting it all together, we obtain (12), (13), which concludes the proof. \square

References

- Ait-Sahalia, Y., & Brandt, M. W. (2001). Variable selection for portfolio choice. *Journal of Finance*, 56(4), 1297–1351.
- Ben-Tal, A., Goryashko, A., Guslitzer, E., & Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming, Series A*, 99, 351–376.
- Bertsekas, D. P., & Tsitsiklis, J. (1996). *Neuro-dynamic programming*. Belmont, MA: Athena Scientific.
- Best, M. J., & Grauer, R. R. (1991). Sensitivity analysis for mean-variance portfolio problems. *Management Science*, 37(8), 980–989.
- Birge, John R. (1997). Stochastic programming computation and applications. *INFORMS Journal on Computing*, 9(2), 111–133.
- Birge, J. R., & Louveaux, F. (1997). *Introduction to stochastic programming*. Springer Verlag.
- Brandt, M. W. (1999). Estimating portfolio and consumption choice: A conditional Euler equations approach. *Journal of Finance*, 54(5), 1609–1645.
- Brennan, M. J., Schwartz, E. S., & Lagnado, R. (1997). Strategic asset allocation. *Journal of Economic Dynamics and Control*, 21, 1377–1403.
- Calafiore, G., & Campi, M. C. (2005). On two-stage portfolio allocation problems with affine recourse. In *Joint 44th IEEE conference on decision and control and European control conference*, Seville, Spain.
- Calafiore, G. C., & El Ghaoui, L. (2006). On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, 130(1), 1–22.
- Dantzig, G. B., & Infanger, G. (1993). Multi-stage stochastic linear programs for portfolio optimization. *Annals of Operations Research*, 45, 59–76.
- El Ghaoui, L., Oks, M., & Oustry, F. (2003). Worst-case value-at-risk and robust portfolio optimization: A conic programming approach. *Operations Research*, 51(4), 543–556.
- Ermoliev, Y., & Wets, R. J. (Eds.). (1988). *Numerical techniques for stochastic optimization*. Springer.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Farias, V. F., & Van Roy, B. (2006). Tetris: A study of randomized constraint sampling. In G. C. Calafiore, & F. Dabbene (Eds.), *Probabilistic and randomized methods for design under uncertainty*. London: Springer-Verlag.
- Goldfarb, D., & Iyengar, G. (2003). Robust portfolio selection problems. *Mathematics of Operations Research*, 28(1), 1–38.
- Gulpinar, N., Rustem, B., & Settergren, R. (2002). Multistage stochastic programming in computational finance. In *Decision making, economics and finance: Optimization models* (pp. 33–47). Kluwer Academic Publishers.
- Gulpinar, N., Rustem, B., & Settergren, R. (2004). Optimization and simulation approaches to scenario tree generation. *Journal of Economics, Dynamics and Control*, 28(7), 1291–1315.
- Hibiki, N. (2006). Multi-period stochastic optimization models for dynamic asset allocation. *Journal of Banking and Finance*, 30, 365–390.
- Infanger, G. (2006). Dynamic asset allocation strategies using a stochastic dynamic programming approach. In S. Zenios, & W. T. Ziemba (Eds.), *Handbook of asset liability management*. Elsevier.
- Löfberg, J. (2003). *Minimax approaches to robust model predictive control*. Ph.D thesis, Linköping University, Sweden.
- Löfberg, J. (2004). Yalmip: A toolbox for modeling and optimization in MATLAB. In *Proceedings of the CACSD conference*, Taipei, Taiwan.
- Luenberger, David G. (1997). *Investment science*. Oxford University Press.
- Lynch, A. W., & Balduzzi, P. (2000). Predictability and transaction costs: The impact on rebalancing rules and behavior. *Journal of Finance*, 50(5), 2285–2309.

- Malkiel, B. G. (1987). Efficient market hypothesis. In M. Milgate, J. Eatwell, & P. Newman (Eds.), *The new palgrave: A dictionary of economics: Vol. 2* (pp. 120–123). London and New York: Macmillan & Stockton.
- Markowitz, H. M. (1991). *Portfolio selection*. Blackwell.
- Merton, R. C. (1971). Optimum consumption and portfolio rules in continuous time model. *Journal of Economic Theory*, 3, 373–413.
- Merton, R. C. (1973). An intertemporal asset pricing model. *Econometrica*, 41, 867–887.
- Pflug, G. Ch. (2001). Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming*, 89, 251–271.
- Pinar, M. (2007). Robust scenario optimization based on downside-risk measure for multi-period portfolio selection. *OR Spectrum*, 29, 295–309.
- Rustem, B., & Gulpinar, N. (2007). Worst-case optimal robust decisions for multi-period portfolio optimization. *European Journal of Operational Research*.
- Ruszczynski, A., & Shapiro, A. (Eds.) (2003). *Handbooks in operations research and management science: Vol. 10. Stochastic programming*. Amsterdam: Elsevier.
- Shapiro, A. (2005). On complexity of multistage stochastic programs. Optimization online http://www.optimization-online.org/DB_HTML/2005/01/1046.html.
- Steinbach, M. C. (2001). Markowitz revisited: Mean-variance models in financial portfolio analysis. *SIAM Review*, 43(1), 31–85.
- Takriti, S., & Ahmed, S. (2004). On robust optimization of two-stage systems. *Mathematical Programming, Ser. A*, 99, 109–126.
- Tütüncü, R. H., & Koenig, M. (2004). Robust asset allocation. *Annals of Operations Research*, 132, 157–187.
- Zhou, X., & Li, D. (2000). Continuous time mean-variance portfolio selection: A stochastic LQ framework. *Applied Mathematics and Optimization*, 42, 19–33.
- Ziemba, W. T., & Mulvey, J. M. (Eds.) (1998). *World wide asset and liability management*. Cambridge Univ. Press.



Giuseppe Carlo Calafiore received the “Laurea” degree in Electrical Engineering from Politecnico di Torino in 1993, and the Doctorate degree in Information and System Theory from Politecnico di Torino in 1997. Dr. Calafiore currently serves as Associate Professor at the “Dipartimento di Automatica e Informatica,” Politecnico di Torino, Italy.

He held visiting positions at Information Systems Laboratory, Stanford University, in 1995, at Ecole Nationale Supérieure de Techniques Avancées (ENSTA), Paris, in 1998, and at the University of California at Berkeley, in 1999, 2003 and 2007.

Dr. Calafiore is Associate Editor for the IEEE Transactions on Automation Science and Engineering and for the IEEE Transactions on Systems, Man, and Cybernetics. His research interests are in the field of analysis and control of uncertain systems, convex optimization, probabilistic algorithms and risk management. He is coauthor of 6 books and over 100 publications on international journals and conference proceedings.