

On Optimal Portfolios of Dynamic Resource Allocations

Yingdong Lu, Siva Theja Maguluri, Mark S. Squillante
Chai Wah Wu

February 24, 2017

Abstract

We consider the optimal allocation of generic resources among multiple generic entities of interest over a finite planning horizon, where each entity generates stochastic returns as a function of its resource allocation during each period. The main objective is to maximize the expected return while at the same time managing risk to an acceptable level for each period. We devise a general solution framework and establish how to obtain the optimal dynamic resource allocation.

1 Introduction

The trade-offs among risks and returns are central and fundamental to the planning and management of any collection of generic entities of interest to an organization. At any given time, the organization may support various entities under its control with generic resources of interest. The risks and returns associated with each entity as a function of the amount of allocated resources differ from one entity to the next, as well as across entities. These risks and returns also evolve over time according to the dynamics of each entity, the dynamics across entities, and the dynamics of exogenous factors

Y. Lu, M.S. Squillante and C.W. Wu are with the Mathematical Sciences Department at the IBM Thomas J. Watson Research Center, {yingdong,mss,cwu}@us.ibm.com; S.T. Maguluri is with the H. Milton Stewart School of Industrial and Systems Engineering at the Georgia Institute of Technology, siva.theja@gatech.edu

such as environmental changes. Meanwhile, the organization seeks to optimize its own objectives within the context of various restrictions on what portfolio of resource allocation decisions and actions can be made and when. A commonly used objective is to maximize the expected returns of a collection of entities of interest to the organization over a given planning horizon, while simultaneously maintaining the risk exposure of its portfolio of resource allocation decisions and actions, usually represented by the variability of returns, within acceptable levels. To ensure all objectives are achieved and restrictions are met, the decisions on resource allocations to each entity are periodically reviewed and adjusted over time.

This general dynamic resource allocation problem prominently arises for different types of entities across a broad spectrum of application domains. Examples include applications in cybersecurity, pharmaceuticals, software products, and business management. For cybersecurity applications, the entities can represent different computing devices in distinct states of cybersecurity vulnerability; the resources can include various forms of financial investments, people working on system security, and various cybersecurity applications (e.g., patch); and the decisions can concern the amount of resources to be allocated to each device; an analogous set of entities, resources and decisions exists for other forms of computer system management which may also include reliability and performance. In the pharmaceutical industry, the entities can represent research projects in different stages of development and trials; the resources can include various forms of financial investments, scientists and researchers working on the projects, and various pharmaceutical compounds and materials; and the decisions can concern the amount of resources to be allocated to each project. Turning to the software industry for yet another example, the entities can represent product offerings and releases of versions of product offerings; the resources can include various forms of financial investments, software architects and programmers working on the products, and various equipment and infrastructure support; and the decisions can concern the amount of resources to be allocated to each product. Finally, as a general business management application, the entities can represent different areas of the business organization such as development, marketing and sales; the resources can include various forms of financial investments, people working in the areas, and various infrastructure support; and the decisions can concern the amount of resources to be allocated to each area. In each of these application domains, the organization may want to weight the benefits differently from one time period to the next such as

preferring long-term benefits even at the expense of short-term benefits or preferring short-term benefits even at the expense of long-term benefits.

Motivated by these real-world applications, we consider a mathematical abstraction of the dynamic resource allocation problem and devise a general solution framework. In particular, we mathematically model the resources and returns associated with every entity as quantities that can be measured in terms of certain forms of currency which facilitate the interactions among different entities, where this abstract notion of currency represents any combination of different types of resources such as financial investments, human resources, equipment, materials, and so on. This mathematical abstraction also includes a fundamental connection between entities and financial instruments, even though they are very different from traditional financial instruments, which leads to the representation of the returns from entities as random variables where such returns are functions of the amount of resources allocated to the entities and evolve over time according to the intrinsic dynamics associated with the resource allocations and entities. The allocation of resources are reviewed and adjusted over time to achieve maximum benefit as uncertainty is realized and information is updated. In addition, at each time period, constraints are imposed on the variance of the return random variables such that only allocations satisfying these restrictions are deemed feasible.

Our mathematical models and analysis have certain degrees of resemblance to multi-period mean-variance models in the mathematical finance literature; see, e.g., [4, 5, 8]. This includes an additional entity, with little or no risk, to serve as a reference point for the returns of the original set of entities, as is standard in the literature [4]. In our context, the additional (last) entity can be cash or a Treasury note, thus representing an option that we do not need to invest all of the resources all of the time. Alternatively, in other instances of our models, the additional (last) entity can be a benchmark entity that represents standards against which every entity will be evaluated. There are also important connections between the resource allocations in our models and the investments in classical portfolio optimization, though with differences in semantics, where there is an implicit assumption in our models that the returns from resource allocations at the end of one period are converted into the amount of resources available at the start of the next period.

At the same time, however, there are significant differences between our models and those in the mathematical finance literature. First, it is typ-

ical for the risk constraints to only be imposed on the terminal value in mathematical finance models. This differs from our allocation of resources to entities together with their associated financial assets, where such operational aspects of the problem have a much larger impact on the planning and management of entities and their returns, which in turn require that the associated risk factors be monitored and controlled at a much higher frequency. Second, there appears to be an issue with a couple of key papers in the field [4, 5] where, under the same model that we consider, the amount of investment funds available from one period to the next are not properly restricted to the amount of wealth obtained up to that point in time; i.e., an unlimited amount of “borrowed” funds are allowed at no cost. This differs from our allocation of resources to entities, where allocation beyond current capacity can be much more difficult to do and therefore we explicitly enforce the constraint of the problem formulation on the investment funds available in each period. In addition, our models and analysis support the general case where borrowing is allowed at a given cost rate. Our models and analysis further support the inclusion of weights to differentiate the contributions from one time period to the next.

We note that the research literature includes other mathematical models and methods to address the problem of dynamic resource allocation. In particular, there have been several attempts to directly apply portfolio management models and methods in the planning and management of scientific and development projects. One such approach [1, 3] is based on mathematical programming, including stochastic programming. While some of these models and methods are quite powerful at incorporating constraints, it is not surprising that one of the key features often ignored in these previous studies concerns the correlations among entities, which usually causes the mathematical programming approach to become extremely complex and renders its computational complexity prohibitive [3]. In contrast, as we will show, the mathematical models and methods devised in this paper can readily support many aspects of the various sources of uncertainty, including correlations among entities and their evolution over time. Another approach to addressing the problem of project management is based on decision science methods, mostly involving various structured decision trees. This approach has important limitations when the parameter and state spaces that have to be discretized become very large. Our interests in this paper concern large-scale optimal dynamic allocation of resources among multiple entities of interest over a long-run planning horizon.

The remainder of the paper is organized as follows. Our mathematical models are presented in Section 2, followed by our mathematical analysis in Section 3. Extensions of our general solution framework are discussed in Section 4, followed by concluding remarks.

2 Mathematical Models

Consider the allocation of a set of common resources among $n + 1$ general entities over a discrete-time planning horizon comprising T periods. For each period $t = 1, \dots, T$, the amount of resources allocated to entity $i = 1, \dots, n + 1$ is denoted by u_i^t . Let x_t denote the total return at the end of period t where, with slight abuse of notation, such returns are implicitly converted into the amount of available resources at the start of period $t + 1$. The total amount of resources available at the beginning of the horizon is given by x_0 , and thus $\sum_{i=1}^{n+1} u_i^1 = x_0$. Within each period, random returns are generated proportional to the amount of resources allocated to each entity. More specifically, let us define

- $e_t :=$ return of the $(n + 1)$ -st entity during time period t per unit of resource allocated;
- $e_i^t :=$ return of the i -th entity during time period t per unit of resource allocated;

where these random variables (r.v.s) are independent over time. Under these assumptions, x_t satisfies the dynamics

$$x_t = e_t x_{t-1} + \sum_{i=1}^n (e_i^t - e_t) u_i^t, \quad t = 1, \dots, T, \quad (1)$$

where the $(n + 1)$ -st entity is used as a point of reference and $u_{n+1}^t = x_{t-1} - \sum_{i=1}^n u_i^t$ by definition.

Defining the n -dimensional vectors \mathbf{P}_t and \mathbf{u}_t to be

$$\begin{aligned} \mathbf{P}_t &:= [(e_1^t - e_t), (e_2^t - e_t), \dots, (e_n^t - e_t)]', \\ \mathbf{u}_t &:= (u_1^t, \dots, u_n^t)', \end{aligned}$$

respectively, the system dynamics can be expressed as

$$x_t = e_t x_{t-1} + \mathbf{P}_t' \mathbf{u}_t. \quad (2)$$

We next define the $(n + 1)$ -dimensional vector $\mathbf{e}^t := (e_t, e_1^t, \dots, e_n^t)'$ and introduce the following assumption.

Assumption 1. $\mathbb{E}[\mathbf{e}^t(\mathbf{e}^t)']$ is positive definite for all time periods $t = 1, \dots, T$.

Remark 1. While $\mathbb{E}[\mathbf{e}^t(\mathbf{e}^t)']$ are positive semidefinite matrices by definition, Assumption 1 guarantees they are not degenerate (i.e., $\mathbf{e}^t = 0$, $\forall t$). Further note that our model does not impose strong distributional assumptions on the returns; we only require that they have finite second moments.

The $(n + 1) \times (n + 1)$ matrix $\mathbb{E}[\mathbf{e}^t(\mathbf{e}^t)']$ then can be expressed as

$$\begin{bmatrix} \mathbb{E}[e_t^2] & \mathbb{E}[e_t e_1^t] & \dots & \mathbb{E}[e_t e_n^t] \\ \mathbb{E}[e_1^t e_t] & \mathbb{E}[(e_1^t)^2] & \dots & \mathbb{E}[e_1^t e_n^t] \\ \vdots & \vdots & \ddots & \vdots \\ \mathbb{E}[e_n^t e_t] & \mathbb{E}[e_n^t e_1^t] & \dots & \mathbb{E}[(e_n^t)^2] \end{bmatrix}.$$

Assumption 1 implies that

$$\begin{bmatrix} \mathbb{E}[(e_t)^2] & \mathbb{E}[e_t \mathbf{P}_t'] \\ \mathbb{E}[e_t \mathbf{P}_t] & \mathbb{E}[\mathbf{P}_t \mathbf{P}_t'] \end{bmatrix} = \begin{bmatrix} 1 & 0 & \dots & 0 \\ -1 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ -1 & 0 & \dots & 1 \end{bmatrix} \mathbb{E}[\mathbf{e}^t(\mathbf{e}^t)'] \begin{bmatrix} 1 & -1 & \dots & -1 \\ 0 & 1 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & 1 \end{bmatrix}$$

is also positive definite. In addition, we have

$$\mathbb{E}[\mathbf{P}_t \mathbf{P}_t'] > 0, \quad \forall t, \quad (3)$$

$$\mathbb{E}[(e_t)^2] - \mathbb{E}[e_t \mathbf{P}_t'] \mathbb{E}^{-1}[\mathbf{P}_t \mathbf{P}_t'] \mathbb{E}[e_t \mathbf{P}_t] > 0, \quad \forall t. \quad (4)$$

Our resource allocation goal is to maximize the expected return over time, while maintaining the variance of return within an acceptable range. This is a natural model for many different problem instances across a wide variety of applications. The progress of resource allocations are often closely monitored and periodically reviewed so that they can be adjusted to yield the best return. Hence, the returns and risks of each period need to be factored into the decision making process as forms of operational measurements.

Mathematically, our resource allocation planning can be formulated as the following stochastic optimization problem

$$\max_{(\mathbf{u}_1, \dots, \mathbf{u}_T)} \sum_{t=1}^T w_t \mathbb{E}[x_t] \quad (5)$$

$$\text{s.t.} \quad \text{Var}[x_t] \leq \alpha_t, \quad t = 1, \dots, T, \quad (6)$$

where the weights w_t allow differentiation among the contributions from distinct periods, and the pre-specified risk tolerances α_t reflect levels of risk deemed to be acceptable. Let us denote the set of optimal allocations as $\Pi_1(\mathbf{w}, \alpha)$.

This type of portfolio optimization problem is often investigated through related Lagrangian relaxations. More specifically, solving this problem is equivalent to solving

$$\max_{(\mathbf{u}_1, \dots, \mathbf{u}_T)} \sum_{t=1}^T w_t \mathbb{E}[x_t] - y_t \text{Var}[x_t] \quad (7)$$

for some vectors \mathbf{y} . Denote the set of optimal allocations by $\Pi_2(\mathbf{w}, \mathbf{y})$. This problem remains challenging to solve because of the variance term. Since variance is not linear, and since x_t evolves according to (1), the variance term here couples the decisions across the time slots. However, as in [4, 5], the problem can be solved by considering a simpler set of problems that also renders the optimal solution of (7). To this end, we consider the general problem

$$\max_{(\mathbf{u}_1, \dots, \mathbf{u}_T)} \sum_{t=1}^T \mathbb{E}[a_t x_t - b_t x_t^2]. \quad (8)$$

Let us denote the set of optimal allocations as $\Pi_3(\mathbf{a}, \mathbf{b})$. We henceforth focus on the solution of (8) with respect to a_t and b_t as a result of the following proposition.

Proposition 1. *Any solution to the problem (7) is also a solution to the problem (8) for some vectors \mathbf{a} and \mathbf{b} , i.e., $\Pi_2(\mathbf{w}, \mathbf{y}) \subseteq \Pi_3(\mathbf{a}, \mathbf{b})$ for some vectors \mathbf{a} and \mathbf{b} .*

Proof. The proof of this proposition is by contradiction and similar to that of Theorem 1 in [5], and thus we only provide a brief summary here. Let $U(\cdot)$

denote the objective in (7). Suppose there is a sequence of optimal decisions for (7), $\pi^* = (\mathbf{u}_1^*, \dots, \mathbf{u}_T^*) \in \Pi_2(\mathbf{w}, \mathbf{y})$, and let \mathbf{x}^* denote the corresponding returns. Set $b_t = y_t$ and $a_t = w_t + 2y_t E[x_t^*]$ and further suppose that $(\mathbf{u}_1^*, \dots, \mathbf{u}_T^*) \notin \Pi_3(\mathbf{a}, \mathbf{b})$. Note that we selected $a_t = \frac{\partial U}{\partial E[x_t]}$ under the optimal decisions π^* . There exists a $\tilde{\pi} = (\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_T)$ that is an optimal solution of (8). From this fact together with the fact that $U(\cdot)$ is convex in $E[x_t]$ and $E[x_t^2]$ for all t , one can conclude that using the decisions $\tilde{\pi}$ leads to a better value of $U(\cdot)$ than the decisions π^* , which contradicts the optimality of π^* . \square

3 Mathematical Analysis

We now present our analysis of the above stochastic optimization problem. First, to illustrate the basic approach, we present a detailed analysis of the one-dimensional case where many quantities can be explicitly calculated. Then, we proceed to provide our solution for the general problem, followed by extensions of our analysis to support acquisition of additional resources in each period at a cost.

3.1 One-dimensional Case

Consider the case of $n = 1$. From (2), the total return for any period $t = 1, \dots, T$ is given by

$$R_t = e_t x_{t-1} + P_t u_t = e_1^t u_t + e_t (x_{t-1} - u_t),$$

with u_t denoting the amount of resource allocation. The r.v. e_1^t has first two moments $p_{1,t}$ and $p_{2,t}$, whereas the return of the second entity e_t is assumed to be riskless with mean $r_0 > 0$ and variance 0. This then renders the stochastic optimization problem

$$\max_{(u_1, \dots, u_T)} \sum_{t=1}^T [a_t \mu_t - b_t \nu_t],$$

where $a_t, b_t \geq 0$ are given parameters, $\mu_t = \mathbb{E}[R_t]$, and $\nu_t = \mathbb{E}[R_t^2]$.

Let us next consider the solution of this stochastic dynamic program. At the beginning of period T , since the total amount of available resources is x_{T-1} , we only need to determine the resource allocations u_T for time period

T that maximize $J_T = a_T\mu_T - b_T\nu_T$. If the total allocation is y , the objective can be rewritten as

$$J_T = a_T[yp_{1,T} + r_0(x_{T-1} - y)] - b_T[y^2p_{2,T} + r_0^2(x_{T-1} - y)^2 + 2r_0p_{1,T}y(x_{T-1} - y)].$$

Upon taking the derivative with respect to y , we obtain

$$\begin{aligned} \frac{\partial J_T}{\partial y} &= [a_T(p_{1,T} - r_0) + 2b_Tx_{T-1}(r_0^2 - r_0p_{1,T})] \\ &\quad - 2b_T(p_{2,T} + r_0^2 - 2r_0p_{1,T})y. \end{aligned}$$

Hence, the stationary point is given by

$$y^* = \frac{a_T(p_{1,T} - r_0) + 2b_Tx_{T-1}(r_0^2 - r_0p_{1,T})}{2b_T(p_{2,T} + r_0^2 - 2r_0p_{1,T})}.$$

In contrast to [4, 5], we enforce the available resource constraint x_{T-1} on allocation decisions. More specifically, on one side when $y^* > x_{T-1}$, we have

$$\frac{a_T(p_{1,T} - r_0) + 2b_Tx_{T-1}(r_0^2 - r_0p_{1,T})}{2b_T(p_{2,T} + r_0^2 - 2r_0p_{1,T})} > x_{T-1}$$

which is equivalent to

$$x_{T-1} < \frac{a_T(p_{1,T} - r_0)}{2b_T(p_{2,T} - r_0p_{1,T})}.$$

Then the stationary point is out of reach and the maximum is obtained at x_{T-1} . The optimal policy therefore allocates all resources to the first entity, thus rendering

$$J_T(x_{T-1}) = a_Tx_{T-1}p_{1,T} - b_Tx_{T-1}^2p_{2,T}.$$

Observe that $J_T(x_{T-1})$ has a decreasing derivative as a function of x_{T-1} and that the right derivative of $J_T(x_{T-1})$ evaluated at

$$x_{T-1} = \frac{a_T(p_{1,T} - r_0)}{2b_T(p_{2,T} - r_0p_{1,T})}$$

is equal to, after some simplification,

$$\frac{p_{2,T} - p_{1,T}^2}{p_{2,T} - r_0p_{1,T}}a_Tr_0.$$

On the other hand, if

$$x_{T-1} \geq \frac{a_T(p_{1,T} - r_0)}{2b_T p_{2,T}},$$

then the optimal policy allocates y^* resources to the first entity, in which case the value function bears the form

$$\begin{aligned} J_T(x_{T-1}) = & a_T[y^* p_{1,T} + r_0(x_{T-1} - y^*)] - b_T[(y^*)^2 p_{2,T} \\ & + r_0^2(x_{T-1} - y^*)^2 + 2r_0 p_{1,T} y^*(x_{T-1} - y^*)]. \end{aligned}$$

Straightforward calculations show that $J_T(x_{T-1})$ has a decreasing derivative with respect to x_{T-1} and the left derivative at

$$x_{T-1} = \frac{a_T(p_{1,T} - r_0)}{2b_T(p_{2,T} - r_0 p_{1,T})}$$

is also equal to

$$\frac{p_{2,T} - p_{1,T}^2}{p_{2,T} - r_0 p_{1,T}} a_T r_0.$$

As a result of these two cases, we can conclude that the value function $J_T(x_{T-1})$ is a concave function of x_{T-1} .

With x_{T-2} resources available at the beginning of period $T - 1$, if the allocation to the first entity is set at $y \leq x_{T-2}$, then the objective can be expressed as

$$\max J_{T-1} = a_{T-1} \mu_{T-1} - b_{T-1} \nu_{T-1} + \mathbb{E}[J_T(x_{T-1})].$$

Equivalently, we have

$$\begin{aligned} \max_y \quad & a_{T-1}[y p_{1,T-1} + r_0(x_{T-2} - y)] - b_{T-1}[y^2 p_{2,T-1} + r_0^2(x_{T-2} - y)^2 \\ & + 2r_0 p_{1,T-1}(x_{T-2} - y)] + \mathbb{E}[J_T(r_0(x_{T-2} - y) + e_1^{T-1} y)]. \end{aligned}$$

Meanwhile, from the above analysis, we know that

$$\begin{aligned} & \mathbb{E}[J_T(r_0(x_{T-2} - y) + e_1^{T-1} y)] \\ = & \mathbb{E}[\{a_T[(r_0(x_{T-2} - y) + e_1^{T-1} y) p_{1,T}] \\ & - b_T[(r_0(x_{T-2} - y) + e_1^{T-1} y)^2 p_{2,T}]\} \times \mathbb{I}\{G_{T-1}\}] \\ & + \mathbb{E}[\{a_T[(y^*) p_{1,T} + r_0((x_{T-2} - y) + e_1^{T-1} y - y^*)] \\ & - b_T[(y^*)^2 p_{2,T} + r_0^2(x_{T-2} - y + e_1^{T-1} y - y^*)^2]\} \times \mathbb{I}\{G_{T-1}^c\}]] \end{aligned}$$

with

$$G_{T-1} := \left\{ r_0(x_{T-2} - y) + e_1^{T-1}y \leq \frac{a_T(p_{1,T} - r_0)}{2b_T(p_{2,T} - r_0p_{1,T})} \right\},$$

where $\mathbb{I}\{Z\}$ is the indicator function for event Z .

The above quantity, of course, can be easily calculated for any distributions with finite first two moments $p_{1,t}$ and $p_{2,t}$ for the independent r.v.s e_1^t . Once again, in contrast to [4, 5], we explicitly consider the constraint of available resources x_{T-2} on allocation decisions. Hence, we can repeat the process for the last period by finding the stationary point, optimally allocating the amount of resources $(x_{T-2} \wedge y^*)$ to the first entity, and then exploiting the properties of the function to show that the value function remains concave.

In general, for any time period t , with initial availability of x_{t-1} resources, the objective becomes

$$\begin{aligned} \max_y \quad & a_t[yp_{1,t} + r_0(x_{t-1} - y)] - b_t[y^2p_{2,t} + r_0^2(x_{t-1} - y)^2 \\ & + 2r_0p_{1,t}y(x_{t-1} - y)] + \mathbb{E}[J_{t+1}(r_0(x_{t-1} - y) + e_1^ty)] \end{aligned}$$

where $\mathbb{E}[J_{t+1}(r_0(x_{t-1} - y) + e_1^ty)]$ can be readily calculated as demonstrated above. Then we can obtain the critical value and the optimal policy will again be to allocate the minimum between this threshold and x_{t-1} . To guarantee obtaining the critical value efficiently, which is key in the derivation for the one-dimensional case, we will establish the preservation of the concavity of the objective for any period. Our approach is based on inductive arguments, where the initial case has been discussed fully above. Now, assume the concavity of the objective function holds for any period $t + 1, \dots, T$.

At the beginning of period t , the amount of available resources is x_{t-1} such that $y \leq x_{t-1}$ resources can be allocated to the first entity which yields a total return of

$$e_1^ty + J_{t+1}(e_1^ty + r_0(x_{t-1} - y)).$$

Hence, the optimal decision for period t is determined through the following problem

$$\begin{aligned} \max_y \quad & a_tp_{1,t}y + a_tr_0(x_{t-1} - y) - b_tp_{2,t}y^2 - b_tr_0^2(x_{t-1} - y)^2 \\ & - 2b_tr_0p_{1,t}y(x_{t-1} - y) + \mathbb{E}[J_{t+1}(e_1^ty + r_0(x_{t-1} - y))]. \end{aligned} \quad (9)$$

Taking the derivative of the objective function with respect to y , we obtain

$$a_t(p_{1,t} - r_0) + 2b_tx_{t-1}(r_0^2 - r_0p_{1,t}) - 2b_t(p_{2,t} + r_0^2 - 2r_0p_{1,t})y + \frac{\partial}{\partial y}E[J_{t+1}(e_1^ty + r_0(x_{t-1} - y))]. \quad (10)$$

Let y^* denote the critical value for this function. Further, since y^* is a function of x_{t-1} , let us denote by x_{t-1}^* the solution to $y^*(x_{t-1}) = x_{t-1}$. Then, if $x_{t-1} < x_{t-1}^*$, the optimal decision for period t is to allocate all x_{t-1} available resources to the first entity; otherwise, the optimal decision is to allocate x_{t-1}^* resources to the first entity.

Hence, when $x_{t-1} < x_{t-1}^*$, the value function can be expressed as

$$a_tp_{1,t}x_{t-1} - b_tp_{2,t}x_{t-1}^2 + E[J_{t+1}(e_1^tx_{t-1})],$$

which is clearly a concave function of x_{t-1} . Moreover, the left derivative evaluated at $x_{t-1} = y^*$ is given by

$$a_tp_{1,t} - 2b_tp_{2,t}y^* + \frac{\partial}{\partial y}E[J_{t+1}(e_1^ty^*)]. \quad (11)$$

By definition of y^* , the above expression is equivalent to

$$1 - \frac{\partial}{\partial y}E[J_{t+1}(e_1^ty + r_0(x_{t-1} - y))] + \frac{\partial}{\partial y}E[J_{t+1}(e_1^ty^*)],$$

which tends to approach 1 as $x_{t-1} \rightarrow x_{t-1}^*$. In summary, the derivative remains positive and decreases when $x_{t-1} < x_{t-1}^*$.

On the other hand, when $x_{t-1} \geq x_{t-1}^*$, the value function and its derivative are given by (9) and (10), respectively. From the definition of x_{t-1}^* , we know that the derivative is less than or equal to (11) and non-increasing. Concavity is therefore preserved and the desired result follows by induction.

3.2 General Case

We proceed along similar lines in our development of an algorithmic framework to solve the general problem. Starting at period T , with initial available resources x_{T-1} , the objective that is to be maximized has the following form

$$J_T = E[-b_Tx_T^2 + a_Tx_T \mid x_{T-1}].$$

Substituting the expression (2) for x_T , we obtain

$$J_T = \{-b_T \mathbb{E}[e_T^2] x_{T-1}^2 + a_T \mathbb{E}[e_T] x_{T-1}\} + \{a_T \mathbb{E}[\mathbf{P}'_T] - 2b_T \mathbb{E}[e_T \mathbf{P}'_T] x_{T-1}\} \mathbf{u}_T - b_T \mathbf{u}'_T \mathbb{E}[\mathbf{P}_T \mathbf{P}'_T] \mathbf{u}_T,$$

and therefore

$$\frac{\partial J_T(\mathbf{u}_T \mid x_{T-1})}{\partial \mathbf{u}_T} = a_T \mathbb{E}[\mathbf{P}'_T] - 2b_T \mathbb{E}[e_T \mathbf{P}'_T] x_{T-1} - 2b_T \mathbb{E}[\mathbf{P}_T \mathbf{P}'_T] \mathbf{u}_T.$$

Upon examining the system of equations

$$\frac{\partial J_T(\mathbf{u}_T \mid x_T)}{\partial \mathbf{u}_T} = 0,$$

we derive the solution

$$\mathbf{u}_T^* = \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \{a_T / (2b_T) \mathbb{E}[\mathbf{P}_T] - \mathbb{E}[e_T \mathbf{P}'_T] x_{T-1}\}. \quad (12)$$

Now, if $\mathbf{1}' \mathbf{u}_T^* \leq x_{T-1}$ with $\mathbf{1}$ being the n -dimensional column vector of all ones, i.e., the initial amount of available resources is large enough to accommodate the allocation corresponding to the solution (12), then \mathbf{u}_T^* is in fact the optimal resource allocation decision. Otherwise, we need to introduce the constraint $\mathbf{1}' \mathbf{u}_T = x_{T-1}$, as well as a multiplier ν . Applying this to the optimization problem, we obtain

$$\mathbf{u}_T = \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - \nu, \quad \mathbf{1}' \mathbf{u}_T = x_{T-1},$$

with

$$c_T := \frac{a_T}{2b_T} \mathbb{E}[\mathbf{P}_T] - \mathbb{E}[e_T \mathbf{P}'_T] x_{T-1},$$

which renders

$$\nu = \frac{1}{n} \{ \mathbf{1}' \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - x_{T-1} \}$$

and therefore

$$\mathbf{u}_T^* = \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - \frac{1}{n} \{ \mathbf{1}' \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - x_{T-1} \}. \quad (13)$$

Upon combining the solutions in (12) and (13), we have

$$\mathbf{u}_T^* = \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - \frac{1}{n} \{ \mathbf{1}' \mathbf{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] c_T - x_{T-1} \}^+. \quad (14)$$

Substituting this into the objective function, we derive the optimal cost-to-go for a given x_{T-1} to be

$$\begin{aligned}
J_T^*(x_{T-1}) = & \\
& - b_T \left[\mathbb{E}[e_T^2] - 2 \left(1 - \frac{1}{n} A_T \right) \mathbb{E}[e_T \mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \times \mathbb{E}[e_T \mathbf{P}'_T] - A_T \right] x_{T-1}^2 \\
& + a_T \left[\mathbb{E}[e_T] - \left(1 - \frac{1}{n} A_T \right) \mathbb{E}[\mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \mathbb{E}[e_T \mathbf{P}'_T] \right] x_{T-1} \\
& + \left(1 - \frac{1}{n} A_T \right) \frac{a_T^2}{4b_T} \mathbb{E}[\mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \mathbb{E}[\mathbf{P}_T]
\end{aligned}$$

where

$$A_T = \mathbb{I}\{\mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T][c_T - \mathbf{1}x_{T-1}] \geq 0\}.$$

Observe that the above expression for $J_T^*(x_{T-1})$ is a joint combination of two quadratic functions of x_{T-1} . Furthermore, by the following known result in convex analysis (see, e.g., [6]), $J_T^*(x_{T-1})$ remains a concave function of x_{T-1} .

Theorem 1. *The solution of a convex programming problem with one knapsack constraint is convex with respect to the constrained value.*

Next, proceeding to period $T-1$ with an initial amount of x_{T-2} resources, the objective can be expressed as

$$\max_{\mathbf{u}_{T-1}} a_{T-1} \mathbb{E}[x_{T-1}] - b_{T-1} \mathbb{E}[x_{T-1}^2] + \mathbb{E}[J_T^*(x_{T-1})].$$

We can write the objective function at time $T-1$ in the following nominal form

$$\hat{a}_{T-1} x_{T-1} - \hat{b}_{T-1} x_{T-1}^2 + \hat{\gamma}_{T-1} \mathbb{E}[\mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \mathbb{E}[\mathbf{P}_T]$$

with

$$\begin{aligned}
\hat{a}_{T-1} &= a_{T-1} + a_T \left[\mathbb{E}[e_T] - \left(1 - \frac{1}{n} A_T \right) \mathbb{E}[\mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \mathbb{E}[e_T \mathbf{P}'_T] \right], \\
\hat{b}_{T-1} &= b_{T-1} + b_T \left[\mathbb{E}[e_T^2] - 2 \left(1 - \frac{1}{n} A_T \right) \mathbb{E}[e_T \mathbf{P}'_T] \mathbb{E}^{-1}[\mathbf{P}_T \mathbf{P}'_T] \mathbb{E}[e_T \mathbf{P}'_T] - A_T \right], \\
\hat{\gamma}_{T-1} &= \left(1 - \frac{1}{n} A_T \right) \frac{a_T^2}{4b_T}.
\end{aligned}$$

Once again, we follow the same approach as presented above. First, we obtain the derivatives with respect to the allocation \mathbf{u}_{T-1} at beginning of time period

$T - 1$ and solve for the stationary point. Note that the derivatives will be piecewise linear functions of \mathbf{u}_{T-1} . Then we check whether the allocation exceeds the initial amount of available resources; whenever this is the case, we introduce the Lagrangian multiplier and obtain the corresponding optimal allocation.

This method can be carried out for any time period. In general, for time period t , the objective will be

$$\hat{a}_t x_t - \hat{b}_t x_t^2 + \sum_{s=t}^{T-1} \hat{\gamma}_s \mathbb{E}[\mathbf{P}'_{t+1}] \mathbb{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbb{E}[\mathbf{P}_{t+1}]$$

with

$$\begin{aligned} \hat{a}_t &= a_t + a_{t+1} \left[\mathbb{E}[e_{t+1}] - \left(1 - \frac{1}{n} A_{t+1}\right) \mathbb{E}[\mathbf{P}'_{t+1}] \mathbb{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbb{E}[e_{t+1} \mathbf{P}'_{t+1}] \right], \\ \hat{b}_t &= b_t + b_{t+1} \left[\mathbb{E}[e_{t+1}^2] - 2 \left(1 - \frac{1}{n} A_{t+1}\right) \mathbb{E}[e_{t+1} \mathbf{P}'_{t+1}] \right. \\ &\quad \left. \times \mathbb{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbb{E}[e_{t+1} \mathbf{P}'_{t+1}] - A_{t+1} \right], \\ \hat{\gamma}_t &= \left(1 - \frac{1}{n} A_{t+1}\right) \frac{\hat{a}_{t+1}^2}{4 \hat{b}_{t+1}}, \end{aligned}$$

and

$$A_t = \mathbb{I}\{\mathbb{E}^{-1}[\mathbf{P}_t \mathbf{P}'_t] [\hat{a}_t / (2 \hat{b}_t) \mathbb{E}[\mathbf{P}_t] - (\mathbb{E}[e_t \mathbf{P}'_t] + \mathbf{1}) x_{t-1}] \geq 0\}.$$

The optimal allocation then takes the form

$$\begin{aligned} \mathbf{u}_t^* &= \mathbb{E}^{-1}[\mathbf{P}_t \mathbf{P}'_t] \left[\frac{\hat{a}_{t+1}}{2 \hat{b}_{t+1}} \mathbb{E}[\mathbf{P}_t] - \mathbb{E}[e_t \mathbf{P}'_t] x_t \right] \\ &\quad - \frac{1}{n} \left\{ \mathbb{E}^{-1}[\mathbf{P}_t \mathbf{P}'_t] \left[\frac{\hat{a}_{t+1}}{2 \hat{b}_{t+1}} \mathbb{E}[\mathbf{P}_t] - \mathbb{E}[e_t \mathbf{P}'_t] x_t \right] - x_t \right\}^+. \end{aligned}$$

Another important fact is that at time t , the objective function will be the joint combination of at most $T - t + 1$ quadratic functions. As seen at each step of the backward induction, we are introducing at most one break point, which increases the number of quadratic functions by at most one. This yields the following result.

Proposition 2. *The objective function at time t is a piecewise quadratic function with at most $T - t + 1$ pieces.*

Our algorithmic framework for solving the stochastic dynamic programming problem can be summarized as follows:

At the beginning of period t , the initial amount of resources is x_{t-1} and the optimal allocation of resources is given by the solution of the convex programming problem for period t . Moreover, a function $f_t(x_{t-1})$ is obtained as the localization of the objective function at period t , where this function is increasing and concave.

3.3 Extensions on Resource Flexibility

Our general solution framework thus far has strictly enforced the available resource constraint x_t on allocation decisions at any time t . We now relax this constraint and consider the more general case that allows additional resources to be acquired at a per-unit cost of \tilde{c}_t . While this introduces a more complex cost structure, the formulation and solution gain more flexibility and agility in allocating resources for each time period. We therefore adapt our algorithmic framework to address the following objective function at each time t

$$\tilde{a}_t x_t - \tilde{b}_t x_t^2 - \tilde{c}_t [\mathbf{1}' \mathbf{u}_t - x_{t-1}]^+ + \sum_{s=t}^{T-1} \tilde{\gamma}_s \mathbf{E}[\mathbf{P}'_{t+1}] \mathbf{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbf{E}[\mathbf{P}_{t+1}],$$

where

$$\begin{aligned} \tilde{a}_t &= a_t + a_{t+1} \left\{ \mathbf{E}[e_{t+1}] - (1 - \tilde{c}_t A_{t+1}) \mathbf{E}[\mathbf{P}'_{t+1}] \mathbf{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbf{E}[e_{t+1} \mathbf{P}'_{t+1}] \right\}, \\ \tilde{b}_t &= b_t + b_{t+1} \left\{ \mathbf{E}[e_{t+1}^2] - 2(1 - \tilde{c}_t A_{t+1}) \mathbf{E}[e_{t+1} \mathbf{P}'_{t+1}] \right. \\ &\quad \left. \times \mathbf{E}^{-1}[\mathbf{P}_{t+1} \mathbf{P}'_{t+1}] \mathbf{E}[e_{t+1} \mathbf{P}'_{t+1}] - A_{t+1} \right\} \\ \tilde{\gamma}_t &= (1 - \tilde{c}_t A_{t+1}) \frac{\tilde{a}_{t+1}^2}{4\tilde{b}_{t+1}}. \end{aligned}$$

The optimal solution then takes the analogous form

$$\mathbf{u}_t^* = \mathbf{E}^{-1}[\mathbf{P}_t \mathbf{P}'_t] \left[\frac{\tilde{a}_{t+1}}{2\tilde{b}_{t+1}} \mathbf{E}[\mathbf{P}_t] - \mathbf{E}[e_t \mathbf{P}'_t] x_t \right].$$

4 Extensions on General Formulations

The stochastic dynamic programming problem considered above is most related to the various real-world applications motivating this study. Our algorithmic solution framework derived herein, however, is quite general and can be applied to a much wider range of stochastic dynamic program formulations, including those in which different objective functions are used to reflect distinct aspects of the problem.

One performance criteria often used to measure the success of entities of interest concerns the likelihood of achieving certain pre-specified targets. This class of problems can be readily modeled by a chance constraint, or equivalently by an objective function with a chance element. More specifically, we can have the alternative objective function

$$\max_{(u_1, \dots, u_T)} \sum_{t=1}^T w_t \mathbf{P}[x_t > d_t],$$

where the d_t are any type of target return values that the decision makers usually consider to be a primary indicator of the progress and success for each of the various entities, and the w_t are user-specified weights that differentiate the importance of each chance element over time. From Tchebyshev's inequality, we know that, when x_t is large,

$$\mathbf{P}[x_t > d_t] \leq \frac{\mathbf{Var}[x_t]}{(d_t - \mathbf{E}[x_t])^2},$$

and therefore we can again obtain a proper weighted objective function of the mean and variance of x_t to formulate an equivalent problem to which our solution framework applies.

More generally, variations of our algorithmic solution framework can be applied to address the following generic class of dynamic utility maximization problems:

$$\max_{(u_1, \dots, u_T)} \sum_{t=1}^T U_t(\mathbf{E}[x_t], \mathbf{Var}[x_t])$$

where the $U_t(\cdot, \cdot)$ represent various utility functions of the mean and variance of the returns which exhibit properties that are increasing in $\mathbf{E}[x_t]$ and decreasing in $\mathbf{Var}[x_t]$. This includes a wide range of risk measures, especially

those exhibiting convexity, such as the entire class of value-at-risk measures. More recent measures such as least partial moments can also be incorporated within our general approach.

Lastly, there are many studies related to continuous time mean-variance portfolio analysis that are also of interest. Our algorithmic solution framework derived herein can be similarly applied to address this general class of problems.

5 Conclusions

We devised a general solution framework to support the dynamic allocation of generic resources among generic entities of interest over a finite horizon during which entities generate random returns. With a primary objective to maximize the expected return while restricting the risk to a manageable level for each period, we derived an algorithmic solution framework for determining the optimal portfolio of resource allocations. Our general solution framework can be applied to address the planning and management of entities of interest across a broad spectrum of application domains.

References

- [1] BARDHAN, I., KUAFFMAN, R., AND NARANPANNAWE, S. IT Project Portfolio Optimization: A Risk Management Approach to Software Development Governance, *IBM Journal of Research & Development*, 2010.
- [2] GAO, J., ZHU, K., LI, D., AND CAO, X. Dynamic Mean-LPM and Mean-CVaR Portfolio Optimization in Continuous-time, <http://arxiv.org/abs/1402.3464>, 2015.
- [3] GUSTAFFSON, J. AND SELO, A. Contingent Portfolio Programming for the Management of Risky Projects, *Operations Research*, 2005.
- [4] LI, D., CHAN, T.F., AND NG, W.L. Safety-First Dynamic Portfolio Selection, *Dynamics of Continuous, Discrete and Impulsive Systems*, 4(4): 585-600, 1998.

- [5] LI, D. AND NG, W.L. Optimal Dynamic Portfolio Selection: Multi-Period Mean-Variance Formulation, *Mathematical Finance*, 10(3): 387-406, 2000.
- [6] ROCKAFELLAR, T. *Convex Analysis*, Princeton Univ. Press, 1996.
- [7] SMITH, J.E. AND NAU, R.F. Valuing Risky Projects: Option Pricing Theory and Decision Analysis, *Management Science*, 41(5): 795-816.
- [8] ZHOU, X. AND LI, D. Continuous Time Mean-Variance Portfolio Selection: A Stochastic LQ Framework, *Applied Mathematics and Optimization*, 42: 19-33, 2000.