

Optimal Dynamic Asset Allocation with Lower Partial Moments Criteria and Affine Policies

Original

Optimal Dynamic Asset Allocation with Lower Partial Moments Criteria and Affine Policies / Calafiore, Giuseppe Carlo. -
In: INTERNATIONAL JOURNAL OF FINANCIAL ENGINEERING AND RISK MANAGEMENT. - ISSN 2049-0909. -
STAMPA. - 2:2(2015), pp. 87-108. [10.1504/IJFERM.2015.074040]

Availability:

This version is available at: 11583/2634570 since: 2016-02-23T10:31:18Z

Publisher:

Inderscience

Published

DOI:10.1504/IJFERM.2015.074040

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Optimal Dynamic Asset Allocation with Lower Partial Moments Criteria and Affine Policies*

Giuseppe Carlo Calafiore[†]

Abstract

This paper discusses an optimization-based approach for solving multi-period dynamic asset allocation problems using empirical asymmetric measures of risk. Three features distinguish the proposed approach from the mainstream ones. First, our approach is *non parametric*, in the sense that it does not require explicit estimation of the parameters of a statistical model for the returns distribution: the approach relies directly on *data* (the scenarios) generated by an oracle which may include expert knowledge along with a standard stochastic return model. Second, it employs affine decision policies, which make the multi-period formulation of the problem amenable to an efficient convex optimization format. Third, it uses asymmetric, unilateral measures of risk which, unlike standard symmetric measures such as variance, capture the fact that investors are usually not averse to return deviations from the expected target, if these deviations actually exceed the target.

1 Introduction

The dynamic asset allocation methodology discussed in this paper belongs to the class of data-driven techniques, which use return data directly in order to numerically compute the optimal investment decisions. Mainstream model-based approaches, derived from the classical Markowitz setup (see, e.g., [22]), focus on parameters of the return distribution (such as expected returns and covariance) that need to be estimated, and then derive the optimal investment decisions on the basis of these parameters. It is for instance well known (see, e.g., [3], [14]) that the allocation resulting from “classical” approaches is sensitive to the estimated model parameters (e.g., expected returns and covariances, or other parameters of the elicited return distribution). As a consequence, due to model estimation errors, a portfolio designed to have certain desired characteristics may well fail to provide the expected performance “out of sample,” that is on new, future scenarios that have not been accounted for at the model estimation stage.

*This work has been funded in part by Fondaco SGR S.p.A.

[†]Full professor at Dipartimento di Automatica e Informatica, Politecnico di Torino, Italy.
E-mail: giuseppe.calafiore@polito.it

Data-driven approaches (see, e.g., [10, 18, 20, 21], for applications in single-period financial problems), in contrast, do not necessarily or explicitly require the intermediate step of estimating a statistical model of the returns: they focus instead on *data*, and aim at determining the optimal allocations using the data directly, as illustrated schematically in Figure 1. We observe that while it is true that the data (scenarios) can (and typically are) generated by a simple parametric statistical model of the returns, this is not necessary: scenarios can be generated by a mixture of parametric statistical models and expert knowledge, and may thus represent situations that are possibly more rich and realistic than those captured by, say, a parametric model based on the first two moments of the returns distribution.




Figure 1: Model based approach (top) vs. data-driven approach (bottom).

Typical multi-period decision problems are initially cast as model-based optimization problems, but are later transformed into data-driven ones, at the stage when numerical solution is required. Indeed, the usual computational approach to solve recursive decision problems in presence of uncertainty is given by multi-stage stochastic programming, see, e.g., [7, 8, 19, 24] and the many references therein. However, while stochastic programming provides a conceptually sound framework for posing multi-stage decision problems, it is hard to solve numerically (see, e.g., [25]); the key difficulty stemming from the fact that exponentially growing “scenario trees” need to be introduced in order to model approximately the conditional nature of the decision problem. This limits in practice the applicability of stochastic programming techniques to financial decision problems with few decision stages.

In this paper, we follow an alternative approximation approach, based on restricting the reaction policies to have a prescribed structure (in particular, an affine structure). This idea emerged in the context of robust optimization (see, e.g., [5, 15]), and has been successfully applied in the context of multi-stage financial decision problems in, e.g., [9] and [12, 13]. In these latter references, the author uses a model-based approach, and decision criteria based on symmetric measures of risk (variance). The contribution of the present work is to extend such a model to a data-driven framework, with possibly asymmetric risk measures. The decision model we explore is based on the use of first and second order empirical lower partial moments (LPM), which capture the risk averse-

ness of the investor only for negative deviations of the returns from the target return. More precisely, we consider a decision problem over T periods, where at each period we have the opportunity of rebalancing our portfolio allocation, with the objective of obtaining a minimum level of a suitable cost function at the final stage, while guaranteeing satisfaction of portfolio constraints at each stage. The cost function we consider is of the form

$$\text{LPM}_\nu \doteq \frac{1}{N} \sum_{i=1}^N \left(\max(0, \gamma - \varrho^{(i)}) \right)^\nu, \quad (1)$$

where ν is either 1 (first-order LPM model) or 2 (second-order LPM model), γ is a given desired gain level for our investment strategy at the final period T , and $\varrho^{(i)}$ is the final gain under the i -th scenario; see Section 3.2 for a precise definition of all terms. The gains $\varrho^{(i)}$ depend on our stage decisions, and the objective is to devise such decisions so to make the cost LPM_ν minimal. For $\nu = 1$, the cost LPM_1 penalizes strategies resulting in gains that, on average, fall short of the target gain γ . Contrary to standard, symmetric measures of risk, LPM_1 does not penalize gains *above* the target, and this should model the fact that such excess gains are indeed usually welcome by investors. Cost LPM_2 acts in a similar way, but now the negative deviations from the target are squared: this makes the cost more sensitive to large deviations and thus tends to provide more “cautious” strategies. Full justification for the use of such asymmetric measures is given, e.g., in [2, 17, 23]; see also [4].

Whether an LPM_1 or an LPM_2 criterion should be employed is a matter related to the risk averseness of the decision maker. In this paper, we propose efficient numerical solution models for both criteria. In particular, we shall develop such decision models under two frameworks that we name *open loop* and *closed loop*. These two approaches are described in detail in Section 3.2 and in Section 4, respectively. In the open-loop approach all stage decisions are virtually taken at the initial time, while in the closed-loop approach the stage decision may change according to an affine policy, thus adapting to market fluctuations as time moves towards the final stage. The use of affine policies for multi-stage financial allocation problems in conjunction with asymmetric risk measures is a novelty in the literature. Our terminology, referring to “open-loop” and “closed-loop” approaches, is perhaps nonstandard in the financial literature: it is borrowed from control systems engineering, where the closed-loop approach is related to *feedback* (see, e.g., [1]), which is a mechanism whereby system’s outputs are measured and used to devise suitable inputs for the system itself. In the present context, closed-loop indeed refers, by analogy, to an approach where “measurements” of the returns that become available during the optimization horizon are dynamically used in order to compute control actions (i.e., portfolio adjustments), while open-loop refers to the approach where such information is neglected.

We remark that we do not consider transaction costs in this work. A simple scheme of proportional transaction costs, however, can be easily introduced in our model, while more elaborate cost structures are likely to destroy the convexity of the model.

2 Definitions and preliminaries

2.1 Return and gain vectors

We denote with a_1, \dots, a_n , a collection of n assets, and with $p_i(k)$ the market price of a_i at time $k\Delta$, where k is an integer, and Δ is a fixed period of time. The simple return of an investment in asset i over the k -th period, from $(k-1)\Delta$ to $k\Delta$, is

$$r_i(k) \doteq \frac{p_i(k) - p_i(k-1)}{p_i(k-1)}, \quad i = 1, \dots, n; \quad k = 1, 2, \dots,$$

and the corresponding gross return, or *gain*, is defined as

$$g_i(k) \doteq 1 + r_i(k), \quad i = 1, \dots, n; \quad k = 1, 2, \dots$$

We denote with $r(k) \doteq [r_1(k) \ \dots \ r_n(k)]^\top$ the vector of assets' returns over the k -th period, and with $g(k)$ the corresponding vector of gains. The notation $G(k) = \text{diag}(g(k))$ indicates a diagonal matrix having the elements of $g(k)$ in the diagonal.

The return and gain vectors are assumed to be random quantities, and we denote with $\mathbb{P}_{1,2,\dots}$ the probability distribution of $\{r(1), r(2), \dots\}$, given the past $\{\dots, r(-1), r(0)\}$, where $k = 0$ denotes the current time, at which the portfolio allocation decision is to be taken. We let $T \geq 1$ denote the number of forward periods over which the allocation decisions need to be taken, and we let \mathbb{P} denote the marginal joint probability distribution of $\{r(1), \dots, r(T)\}$ given the past. We further set, for compactness of notation, $\mathcal{T} \doteq \{0, \dots, T-1\}$, and $\mathcal{N} \doteq \{1, \dots, N\}$, where N is the number of scenarios.

2.2 Scenario-generating oracle

We do not assume that \mathbb{P} is known. We only assume that there is available a *scenario-generating oracle* (SGO), which is capable of generating independent and identically distributed (iid) samples of the forward return streams $\{r(1), \dots, r(T)\}$, according to \mathbb{P} . It is important to observe that the mechanism inside the SGO may include both a classical, statistically estimated model of the returns, as well as any type of “expert knowledge,” which substantiates in the inclusion, with a given probability, of certain specific return paths in the SGO.

2.3 Portfolio vector and constraints

A *portfolio* of assets a_1, \dots, a_n is defined by a vector $x(k) \in \mathbb{R}^n$ whose entry $x_i(k)$, $i = 1, \dots, n$, describes the (signed) amount of an investor's wealth invested in asset a_i at time $k \in \mathcal{K}$, where $x_i(k) \geq 0$ denotes a “long” position, and $x_i(k) < 0$ denotes a “short” position. In portfolio design, the portfolio vector $x(k)$ is typically subject to various constraints, reflecting the investor's a-priori policies and bindings. For example, short-selling might be forbidden, in which

case the components of $x(k)$ must be nonnegative (which we write as $x(k) \geq 0$, with element-wise inequality), or the portfolio should be self-financing (the sum of portfolio entries must be equal to a constant), or yet constraints may include minimum and maximum exposure in an individual asset, or limits in the exposure over classes of assets, etc. In this paper, we shall treat the problem in reasonable generality by assuming that the portfolio vector is constrained in a polytope (a bounded polyhedron) $\mathcal{X}(k)$.

3 Open-loop portfolio dynamics

We consider a decision problem over T periods (or *stages*), where at each period we have the opportunity of rebalancing our portfolio allocation, with the objective of obtaining a minimum level of a suitable cost function (to be discussed later) at the final stage, while guaranteeing satisfaction of portfolio constraints at each stage.

Consider a decision horizon of T periods, where the k -th period starts at time $k - 1$ and ends at time k , see Figure 2.




Figure 2: Investment periods.

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 portfolio at time k is the vector

$$x(k) \doteq [x_1(k) \quad \cdots \quad x_n(k)]^\top.$$

The investor's total wealth at time k is

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

where $\mathbf{1}$ denotes a vector of ones. Let $x(0)$ be the given initial portfolio composition at time $k = 0$ (for example, one may assume that $x(0)$ is all zeros, except for one entry representing the initial available amount of cash). At $k = 0$, we have the opportunity of conducting transactions on the market and therefore adjusting the portfolio by increasing or decreasing the amount invested in each asset. Just after transactions, the adjusted portfolio is $x^+(0) = x(0) + u(0)$, where $u_i(0) > 0$ if we increase the position on the i -th asset, $u_i(0) < 0$ if we decrease it, and $u_i(0) = 0$ if we leave it unchanged. Suppose now that the portfolio is held fixed 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 $G(1) = \text{diag}(g_1(1), \dots, g_n(1))$ 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 $u(1)$ 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 $(k + 1)$, for $k \in \mathcal{K}$

$$x(k + 1) = G(k + 1)x(k) + G(k + 1)u(k), \quad (2)$$

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

$$x^+(k) = x(k) + u(k). \quad (3)$$

From (2) it results that the (random) portfolio composition at time $k = 1, \dots, T$, is

$$\begin{aligned} x(k) &= \Phi(1, k)x(0) + \\ &\quad \left[\begin{array}{cccc} \Phi(1, k) & \cdots & \Phi(k-1, k) & \Phi(k, k) \end{array} \right] \begin{bmatrix} u(0) \\ \vdots \\ u(k-2) \\ u(k-1) \end{bmatrix} \\ &= \Phi(1, k)x(0) + \Omega_k \mathbf{u}, \end{aligned} \quad (4)$$

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

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

and

$$\begin{aligned} \mathbf{u} &\doteq \left[u(0)^\top \quad \cdots \quad u(T-2)^\top \quad u(T-1)^\top \right]^\top, \\ \Omega_k &\doteq \left[\begin{array}{cccc|ccc} \Phi(1, k) & \cdots & \Phi(k-1, k) & \Phi(k, k) & 0 & \cdots & 0 \end{array} \right]. \end{aligned}$$

We thus have for the total wealth

$$w(k) = \mathbf{1}^\top x(k) = \phi(1, k)^\top x(0) + \omega_k^\top \mathbf{u},$$

where $\phi(\eta, k)^\top \doteq \mathbf{1}^\top \Phi(\eta, k)$, and

$$\omega_k^\top \doteq \mathbf{1}^\top \Omega_k = [\phi(1, k)^\top \quad \cdots \quad \phi(k-1, k)^\top \quad \phi(k, k)^\top \mid 0 \quad \cdots \quad 0].$$

We consider the portfolio to be self-financing, that is

$$\sum_{i=1}^n u_i(k) = 0, \quad k \in \mathcal{K},$$

and we include generic linear constraints in the model by imposing that the updated portfolios $x^+(k)$ lie within a given polytope $\mathcal{X}(k)$. The cumulative gross return of the investment over the whole horizon is

$$\varrho(\mathbf{u}) \doteq \frac{w(T)}{w(0)} = \frac{\mathbf{1}^\top x(T)}{\mathbf{1}^\top x(0)} = \frac{\phi(1, T)^\top x(0)}{\mathbf{1}^\top x(0)} + \frac{1}{\mathbf{1}^\top x(0)} \omega_T^\top \mathbf{u}. \quad (5)$$

We see that $\varrho(\mathbf{u})$ is an affine function of the decision variables \mathbf{u} , with a random vector ω_T of coefficients that depends on the random gains over the T periods.

3.1 Scenarios and cost criteria

Suppose that N iid samples (scenarios) $\{G^{(i)}(k), k = 1, \dots, T\}$, $i \in \mathcal{N}$, of the period gains are available from a scenario generating oracle. These samples produce in turn N scenarios for each of the Ω_k matrices, $k = 1, \dots, T$, and hence of the ω_k and $\phi(1, k)$ vectors. We denote such scenarios with $\Omega_k^{(i)}$, $\omega_k^{(i)}$, $\phi^{(i)}(1, k)$, $i \in \mathcal{N}$, and with $x^{(i)}(k)$, $w^{(i)}(k)$, $\varrho^{(i)} = \varrho^{(i)}(\mathbf{u})$, respectively, the portfolio composition at time k , the total wealth at time k , and the cumulative final gain, under the i -th scenario. Let further γ denote a given desired level of gain at the final stage: we define the following two empirical partial moments (LPM) for the gross return distribution at the final stage:

$$\text{LPM}_1 \doteq \frac{1}{N} \sum_{i=1}^N \max(0, \gamma - \varrho^{(i)}) \quad (6)$$

$$\text{LPM}_2 \doteq \frac{1}{N} \sum_{i=1}^N \left(\max(0, \gamma - \varrho^{(i)}) \right)^2. \quad (7)$$

Clearly, LPM_1 is a cost function measuring the empirical average of the return values $\varrho^{(i)}$ falling below level γ , while the LPM_2 cost measures the average of the squares of the same deviations. The choice of the first or second order LPM cost depends on the level of *risk aversion* of the investor, the higher degree in the LPM reflecting higher levels of risk aversion, due to the fact that large residuals are squared, and hence weight more on the cost, in LPM_2 .

3.2 Optimal open-loop allocation

Our open-loop multi-stage allocation strategy is determined by finding the adjustments $\mathbf{u} = (u(0), \dots, u(T-1))$ that minimize either LPM_1 or LPM_2 , subject to given portfolio composition constraints at each period of the investment horizon. This strategy is denoted as “open loop” since all decisions

$\mathbf{u} = (u(0), \dots, u(T-1))$ are computed at “time 0” and, in principle, they should next be executed without observing the market behavior during the $\{0, T-1\}$ horizon. Notice that this will hardly happen in practice, since these decisions are typically re-computed at the beginning of each period and executed in a receding-horizon fashion (see Section 4.3); however, we name this strategy “open loop” in order to distinguish it from the “closed-loop” strategy discussed in Section 4, where we optimize over a class of *policies* rather than on direct actions. In the open-loop approach we thus need to solve a problem of the form

$$\begin{aligned} \text{LPM}_\nu^{\text{ol}}(\gamma) = \min_{\mathbf{u}} \quad & \frac{1}{N} \sum_{i=1}^N (\max(0, \gamma - \varrho^{(i)}(\mathbf{u})))^\nu & (8) \\ \text{s.t.} \quad & x^{(i)+}(k) \in \mathcal{X}(k), \quad k \in \mathcal{T}, \quad i \in \mathcal{N} \\ & \mathbf{1}^\top u(k) = 0, \quad k \in \mathcal{T}, \end{aligned}$$

where either $\nu = 1$ (for the LPM_1 cost) or $\nu = 2$ (for the LPM_2 cost), $\varrho^{(i)}(\mathbf{u})$ is given by (5) under the i -th scenario, and where $x^{(i)+}(k)$ is given by (3), (4), under the i -th scenario. These optimal allocations may be determined in a numerically efficient way by solving, respectively, a linear programming or a convex quadratic programming problem, as detailed in the following propositions.

Proposition 1 *For $\nu = 1$ (LPM_1 cost), problem (8) is equivalent to the following linear programming problem in the variables $\mathbf{u} \in \mathbb{R}^{Tn}$ and $z \in \mathbb{R}^N$*

$$\begin{aligned} \text{LPM}_1^{\text{ol}}(\gamma) = \min_{\mathbf{u}, z} \quad & \frac{1}{N} \sum_{i=1}^N z_i & (9) \\ \text{s.t.} \quad & \Phi^{(i)}(1, k)x(0) + \Omega_k^{(i)} \mathbf{u} + u(k) \in \mathcal{X}(k), \quad k \in \mathcal{T}, \quad i \in \mathcal{N} \\ & \mathbf{1}^\top u(k) = 0, \quad k \in \mathcal{T}, \\ & z_i \geq 0, \quad i \in \mathcal{N}, \\ & z_i \geq \gamma - \varrho^{(i)}(\mathbf{u}), \quad i \in \mathcal{N}, \end{aligned}$$

where $\varrho^{(i)}(\mathbf{u}) = \frac{\phi^{(i)}(1, T)^\top x(0)}{\mathbf{1}^\top x(0)} + \frac{1}{\mathbf{1}^\top x(0)} \omega_T^{(i)\top} \mathbf{u}$.

Proposition 2 *For $\nu = 2$ (LPM_2 cost), problem (8) is equivalent to the following convex quadratic programming problem in the variables $\mathbf{u} \in \mathbb{R}^{Tn}$ and $z \in \mathbb{R}^N$*

$$\begin{aligned} \text{LPM}_2^{\text{ol}}(\gamma) = \min_{\mathbf{u}, z} \quad & \frac{1}{N} \sum_{i=1}^N z_i^2 & (10) \\ \text{s.t.} \quad & \Phi^{(i)}(1, k)x(0) + \Omega_k^{(i)} \mathbf{u} + u(k) \in \mathcal{X}(k), \quad k \in \mathcal{T}; \quad i \in \mathcal{N} \\ & \mathbf{1}^\top u(k) = 0, \quad k \in \mathcal{T}, \\ & z_i \geq 0, \quad i \in \mathcal{N}, \\ & z_i \geq \gamma - \varrho^{(i)}(\mathbf{u}), \quad i \in \mathcal{N}. \end{aligned}$$

The statements of the previous two propositions follow easily by adding slack variables and applying an epigraphic transformation to the objective in (8), see, e.g., Section 8.3.4.4 in [11].

4 Closed-loop portfolio dynamics

The open-loop strategy discussed in the previous section is suboptimal, since all adjustment decisions $u(0), \dots, u(T-1)$ are computed at time $k=0$ and then executed forward without feedback from the actual market behavior. While the first decision $u(0)$ must be immediately implemented (here-and-now variable), the future decisions may well wait-and-see the actual outcomes of the returns in the forward periods, and hence benefit from the uncertainty reduction that comes from these observations, see, e.g., [26]. For example, at time $k \geq 1$, when we need to implement $u(k)$, we have *observed* a realization of the asset returns over the periods from 1 to k . Hence, we would like to exploit this information, by considering *conditional* allocation decisions $u(k)$, that may react to the returns observed over the previous periods. This means that, instead of focusing on fixed decisions $u(k)$, we wish to determine suitable *policies* that prescribe what the actual decision should be, in dependence of the observed returns from 1 up to k . In determining the structure of the decision policy one should evaluate a tradeoff between generality and numerical viability of the ensuing optimization problems. In some recent papers, see, e.g., [6, 15, 12, 13] it has been observed that linear or affine policies do provide an interesting tradeoff by allowing reactive policies to be efficiently computed via convex optimization techniques. In this paper, we follow this route, and consider decisions prescribed by affine policies of the following form

$$u(k) = \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k)), \quad k = 1, \dots, T-1 \quad (11)$$

and $u(0) = \bar{u}(0)$, where $\bar{u}(k) \in \mathbb{R}^n$, $k \in \mathcal{K}$ are “nominal” allocation decision variables, $g(k)$ is the vector of gains over the k -th period, $\bar{g}(k)$ is a given estimate of the expected value of $g(k)$, and $\Theta(k) \in \mathbb{R}^{n,n}$, $k = 1, \dots, T-1$, are the policy “reaction matrices.” Notice that the nominal decisions $\bar{u}(k)$ are the decisions that would be executed if the realized gains $g(k)$ coincide with their expectations $\bar{g}(k)$; the role of the reaction matrices $\Theta(k)$ is to adjust the nominal allocation with a term proportional to the deviation of the gain $g(k)$ from its expected value (we fix henceforth $\Theta(0) \doteq 0$). Since the budget conservation constraint $\mathbf{1}^\top u(k) = 0$ must hold for any realization of the gains, we shall impose the restrictions

$$\mathbf{1}^\top \bar{u}(k) = 0, \quad \mathbf{1}^\top \Theta(k) = 0, \quad k \in \mathcal{K}.$$

4.1 Portfolio dynamics under affine policies

Applying the adjustment policy (11) to the portfolio dynamics equations (2), (3), we have

$$x^+(k) = x(k) + \bar{u}(k) + \Theta(k)(g(k) - \bar{g}(k)) \quad (12)$$

$$x(k+1) = G(k+1)x^+(k), \quad k \in \mathcal{K}, \quad (13)$$

with $\Theta(0) \doteq 0$. From repeated application of (12), (13) we obtain the expression for the portfolio composition at a generic instant $k = 1, \dots, T$:

$$x(k) = \Phi(1, k)x(0) + \Omega_k \bar{\mathbf{u}} + \sum_{t=1}^k \Phi(t, k)\Theta(t-1)\tilde{g}(t-1), \quad (14)$$

where $\bar{\mathbf{u}}^\top \doteq [\bar{u}(0)^\top \ \dots \ \bar{u}(T-2)^\top \ \bar{u}(T-1)^\top]$, and $\tilde{g}(k) \doteq g(k) - \bar{g}(k)$, for $k = 1, \dots, T$. A key observation is that $x(k)$ is an affine function of the decision variables $\bar{u}(k)$ and $\Theta(k)$, $k \in \mathcal{K}$. The cumulative gross return of the investment over the whole horizon is then

$$\begin{aligned} \varrho(\bar{\mathbf{u}}, \Theta) &= \frac{w(T)}{w(0)} = \frac{\mathbf{1}^\top x(T)}{\mathbf{1}^\top x(0)} \\ &= \frac{\phi(1, T)^\top x(0) + \omega_T^\top \bar{\mathbf{u}} + \sum_{t=1}^T \Phi(t, T)\Theta(t-1)\tilde{g}(t-1)}{\mathbf{1}^\top x(0)}, \end{aligned} \quad (15)$$

which is again affine in the variables $\bar{\mathbf{u}}$ and $\Theta \doteq [\Theta(1) \ \dots \ \Theta(T-1)]$.

4.2 Optimal closed-loop allocation with affine policies

Given N iid samples (scenarios) of the period gains $\{G(k), k = 1, \dots, T\}$, generated by a scenario generating oracle, we can determine optimal policies that minimize the empirical LPM₁ or LPM₂ cost by solving a problem similar to (8):

$$\begin{aligned} \text{LPM}_\nu^{\text{cl}}(\gamma) &= \min_{\bar{\mathbf{u}}, \Theta} \frac{1}{N} \sum_{i=1}^N (\max(0, \gamma - \varrho^{(i)}(\bar{\mathbf{u}}, \Theta)))^\nu \\ \text{s.t.} \quad &x^{(i)+}(k) \in \mathcal{X}(k), \quad k \in \mathcal{K}; \quad i \in \mathcal{N} \\ &\mathbf{1}^\top \bar{u}(k) = 0, \quad \mathbf{1}^\top \Theta(k) = 0, \quad k \in \mathcal{K}. \end{aligned} \quad (16)$$

where either $\nu = 1$ (for the LPM₁ cost) or $\nu = 2$ (for the LPM₂ cost), and where $x^{(i)+}(k)$ is given by (12), with $x(k)$ as in (14), under the i -th sampled scenario. These optimal allocations may be determined in a numerically efficient way by solving, respectively, a linear programming or a convex quadratic programming problem, as detailed in the following propositions.

Proposition 3 *For $\nu = 1$ (LPM₁ cost), problem (16) is equivalent to the following linear programming problem in the variables $\bar{\mathbf{u}} \in \mathbb{R}^{Tn}$, $\Theta \in \mathbb{R}^{n, (T-1)n}$, and $z \in \mathbb{R}^N$*

$$\begin{aligned} \text{LPM}_1^{\text{cl}}(\gamma) &= \min_{\bar{\mathbf{u}}, \Theta, z} \frac{1}{N} \sum_{i=1}^N z_i \\ \text{s.t.} \quad &x^{(i)}(k) + \bar{u}(k) + \Theta(k)\tilde{g}^{(i)}(k) \in \mathcal{X}(k), \quad k \in \mathcal{K}; \quad i \in \mathcal{N} \\ &\mathbf{1}^\top \bar{u}(k) = 0, \quad \mathbf{1}^\top \Theta(k) = 0, \quad k \in \mathcal{K}, \\ &z_i \geq 0, \quad i \in \mathcal{N}, \\ &z_i \geq \gamma - \varrho^{(i)}(\bar{\mathbf{u}}, \Theta), \quad i \in \mathcal{N}, \end{aligned} \quad (17)$$

where $x^{(i)}(k)$ is given by (14) on the i -th scenario, and $\varrho^{(i)}(\bar{\mathbf{u}}, \Theta)$ is given by (15).

Proposition 4 For $\nu = 2$ (LPM_2 cost), problem (16) is equivalent to the following convex quadratic programming problem in the variables $\bar{\mathbf{u}} \in \mathbb{R}^{Tn}$, $\Theta \in \mathbb{R}^{n, (T-1)n}$, and $z \in \mathbb{R}^N$

$$\begin{aligned}
LPM_2^{\text{cl}}(\gamma) &= \min_{\bar{\mathbf{u}}, \Theta, z} \frac{1}{N} \sum_{i=1}^N z_i^2 & (18) \\
\text{s.t.} &: x^{(i)}(k) + \bar{u}(k) + \Theta(k)\tilde{g}^{(i)}(k) \in \mathcal{X}(k), & k \in \mathcal{K}; i \in \mathcal{N} \\
& \mathbf{1}^\top \bar{u}(k) = 0, \quad \mathbf{1}^\top \Theta(k) = 0, & k \in \mathcal{K}, \\
& z_i \geq 0, & i \in \mathcal{N}, \\
& z_i \geq \gamma - \varrho^{(i)}(\bar{\mathbf{u}}, \Theta), & i \in \mathcal{N},
\end{aligned}$$

The statements of the previous two propositions follow easily by adding slack variables and applying an epigraphic transformation to the objective in (16), see, e.g., Section 8.3.4.4 in [11].

Remark 1 (γ - LPM_ν tradeoff curve) Problem (8) (resp. (16)) may be solved repeatedly for increasing values of γ , in order to trace a tradeoff frontier of the optimal cost value $LPM_\nu^{\text{ol}}(\gamma)$ (resp. $LPM_\nu^{\text{cl}}(\gamma)$) as a function of γ . Values of γ can be chosen in an interval $[\gamma_{\min}, \gamma_{\max}]$, where $\gamma_{\min} = \min_{i \in \mathcal{N}} \min_{j=1, \dots, n} \phi_j^{(i)}(1, T)$, and $\gamma_{\max} = \max_{i \in \mathcal{N}} \max_{j=1, \dots, n} \phi_j^{(i)}(1, T)$.

4.3 Receding/shrinking horizon implementation

In practical application of the method, either (8) or (16) are solved at time $k = 0$ mainly to obtain a good here-and-now decision $u(0)$. The future decisions or policies are rarely implemented in practice. Rather, at time $k = 1$ the decision maker collects the information coming from the realization of $G(1)$, which can be used to update the model used in the scenario generation oracle (and possibly the target return and constraints), and solves the whole problem again over a forward-shifted interval, to obtain $u(1)$. The same process is iterated for all subsequent periods. At each decision time k , the optimization interval can either be held fixed (receding horizon), if the investment is to be iterated indefinitely in time, or shrunk in duration by one period (shrinking horizon), if the final investment time T is fixed.

5 Numerical experiments

We next present a numerical test based on real financial data. We considered an allocation problem involving $n = 10$ asset types over $T = 12$ periods, each period having the duration of one month. Monthly historical return data have been kindly provided to us by Fondaco SGR, covering the period from 12/2002 to 12/2012, for the following ten assets (all data are converted to Euro currency):

1. MSDEWIN - MSCI Daily Equity Total Return (TR) World Index
2. NDUUEGF - MSCI Daily Equity Total Return Emerging Markets Index
3. SBEGEU - Citigroup EMU Europe Government Bond Index
4. SBWGEU - Citigroup WGBI World Government Bond Index
5. BCIW1E - Barclays World Inflation Linked Bonds TR Hedged
6. JGENVUUG - JPMorgan GBI-EM Emerging Government Bond Index
7. LGDRTRUH - Barclays Global Aggregate Credit Index
8. LG30TRUH - Barclays Global High Yield Total Return Index
9. DJUBSTR - Dow Jones UBS Commodity Index
10. SBWMEU3L - Citigroup 3 Month EUR Deposit Index.

Table 1 reports the historical monthly mean return and standard deviation for the considered assets.

asset n.	1	2	3	4	5	6	7	8	9	10
% ret.	0.096	0.242	0.089	0.069	0.108	0.176	0.058	0.160	0.023	0.042
% st. dev.	2.426	3.077	0.581	0.975	0.729	1.234	1.453	1.543	2.380	0.027

Table 1: Mean return and standard deviation of the considered assets, in the period from 12/2002 to 12/2012 (monthly data).

5.1 In-sample results

In this first test, we fixed the date of Jan. 2011 as the initial time of our investment horizon. The scenario generating oracle was in these experiments simply set up as a bootstrap resampler from the historical data preceding the selected initial date (sampling with replacement from the historical returns, see, e.g., [16]), and we used $N = 100$ scenarios per period (each period, of duration one month, is represented via $N = 100$ scenarios, for a total of $NT = 1200$ scenarios per simulation). The constraints were set to $\mathcal{X}(k) = \{x : x \geq 0\}$, for all k , and the target return was set to $\gamma = 1.08$ (i.e., a target 8% final return at the end of a 12 month investment period).

5.1.1 LPM₁ criterion

Figure 3 and Figure 4 show the result of open-loop and closed-loop (with affine policies) simulations, for the LPM₁ criterion.

As expected, while average results from open-loop and closed-loop approaches are similar (right panel in Figure 4), the closed-loop approach manages to better reduce the down-sided deviation from the target gain, as shown in the histogram in Figure 3; this effect is also visible in the left panel of Figure 4, in which open-loop wealth trajectories show larger dispersion with respect to closed-loop ones. Indeed, in this simulation we obtained $\text{LPM}_1^{\text{ol}}(\gamma) = 0.0656$, and $\text{LPM}_1^{\text{cl}}(\gamma) = 0.0431$, that is a 34% reduction on the objective criterion.




Figure 3: Histogram of final gains on scenarios (LPM₁ criterion). Light (green) bars show open-loop results, darker (blue) bars show closed-loop results.




Figure 4: Left: wealth profiles on the random scenarios (LPM₁ criterion). Right: average wealth profile over the scenarios. Light (green) lines show open-loop results, darker (blue) lines show closed-loop results.

Figure 5 shows the evolution of the optimal portfolio composition over time, for the open-loop approach.

Figure 6 shows the evolution of the *nominal* portfolio composition over time, for the closed-loop approach. The actual composition in the closed-loop approach depends on the realization of the returns during the execution period; the figure shows only the nominal composition, that is the composition obtained applying only the adjustments $\bar{u}(k)$ in eq. (12), neglecting the reactive term $\Theta(k)(g(k) - \bar{g}(k))$.




Figure 5: Composition of the optimal portfolios over time (LPM_1 criterion with open-loop approach). The bottom area of the figure represents the share of the first asset (MSDEWIN) in the portfolio, the top area represents the share of the tenth asset (SBWMEU3L), the other components are represented by the areas in between the bottom and top areas.




Figure 6: Composition of the nominal portfolios over time (LPM_1 criterion with closed-loop approach); shares are ordered as in Figure 5.

For illustration purposes, we also report below the first optimal reaction matrix $\Theta(1)$:

$$\Theta(1) = \begin{bmatrix} -0.0757 & -0.0941 & 0.0122 & 0.1020 & 0.2689 & 0.0402 & -0.0837 & -0.0846 & -0.1269 & 0.0417 \\ -0.0941 & 0.1875 & -0.0885 & -0.0075 & -0.1139 & 0.0367 & 0.1153 & 0.0352 & -0.1391 & 0.0685 \\ 0.0122 & -0.0885 & 0.5357 & -0.0175 & -0.0023 & -0.0924 & -0.1132 & -0.2133 & -0.0917 & 0.0710 \\ 0.1020 & -0.0075 & -0.0175 & -0.0979 & 0.0403 & -0.0136 & -0.0166 & 0.0239 & 0.0396 & -0.0528 \\ 0.2689 & -0.1139 & -0.0023 & 0.0403 & -0.0993 & 0.0839 & 0.0233 & 0.1158 & -0.3423 & 0.0257 \\ 0.0402 & 0.0367 & -0.0924 & -0.0136 & 0.0839 & -0.0018 & 0.0748 & -0.1694 & 0.1197 & -0.0780 \\ -0.0837 & 0.1153 & -0.1132 & -0.0166 & 0.0233 & 0.0748 & -0.3259 & -0.0217 & 0.2851 & 0.0627 \\ -0.0846 & 0.0352 & -0.2133 & 0.0239 & 0.1158 & -0.1694 & -0.0217 & 0.4050 & 0.0561 & -0.1469 \\ -0.1269 & -0.1391 & -0.0917 & 0.0396 & -0.3423 & 0.1197 & 0.2851 & 0.0561 & 0.1911 & 0.0083 \\ 0.0417 & 0.0685 & 0.0710 & -0.0528 & 0.0257 & -0.0780 & 0.0627 & -0.1469 & 0.0083 & -0.0002 \end{bmatrix}.$$

5.1.2 LPM₂ criterion

Similar results follow by using the LPM₂ objective, as shown in Figure 7 and Figure 8. Comments similar to the ones exposed for the LPM₁ case apply. In this case, we obtained $\text{LPM}_2^{\text{ol}}(\gamma) = 0.0070$, and $\text{LPM}_2^{\text{cl}}(\gamma) = 0.0034$, that is a 52% reduction on the objective criterion.




Figure 7: Histogram of final gains on scenarios (LPM₂ criterion). Light (green) bars show open-loop results, darker (blue) bars show closed-loop results.




Figure 8: Left: wealth profiles on the random scenarios (LPM₂ criterion). Right: average wealth profile over the scenarios. Light (green) lines show open-loop results, darker (blue) lines show closed-loop results.

5.2 Multi-period efficient frontier

The numerical efficiency of the proposed method also permits to easily obtain discretized plots of the multi-period efficient frontier, representing the optimal tradeoff curve of the minimal LPM risk level obtainable for a given value of γ , as discussed in Remark 1. For example, discretizing 11 values of γ in the range $[1, 1.2]$, we obtained the plot shown in Figure 9, for the LPM₁ objective under the closed-loop policy. The whole frontier was computed in about 3.75 minutes under Matlab on a standard Xeon workstation. Similarly, Figure 10 shows the result for the LPM₂ criterion.




Figure 9: LPM₁ 12-period risk/return frontier.




Figure 10: LPM₂ 12-period risk/return frontier.

5.3 Out-of-sample shrinking-horizon results

In this section, we test the proposed methodology in a more realistic simulation setting. We consider a horizon of $T = 12$ periods (one month each); starting at time $k = 0$ (i.e., the beginning of the first period), we solve the optimization problem over the 12 periods ahead, we determine the first optimal portfolio allocation $u(0)$, and we implement it. Then, we observe the actual (real) market outcome over the first period, and we mark our performance against this outcome, as well as against a number N_{out} of test outcomes randomly generated by the SGO. Then, at $k = 1$, we solve the optimization again over a reduced horizon of 11 periods, obtaining and implementing the here-and now decision $u(1)$, and marking our performance against both the real and the simulated outcomes of the market over the second period. At $k = 2$ we repeat the same process, over a forward optimization horizon of 10 periods, and so on, until at $k = 11$ we solve a problem over a single period for determining the last portfolio allocation $u(11)$. Therefore, in this simulation, at each $k = 0, 1, \dots, 11$, we solve an optimization problem P_k having a horizon of $T - k$ periods.

In the described setup, we set the target end-of-horizon gain to $\gamma_0 = 1.1$ (i.e., a 10% yearly return), for the first optimization problem P_0 over the 12 periods horizon. Then, for the subsequent problems with progressively shrinking horizon, we set the target end-of-horizon gain to $\gamma_k = \gamma_0^{(T-k)/T}$, $k = 1, \dots, T - 1$. When solving problem P_k , we use historical returns on a look-back period of 250 weeks preceding time k in the SGO. Thus, at each k , the SGO produces N_{in} scenarios of the forward (monthly) gains, by bootstrapping the weekly returns in the look-back window, and composing them monthly (one month is set equal to four weeks). These scenarios are used in the optimization, and the

ensuing allocation decision $u(k)$ is then evaluated on N_{out} new out-of-samples scenarios representing possible gains in the subsequent period (the decision is also evaluated on the actual market outcome in the subsequent period). In the simulations, we set $N_{\text{in}} = 300$, $N_{\text{out}} = 200$.

We performed the described experiment on three simulation periods, each starting at the beginning of January, for the years 2009, 2010, 2011. At all periods, we impose the constraints $x(k) \geq 0$ on the portfolios. For all experiments, we used the LPM_1 cost criterion. For each simulation period, we compared the out-of-sample performance of (a) the optimal open-loop strategy, (b) the optimal closed-loop strategy with affine policies, and (c) the naive $1/n$ fixed portfolio strategy (i.e., one in which an equally weighted portfolio is held fixed through the investment horizon).

The results for year 2009 are reported in Figure 11. Each panel in this figure




Figure 11: Year 2009 simulation: histograms of out-of-samples final gains for the open-loop (top), the closed-loop (middle), and the $1/n$ strategy (bottom). Red vertical line: desired target gain; green vertical line: achieved gain.

shows the out-of-sample histogram of the final gain achieved by each strategy: the red vertical line shows the given target gain $\gamma_0 = 1.1$, the green vertical line shows what the actual gain would have been on the actual realization of the market, the other blue bars show the distribution of the gain for $N_{\text{out}} = 200$ out-of-sample simulated possible behaviors of the market. The same type of

data is displayed in Figure 12 for the year 2010, and in Figure 13 for the year 2011. A summary of some resulting numerical indicators is reported in Table 2.

Year	LPM ₁ OL	LPM ₁ CL	LPM ₁ 1/ <i>n</i>	avg. OL	avg. CL	avg. 1/ <i>n</i>
2009	0.4166	0.4162	0.8248	1.1074	1.1065	1.0928
2010	0.2324	0.2316	0.4536	1.1496	1.1489	1.1608
2011	0.2575	0.2718	0.4883	1.1552	1.1456	1.1576

Table 2: Out-of-sample results for shrinking-horizon simulations with LPM₁ cost, $\gamma = 1.1$ target gain, and $T = 12$ periods.

The results show that, consistently over the simulation years, both the open-loop (OL) and closed-loop (CL) strategies provide average gains that are similar to those obtained via the 1/*n* strategy. However, the downside LPM₁ risk with respect to the target gain $\gamma = 1.1$ is about one half of that provided by the equally-weighted strategy. This indeed suggests that both the OL and CL strategies are superior to the 1/*n* strategy in controlling the downside deviation from the target. The OL and CL strategies yielded similar performances on these simulations: in the considered setting there appears to be little advantage to be gained by considering affine reactive policies instead of static OL policies. This fact, however, may be due to the nature of the SGO implemented in this example. Indeed, the focus of this work is on the optimization method, and not on the design of the SGO; therefore, we implemented a simple SGO based on bootstrap resample in the look-back period. This approach neglects possible inter-temporal correlations in the returns, and may thus make the advantage of a reactive policy less evident in the results. The development of a more sophisticated SGO for multi-period allocation problems (e.g., based on copula estimators and resampling) deserves a study on its own right, and should be the subject of future investigation.




Figure 12: Year 2010 simulation: histograms of out-of-samples final gains for the open-loop (top), the closed-loop (middle), and the $1/n$ strategy (bottom). Red vertical line: desired target gain; green vertical line: achieved gain.




Figure 13: Year 2011 simulation: histograms of out-of-samples final gains for the open-loop (top), the closed-loop (middle), and the $1/n$ strategy (bottom). Red vertical line: desired target gain; green vertical line: achieved gain.

6 Conclusions

In this paper we proposed a computationally efficient suboptimal approach for solving multi-period asset allocation problems under first and second-order lower-partial moment criteria. The key ingredients of the method are: (i) a scenario generating oracle (SGO), providing N scenarios, where each scenario is a path $\{r(1), \dots, r(T)\}$ representing a realization of the random return process over the whole planning horizon of T periods; (ii) an affine recourse policy of the form (11) for the forward decisions; and (iii) an ensuing formulation of the decision problem in the form of an efficiently solvable linear program (in the case of the LPM₁ objective), or convex quadratic problem (in the case of the LPM₂ objective). We observe that the design of the SGO is to be tailored to the needs of the decision maker, and it allows for inclusion of any type of inter-temporal stochastic dependence on the returns, as well as of expert knowledge in the form of specific paths of scenarios. This methodology has the potential of enabling solution of multi-period decision problems with inter-temporal return correlation and realistically-sized number of periods, where alternative stochastic programming approaches tend to be either too coarse, or prohibitively heavy from a computational point of view.

References

- [1] K.J. Åström and R.M. Murray. *Feedback Systems: An Introduction for Scientists and Engineers*. Princeton University Press, 2008.
- [2] V.S. Bawa. Optimal rules for ordering uncertain prospects. *Journal of Financial Economics*, 2:95–121, 1975.
- [3] V.S. Bawa, S.J. Brown, and R.W. Klein. *Estimation Risk and Optimal Portfolio Choice*. North-Holland, Amsterdam, 1979.
- [4] V.S. Bawa and E.B. Lindenberg. Capital market equilibrium in a mean-lower partial moment framework. *J. of Financial Economics*, 5(2):189–200, 1977.
- [5] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Math. Programming, Ser. A*, 99:351–376, 2004.
- [6] D. Bertsimas and A. Thiele. Robust and data-driven optimization: modern decision making under uncertainty. *INFORMS Tutorials on Operations Research*, pages 195–122, 2006.
- [7] J.R. Birge. Stochastic programming computation and applications. *INFORMS J. Comput.*, 9(2):111–133, 1997.
- [8] J.R. Birge and F. Louveaux. *Introduction to Stochastic Programming*. Springer Verlag, 1997.

- [9] S. Boyd, M. T. Mueller, B. O’Donoghue, and Y. Wang. Performance bounds and suboptimal policies for multi-period investment. *Foundations and Trends in Optimization*, 1(1):1–69, 2014.
- [10] M. W. Brandt, A. Goyal, P. Santa-Clara, and J. R. Stroud. A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies*, 18(3):831–873, 2005.
- [11] G. Calafiore and L. El Ghaoui. *Optimization Models*. Cambridge Univ. Press, 2014.
- [12] G.C. Calafiore. Multi-period portfolio optimization with linear control policies. *Automatica*, 44(10):2463–2473, 2008.
- [13] G.C. Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.*, 48(8):2254–2274, 2009.
- [14] V.K. Chopra and W.T. Ziemba. The effect of errors in means, variances and covariances on optimal portfolio choice. *Journal of Portfolio Management*, (Winter):6–11, 1993.
- [15] P.A. Parrilo D. Bertsimas, D.A. Iancu. Optimality of affine policies in multistage robust optimization. *Mathematics Of Operations Research*, 35:363–394, 2010.
- [16] B. Efron. Bootstrap methods: Another look at the jackknife. *The Annals of Statistics*, 7(1):1–26, 1979.
- [17] P.-C. Fishburn. Mean risk analysis with below target returns. *The American Economic Review*, 67(2):116–126, 1977.
- [18] M. Gilli, E. K ellezi, and H. Hysi. A data-driven optimization heuristic for downside risk minimization. *J. of Risk*, pages 1–19, 2006.
- [19] N. Gulpinar, B. Rustem, and R. Settergren. Multistage stochastic programming in computational finance. In *Decision Making, Economics and Finance: Optimization Models*, pages 33–47. Kluwer Academic Publishers, 2002.
- [20] C. Jabbour, J.F. Pena, J.C. Vera, and L.F. Zuluaga. An estimation-free, robust conditional value-at-risk portfolio allocation model. *J. of Risk*, 11(1):1–22, 2008.
- [21] H. Konno and H. Yamazaki. Mean absolute deviation portfolio optimization model and its application to Tokyo stock market. *Management Science*, 37:519–531, 1991.
- [22] H.M. Markowitz. Portfolio selection. *Journal of Finance*, 7:77–91, 1952.

- [23] David N. Nawrocki. Optimal algorithms and lower partial moment: ex post results. *Applied Economics*, 23:465–470, 1991.
- [24] A. Ruszczyński and A. Shapiro, editors. *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*. Elsevier, Amsterdam, 2003.
- [25] A. Shapiro. On complexity of multistage stochastic programs. *Operations Research Letters*, 34(1):1–8, 2006.
- [26] Roger J-B. Wets. Stochastic programming models: Wait-and-see versus here-and-now. In C. Greengard and A. Ruszczyński, editors, *Decision Making Under Uncertainty*, volume 128 of *The IMA Volumes in Mathematics and its Applications*, pages 1–15. Springer, 2002.