# Mean-semivariance optimal portfolios in discrete time using a game-theoretic approach

Kristoffer Lindensjö[1] and Vilhelm Niklasson*[1]

[1]Department of Mathematics, Stockholm University, Sweden

**Abstract**

This paper introduces a novel recursive scheme for optimal asset allocation based on a mean-semivariance reward functional and a game-theoretic approach in a discrete-time setting. Unlike established frameworks that can handle variance as a risk measure, this study shifts focus to semivariance, which cannot be handled by existing theory due to aspects of its definition, including the use of an indicator function. To address this problem and the corresponding challenges of time-inconsistency in multi-period investment decisions, we propose an extended Bellman equation to find a Nash equilibrium. The primary contributions of this paper are the development of a new theoretical framework and a comprehensive numerical study that investigates its properties. The results of the numerical study indicate that our proposed method shows potential in achieving favorable investment outcomes.

***Keywords*:** Time inconsistency; Optimal portfolio; Semivariance; Equilibrium control

## 1 Introduction

Stochastic optimal control constitutes a fundamental model for dynamic decision-making under uncertainty and is widely applied in fields ranging from engineering to economics. Central to this field is the discrete-time controlled Markov process $X = (X_n)$, with $X_n$ representing the state at discrete times $n \in \{0, 1, 2, \dots, T\}$ and $\boldsymbol{u}$ as the corresponding feedback control law, meaning that the control applied only depends on the current state of $X$ and time. We use the notation $X^{\boldsymbol{u}} = (X_n^{\boldsymbol{u}})$ to denote this relationship (see Section 2 for details). In this framework, a conventional optimization problem aims to maximize the reward functional

$$J_n(x, \boldsymbol{u}) = \mathbb{E}_{n,x}\left[\sum_{k=n}^{T-1} H_k(X_k^{\boldsymbol{u}}, \boldsymbol{u}_k(X_k^{\boldsymbol{u}})) + F(X_T^{\boldsymbol{u}})\right], \tag{1.1}$$

*Corresponding author: niklasson@math.su.se

where $\mathbb{E}_{n,x}[\cdot]$ denotes the expected value conditioned on the state $X_n = x$, and $H_k$ and $F$ are predefined real-valued functions. These problems are inherently time-consistent, aligning with the Bellman optimality principle: an optimal control law for the interval $\{n, \dots, T\}$ remains optimal for any subinterval $\{m, \dots, T\}$, with $n \leq m$ [2]. This principle is foundational in deriving the so-called Bellman equation and formulating recursive strategies for optimal control.

However, this standard approach for solving stochastic control problems generally fails when the reward functional cannot be formulated according to (1.1), which is a common occurrence in mathematical finance [5]. A notable instance is mean-variance utility optimization in portfolio theory, corresponding to the reward functional

$$J_n(x, \boldsymbol{u}) = \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right] - \frac{\gamma}{2}\mathrm{Var}_{n,x}[X_T^{\boldsymbol{u}}], \tag{1.2}$$

where $X_n^{\boldsymbol{u}}$ corresponds to the value of a financial portfolio at time $n$, $\mathrm{Var}_{n,x}[\cdot]$ denotes the variance conditional on $X_n = x$, the constant $\gamma > 0$ represents the risk aversion of the investor, and the control law $\boldsymbol{u}$ determines the amount invested in a risky asset with the remainder of the wealth in the portfolio being invested in a non-risky asset [5]. The reward functional (1.2) cannot be written in the form (1.1), and the optimal strategy changes over time due to the nonlinear dependency on the expected value of terminal wealth. Specifically, a decision plan that is optimal at one point typically does not remain optimal when reassessed at a later date. This phenomenon, termed time-inconsistency, also arises in cases where $H_k$ or $F$ depend on the initial state $(n, x)$. Consequently, in such cases, the standard Bellman equation cannot be applied straightforwardly [5].

Addressing this challenge necessitates innovative approaches, such as framing the problem in game-theoretic terms (see, e.g, [15, 14, 1, 5, 3]). From this perspective, the entire decision-making process is construed as a noncooperative game involving a distinct player for each time point $n$. Each player aims to maximize their reward functional, anticipating that subsequent players will adopt a similar strategy. Specifically, the player at time $T-1$ addresses a standard optimization problem by optimizing $J_{T-1}(x, u_{T-1})$, resulting in an optimal control $\hat{u}_{T-1}(x)$. Subsequently, the player at $T-2$ formulates their optimization strategy by assuming that the player at $T-1$ adheres to the control $\hat{u}_{T-1}(x)$. In this context, the $T-2$ player seeks to maximize their own reward functional, resulting in the determination of their optimal control $\hat{u}_{T-2}(x)$. This recursive process continues, resulting in a subgame perfect Nash equilibrium. For a further motivation and interpretation of the game-theoretic approach to time-inconsistent problems, refer to [5] and [12]. For an in-depth treatment of time-inconsistent stochastic control and its applications in finance, refer to the monograph [4].

Using this game-theoretic approach, Björk and Murgoci [5] develops a general framework for time-inconsistent stochastic control in discrete time, relying on a so-called extended Bellman equation, for a large class of reward functionals in the form

$$J_n(x, \boldsymbol{u}) = \mathbb{E}_{n,x}\left[\sum_{k=n}^{T-1} H_k(n, x, X_k^{\boldsymbol{u}}, \boldsymbol{u}_k(X_k^{\boldsymbol{u}})) + F(n, x, X_T^{\boldsymbol{u}})\right] + G(n, x, \mathbb{E}_{n,x}[X_T^{\boldsymbol{u}}]), \tag{1.3}$$

where $H_k, F$ and $G$ are known real-valued functions. This framework incorporates the mean-variance reward model (1.2), as well as other types of problems from financial economics, including features of habit formation and non-exponential discounting [5]. Notably, in the case of the mean-variance reward functional (1.2), Björk and Murgoci successfully derives analytical expressions for an equilibrium control strategy.

However, using variance as a risk measure may not always align with practical investment strategies. Investors often prioritize mitigating losses while valuing the potential for significant positive returns. In this context, a more representative measure of risk is semivariance, which specifically targets downside deviations. Semivariance is defined according to

$$\mathrm{SV}_{n,x}\left(X_T^{\boldsymbol{u}}\right) \coloneqq \mathbb{E}_{n,x}\left[\left(X_T^{\boldsymbol{u}} - \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right]\right)^2 I_{X_T^{\boldsymbol{u}} < \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right]}\right]. \tag{1.4}$$

In his seminal 1959 book [13], Markowitz, a key figure in the development of modern portfolio theory, devoted an entire chapter – Chapter IX – specifically to the concept of semivariance. He even claimed that portfolios constructed using semivariance as a risk measure often result in "better portfolios" than those based on variance. However, using semivariance instead of variance is typically more challenging from a computational perspective, a fact acknowledged by Markowitz and corroborated by subsequent research in the field (see, e.g, [10, 11, 8, 9]). Incorporating semivariance into the reward functional in a stochastic optimal control problem also introduces additional complexity. Indeed, replacing the variance in (1.2) with semivariance leads to a reward functional that cannot be expressed in the form (1.3), thus presenting a problem that cannot be addressed within the general framework developed in [5].

In the present paper, we consider a modified version of the standard game-theoretic approach to time-inconsistent stochastic control suitable for mean-semivariance optimization, i.e., corresponding to the reward functional

$$J_n(x, \boldsymbol{u}) = \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right] - \frac{\gamma}{2}\mathrm{SV}_{n,x}\left[X_T^{\boldsymbol{u}}\right]. \tag{1.5}$$

The main contributions of this paper are the derivation of an extended Bellman equation for this optimization criterion and a thorough investigation of its properties relative to the allocation strategies based on the mean-variance reward functional. The structure of the remaining sections of this paper is as follows: Section 2 presents the setup and the game-theoretic formulation; Sections 3 and 4 discuss applications to one-period and multi-period settings, respectively; and Section 5 offers concluding remarks, highlighting the implications and potential applications of our findings.

## 2 Setup and game-theoretic formulation

We study a discrete-time controlled Markov process $X^{\boldsymbol{u}} = (X_n^{\boldsymbol{u}})$, $n \in \{0, 1, \dots, T\}$, evolving in the state space $\mathbb{R}$. The control process $\boldsymbol{u}$, which takes values in the same space, is defined as a feedback control law where $\boldsymbol{u}_n = \boldsymbol{u}_n(X_n^{\boldsymbol{u}})$ for each $n < T$ according to the standard construction (see, e.g., [5]). In Sections 3 and 4 we study two specific models for specifying an explicit relationship between $\boldsymbol{u}$ and $X^{\boldsymbol{u}}$. We also make the following assumption about the controlled Markov process.

**Assumption 2.1.** *For each control strategy $\boldsymbol{u}$ and each $n$, the random variable $X_n^{\boldsymbol{u}}$ has a discrete distribution with finite mean and finite variance.*

Inspired by the framework established by Björk and Murgoci [5], this paper proposes a modification of the mean-variance reward functional (1.2). The primary objective is to explore the maximization of the mean-semivariance reward functional (1.5) within a game-theoretic framework, striving to identify an equilibrium control.

The following equilibrium definition aligns with the standard equilibrium condition for time-inconsistent stochastic control problems in discrete time (see [5] for an interpretation).

**Definition 2.2** (Equilibrium control)**.** Consider a fixed control law $\hat{\boldsymbol{u}}$ and an arbitrary point $(n, x)$ with $n < T$. For a fixed $u \in \mathbb{R}$, define the control law $\hat{\boldsymbol{u}}^{u,n}$ for any $y \in \mathbb{R}$ according to

$$\hat{\boldsymbol{u}}_k^{u,n}(y) = \begin{cases} \hat{\boldsymbol{u}}_k(y) & \text{for } k = n+1, \dots, T-1, \\ u & \text{for } k = n, \end{cases} \tag{2.1}$$

i.e., $\hat{\boldsymbol{u}}^{u,n}$ deviates from $\hat{\boldsymbol{u}}$ only at $n$ with the value $u$. We say that $\hat{\boldsymbol{u}}$ is a subgame perfect Nash equilibrium control if for every fixed $(n, x)$, with $n < T$, we have

$$\sup_{u \in \mathbb{R}} J_n(x, \hat{\boldsymbol{u}}^{u,n}) = J_n(x, \hat{\boldsymbol{u}}). \tag{2.2}$$

Moreover, if such $\hat{\boldsymbol{u}}$ exists, the equilibrium value function $V$ is defined by

$$V_n(x) \coloneqq J_n(x, \hat{\boldsymbol{u}}). \tag{2.3}$$

In a discrete-time framework with a finite horizon $T$, equilibrium strategies are typically derived through backward induction. Moreover, for many time-inconsistent problems, it is feasible to formulate a system of equations that extend the traditional Bellman equation, facilitating the determination of the value function $V$. In particular, Björk and Murgoci developed such an adaptation of the traditional Bellman equation to encompass a broad spectrum of reward functionals in their foundational study [5]. However, as noted in Section 1, their framework does not encompass models where the reward functional involves a term defined by the expectation of a function, wherein one of the variables is the expectation of the terminal value, explicitly $\mathbb{E}_{n,x}\left[C(X_T^{\boldsymbol{u}}, \mathbb{E}_{n,x}(X_T^{\boldsymbol{u}}))\right]$, where $C$ is a real-valued function. Hence, the framework of [5] cannot be used to find an equilibrium for the mean-semivariance reward functional (1.5). The goal of the present paper is develop methods for the study of such an equilibrium.

To this end, in the remainder of this section, we present a recursive scheme for the calculation of the mean-semivariance reward, which enables us to establish a method for finding a Nash equilibrium. We first introduce two useful function sequences.

**Definition 2.3.** For any $(n, x)$, with $n < T$, control law $\boldsymbol{u}$, and $z \in \mathbb{R}$ we define the function sequences $(g_n^{\boldsymbol{u}})$ and $(h_n^{\boldsymbol{u}})$ by

$$g_n^{\boldsymbol{u}}(x) \coloneqq \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right], \tag{2.4}$$

and

$$h_n^{\boldsymbol{u}}(x,z) \coloneqq \mathbb{E}_{n,x}\left[I_{X_T^{\boldsymbol{u}}=z}\right] = \mathbb{P}_{n,x}\left[X_T^{\boldsymbol{u}} = z\right], \tag{2.5}$$

where $I$ is the indicator function.

Note that $h_n^{\boldsymbol{u}}(x,z)$ represents the conditional probability of attaining a terminal portfolio value of $z$, given the present value $x$ at time-step $n$ when using the control $\boldsymbol{u}$. Before presenting a lemma relevant to these function sequences, we clarify essential notation used throughout the paper.

*Remark.* The notation $X_{n+1}^{u}$ denotes the stochastic value of the controlled process at time $n+1$ conditioned on $X_n^{\boldsymbol{u}}$, emphasizing that the state process value at $n+1$ depends only on $u = \boldsymbol{u}_n(X_n^{\boldsymbol{u}})$, as viewed from time $n$. With a slight abuse of notation, we can thus describe the deviation from a control law $\boldsymbol{u}$ using the value $u$ at any specific point (see, e.g., Theorem 2.6).

**Lemma 2.4.** *For any $(n, x)$, with $n < T$, and control law $\boldsymbol{u}$, the function sequences $(g_n^{\boldsymbol{u}})$ and $(h_n^{\boldsymbol{u}})$ can be expressed recursively as*

$$g_n^{\boldsymbol{u}}(x) = \mathbb{E}_{n,x}\left[g_{n+1}^{\boldsymbol{u}}(X_{n+1}^{u})\right], \tag{2.6}$$

$$h_n^{\boldsymbol{u}}(x,z) = \mathbb{E}_{n,x}\left[h_{n+1}^{\boldsymbol{u}}(X_{n+1}^{u},z)\right], \tag{2.7}$$

*with terminal conditions $g_T(x) = x$, $h_T(x,z) = I_{x=z}$ and $u = \boldsymbol{u}_n(x)$.*

*Proof.* From the Markov property, it follows, for each $n < T$, that

$$g_n^{\boldsymbol{u}}(x) = \mathbb{E}_{n,x}\left[X_T^{\boldsymbol{u}}\right] = \mathbb{E}_{n,x}\left[\mathbb{E}_{n+1,X_{n+1}^{u}}\left[X_T^{\boldsymbol{u}}\right]\right] = \mathbb{E}_{n,x}\left[g_{n+1}^{\boldsymbol{u}}(X_{n+1}^{u})\right], \tag{2.8}$$

and

$$h_n^{\boldsymbol{u}}(x,z) = \mathbb{E}_{n,x}\left[I_{X_T^{\boldsymbol{u}}=z}\right] = \mathbb{E}_{n,x}\left[\mathbb{E}_{n+1,X_{n+1}^{u}}\left[I_{X_T^{\boldsymbol{u}}=z}\right]\right] = \mathbb{E}_{n,x}\left[h_{n+1}^{\boldsymbol{u}}(X_{n+1}^{u},z)\right]. \tag{2.9}$$

The terminal conditions follow directly from the definitions. □

We are now prepared to formulate the reward functional for our mean-semivariance model (1.5), in terms of the the above defined sequences.

**Lemma 2.5.** *For any $(n, x)$, with $n < T$, and control law $\boldsymbol{u}$, the reward functional* (1.5) *satisfies the recursion*

$$\begin{aligned} J_n(x, \boldsymbol{u}) =& \mathbb{E}_{n,x}\left[g^{\boldsymbol{u}}_{n+1}(X^u_{n+1})\right] \\ & - \frac{\gamma}{2}\sum_{i=0}^{\infty}\left(z_i - \mathbb{E}_{n,x}\left[g^{\boldsymbol{u}}_{n+1}(X^u_{n+1})\right]\right)^2 I_{z_i<\mathbb{E}_{n,x}\left[g^{\boldsymbol{u}}_{n+1}(X^u_{n+1})\right]}\mathbb{E}_{n,x}\left[h^{\boldsymbol{u}}_{n+1}(X^u_{n+1}, z_i)\right], \end{aligned} \tag{2.10}$$

*where $\{z_i : i \in \mathbb{N}\}$ represents the set of all possible terminal portfolio values $X^{\boldsymbol{u}}_T$ under $\boldsymbol{u}$.*

*Proof.* The mean-semivariance reward functional (1.5) corresponds to

$$J_n(x, \boldsymbol{u}) = \mathbb{E}_{n,x}\left[X^{\boldsymbol{u}}_T\right] - \frac{\gamma}{2}\mathbb{E}_{n,x}\left[\left(X^{\boldsymbol{u}}_T - \mathbb{E}_{n,x}\left[X^{\boldsymbol{u}}_T\right]\right)^2 I_{X^{\boldsymbol{u}}_T<\mathbb{E}_{n,x}\left[X^{\boldsymbol{u}}_T\right]}\right]. \tag{2.11}$$

Utilizing the definitions of the function sequences $(g^{\boldsymbol{u}}_n)$ and $(h^{\boldsymbol{u}}_n)$, combined with the definition of expectation, yields

$$J_n(x, \boldsymbol{u}) = g^{\boldsymbol{u}}_n(x) - \frac{\gamma}{2}\sum_{i=1}^{\infty}\left(z_i - g^{\boldsymbol{u}}_n(x)\right)^2 I_{z_i<g^{\boldsymbol{u}}_n(x)} h^{\boldsymbol{u}}_n(x, z_i). \tag{2.12}$$

The result now follows directly from Lemma 2.4, achieved by substituting $g^{\boldsymbol{u}}_n(x)$ and $h^{\boldsymbol{u}}_n(x, z_i)$ with their expressions framed in terms of a forward expectation by one step.

□

Next, deriving an expression for an equilibrium control and value function, corresponding to an extended Bellman equation for the present problem, is straightforward. Indeed, we obtain the following verification theorem.

**Theorem 2.6.** *Suppose there exists a control law $\hat{\boldsymbol{u}}$ which satisfies the recursion*

$$\begin{aligned} \hat{\boldsymbol{u}}_n(x) =& \arg\max_{u\in\mathbb{R}}\left\{\mathbb{E}_{n,x}\left[g^{\hat{\boldsymbol{u}}}_{n+1}(X^u_{n+1})\right]\right. \\ & \left. - \frac{\gamma}{2}\sum_{i=0}^{\infty}\left(z_i - \mathbb{E}_{n,x}\left[g^{\hat{\boldsymbol{u}}}_{n+1}(X^u_{n+1})\right]\right)^2 I_{z_i<\mathbb{E}_{n,x}\left[g^{\hat{\boldsymbol{u}}}_{n+1}(X^u_{n+1})\right]}\mathbb{E}_{n,x}\left[h^{\hat{\boldsymbol{u}}}_{n+1}(X^u_{n+1}, z_i)\right]\right\}, \end{aligned} \tag{2.13}$$

*for each $n < T$ and $x \in \mathbb{R}$. Then $\hat{\boldsymbol{u}}$ is an equilibrium control with value function $V_n(x) = J_n(x, \hat{\boldsymbol{u}})$.*

*Remark.* Note that in (2.13), the variable $u$ is essentially a free variable determined, for each $(n, x)$, by a maximization problem. The interpretation is that it is possible to deviate from the equilibrium control law $\hat{\boldsymbol{u}}$ at any $(n, x)$ by using an alternative control $u$ solely at that point. However, such a deviation is not optimal in the sense of (2.13).

*Proof.* Consider any arbitrary fixed $(n, x)$ and a control which may differ from $\hat{\boldsymbol{u}}$ only at time $n$, which in our notation is denoted by $\hat{\boldsymbol{u}}^{u,n}$, meaning that the differing value is $u$. By (2.13) and from Lemma 2.5, we have

$$J_n(x, \hat{\boldsymbol{u}}^{u,n}) \leq J_n(x, \hat{\boldsymbol{u}}), \tag{2.14}$$

for all $u \in \mathbb{R}$, with equality if $u = \hat{\boldsymbol{u}}_n(x)$. The statement holds for all $(n, x)$, $n < T$, and thus the result follows by Definition 2.2. □

*Remark.* It is straightforward to adapt Theorem 2.6 to incorporate a state-dependent risk aversion $\gamma = \gamma(x)$ coefficient.

## 2.1 Our approach versus a naive Monte Carlo approach

Theorem 2.6 formulates the equilibrium control and the value function recursively using $(g_n^{\boldsymbol{u}})$ and $(h_n^{\boldsymbol{u}})$. The corresponding computation involves evaluating a sum that can be computationally intensive. This complexity renders the calculations more intricate compared to the mean-variance problem of [5], which offers an explicit solution. However, it is important to note that our recursive approach, despite its computational demands, significantly streamlines the computation of the equilibrium control compared to, for example, a naive Monte Carlo approach. Indeed, without the recursion in (2.13), one would naturally resort to simulation, which typically entails a higher computational burden. In contrast, our proposed approach circumvents this by requiring only a one-step look-ahead at each iteration.

Let us illustrate the difference between a Monte Carlo simulation approach and our recursive method for finding equilibrium controls by way of an example. Consider a scenario where we are currently at time $T - 3$. To determine the equilibrium control at $T - 3$ using a Monte Carlo approach, one could proceed as follows:

- **At** $T - 1$:
  - Perform $N$ simulations for each combination of possible values for $x$ and $u$ to find the equilibrium control $\hat{\boldsymbol{u}}_{T-1}(x)$ by selecting the control $u$ that gives the highest average reward based on (1.5) for each $x$. Assuming a grid of size $K$ for $x$ and a grid of the same size for $u$, this step requires $NK^2$ simulations.
- **At** $T - 2$:
  - Perform $2NK^2$ simulations to find the equilibrium control $\hat{\boldsymbol{u}}_{T-2}(x)$, taking the equilibrium control from the previous step as given for the transition from $T - 1$ to $T - 2$.
- **At** $T - 3$:
  - In analogy with the previous step, perform $3NK^2$ simulations to find the equilibrium control $\hat{\boldsymbol{u}}_{T-3}(x)$.

Overall, the Monte Carlo approach requires $NK^2(1 + 2 + 3) = 6NK^2$ simulations to find the equilibrium control $\hat{\boldsymbol{u}}_n(x)$ for $n = T - 3, T - 2, T - 1$. In contrast, our recursive method operates as follows:

- **At** $T-1$:
  - Perform calculations corresponding to (2.13) using a grid of size $K$ for $x$ and a grid of the same size for $u$. For each combination of $x$ and $u$, this requires performing a number of calculation steps, denoted by $L$ (the exact number of steps depends on the specific model). Using a straightforward approach for maximization by varying $u$ for each $x$, it is necessary to perform $LK^2$ calculation steps to find the equilibrium control $\hat{\boldsymbol{u}}_{T-1}(x)$.
- **At** $T-2$:
  - Perform another $LK^2$ calculation steps to find the equilibrium control $\hat{\boldsymbol{u}}_{T-2}(x)$, using the fact that only one-step ahead computations are necessary.
- **At** $T-3$:
  - Again, perform $LK^2$ computations to find the equilibrium control $\hat{\boldsymbol{u}}_{T-3}(x)$.

Overall, our recursive approach requires $3LK^2$ computations to determine the equilibrium control $\hat{\boldsymbol{u}}_n(x)$ for $n = T-3, T-2, T-1$. The difference between the computational complexity for the two methods that we want to stress is:

- **Monte Carlo method**: The computational complexity scales with $(T-n)^2$.
- **Recursive method**: The computational complexity scales linearly with $T-n$.

Figure 1 visually illustrates this disparity.

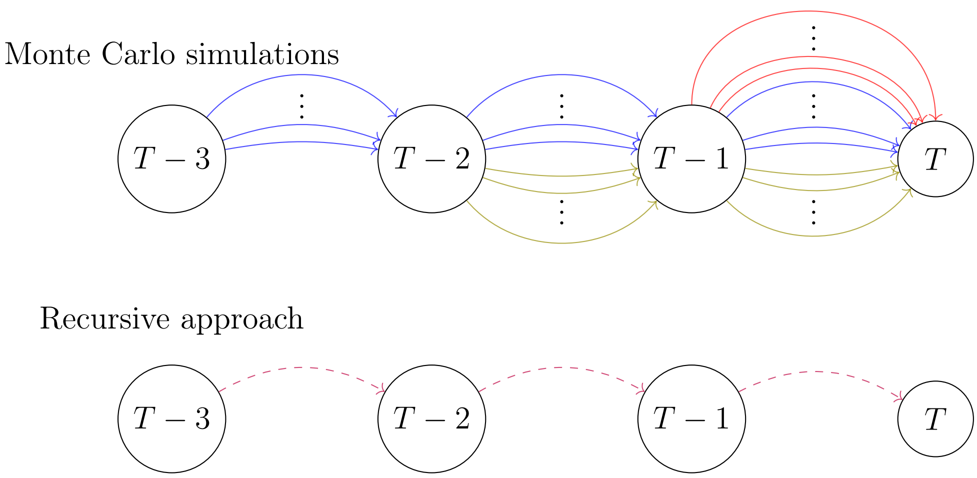


Figure 1: Comparison between our recursive approach and full-path Monte Carlo simulations for finding equilibrium controls.

# 3 One-period setting

Having established the theoretical foundations for our analysis, we now proceed to examine the implications in both one-period and multi-period settings. This exploration will enable us to assess the practicality and effectiveness of the proposed method for different time horizons. The code for the numerical study is available at `https://github.com/vilnik/mean-semivariance`.

In this section, we consider our portfolio selection problem in a one-period model with initial time being 0 and the final time being $T = 1$. It is obvious that time-inconsistency cannot arise in one-period models. In particular, we can simplify our analysis by considering a standard optimization problem without the complexities of game-theoretic considerations. In line with e.g., [5], we consider the portfolio dynamics

$$X_T^u = x(1 + r_f) + u(Y - r_f), \tag{3.1}$$

where $x$ is the initial portfolio value, $r_f$ is the risk-free rate, $Y$ is a random variable corresponding to the return of the risky asset, and the variable $u$ signifies the control action dictating the initial dollar amount invested in the risky asset.

Given the absence of strategic interaction, which is fundamental in the multi-period setting, the term optimal control is more apt than equilibrium control for the one-period setting. We obtain the following result for the optimal control.

*Proposition* 3.1. *The optimal control in the one-period setting, using the portfolio dynamics described by* (3.1) *and the mean-semivariance reward functional* (1.5)*, under the assumption that $Y$ does not follow a degenerate distribution and $\mathbb{E}[Y - r_f] > 0$, is given by*

$$\hat{u} = \frac{\mathbb{E}[Y - r_f]}{\gamma SV(Y)}, \tag{3.2}$$

*and the optimal value function satisfies*

$$V(x) = x(1 + r_f) + \frac{1}{2}\frac{\mathbb{E}^2[Y - r_f]}{\gamma SV(Y)}. \tag{3.3}$$

*Proof.* Given the portfolio dynamics, for any chosen $u$, it follows that

$$\mathbb{E}_{0,x}[X_T^u] = x(1 + r_f) + u\mathbb{E}[Y - r_f] \tag{3.4}$$

and

$$\begin{aligned}
\mathrm{SV}_{0,x}\left(X_T^u\right) =& \mathbb{E}_{0,x}\Big[\big(x(1 + r_f) + u(Y - r_f) - (x(1 + r_f) + u\mathbb{E}_{0,x}[Y - r_f])\big)^2 \\
& I_{x(1+r_f)+u(Y-r_f)<x(1+r_f)+u\mathbb{E}_{0,x}[Y-r_f]}\Big] \\
=& \mathbb{E}_{0,x}\Big[\big(uY - u\mathbb{E}_{0,x}[Y]\big)^2 I_{uY<u\mathbb{E}_{0,x}[Y]}\Big] \\
=& u^2\mathbb{E}_{0,x}\Big[\big(Y - \mathbb{E}_{0,x}[Y]\big)^2 I_{uY<u\mathbb{E}_{0,x}[Y]}\Big].
\end{aligned} \tag{3.5}$$

To address the inequality within the indicator function, three distinct cases based on the value of $u$ are considered: $u > 0$, $u < 0$, and $u = 0$ which we write as $u^+, u^-$ and $u^0$, respectively. For $u > 0$, the expression for the semivariance is

$$\mathrm{SV}_{0,x}(X_T^{u_+}) = u_+^2 \mathbb{E}_{0,x}\Big[\big(Y - \mathbb{E}_{0,x}[Y]\big)^2 I_{Y<\mathbb{E}_{0,x}[Y]}\Big] = u_+^2 \mathrm{SV}(Y). \tag{3.6}$$

In the case where $u < 0$, we have

$$\mathrm{SV}_{0,x}(X_T^{u_-}) = u_-^2 \mathbb{E}_{0,x}\Big[\big(Y - \mathbb{E}_{0,x}[Y]\big)^2 I_{Y>\mathbb{E}_{0,x}[Y]}\Big]. \tag{3.7}$$

When $u = 0$, the result is simply

$$\mathrm{SV}_{0,x}(X_T^{u_0}) = 0. \tag{3.8}$$

Hence the reward functional for $u > 0$ becomes

$$x(1 + r_f) + u_+\mathbb{E}[Y - r_f] - \frac{\gamma}{2} u_+^2 \mathrm{SV}(Y), \tag{3.9}$$

the case $u < 0$ yields

$$x(1 + r_f) + u_-\mathbb{E}[Y - r_f] - \frac{\gamma}{2} u_-^2 \mathbb{E}_{0,x}\Big[\big(Y - \mathbb{E}_{0,x}[Y]\big)^2 I_{Y>\mathbb{E}_{0,x}[Y]}\Big], \tag{3.10}$$

and for $u = 0$ we get

$$x(1 + r_f). \tag{3.11}$$

By setting the derivative of the first two reward functionals with respect to $u$ to be zero, we determine the optimal $u$ in the case $u > 0$ to be

$$\hat{u}_+ = \frac{\mathbb{E}[Y - r_f]}{\gamma \mathrm{SV}(Y)}, \tag{3.12}$$

which is a maximum due to the concavity of the reward functional. Similarly, the case $u < 0$ yields

$$\hat{u}_- = \frac{\mathbb{E}[Y - r_f]}{\gamma u_-^2 \mathbb{E}_{0,x}\Big[\big(Y - \mathbb{E}_{0,x}[Y]\big)^2 I_{Y>\mathbb{E}_{0,x}[Y]}\Big]}. \tag{3.13}$$

This expression, together with $\mathbb{E}[Y - r_f] > 0$, implies that $\hat{u}_-$ is positive, which is a contradiction. Hence, only $\hat{u}_+$ is a valid maximum. Inserting the expression of $\hat{u}_+$ into (3.9) gives the optimal value function (3.3), which is always greater than the value function corresponding to using $u = 0$. □

*Remark.* The mean-semivariance optimal control is invariant to the initial portfolio value $x$. This characteristic aligns with established results for the mean-variance reward functional, as delineated in [5]. Specifically, for the mean-variance framework in a one-period setting, the optimal control $u$ is determined by

$$\hat{u} = \frac{\mathbb{E}[Y - r_f]}{\gamma \operatorname{Var}(Y)}. \tag{3.14}$$

For symmetric distributions, it typically holds that $\operatorname{Var}(Y) = 2\operatorname{SV}(Y)$. In the one-period model with a symmetric distribution, this creates a straightforward relationship between the mean-variance and mean-semivariance allocations, corresponding to investing twice as much in the risky asset in the latter compared to the former case.

## 3.1 Results based on binomial returns

Despite the availability of analytical solutions for the mean-semivariance optimization in the one-period case as seen above, it holds that employing our recursive game-theoretic approach based on Theorem 2.6 is beneficial for illustrative purposes. Indeed, it facilitates a clearer understanding of the numerical efficacy of the recursive method in approximating the analytical solution.

To effectively implement the game-theoretic approach, a pivotal aspect lies in establishing a grid of potential portfolio values at different time points. This grid serves as the foundation for computing the function sequences $(g_n^{\boldsymbol{u}})$ and $(h_n^{\boldsymbol{u}})$ at various nodes. Linear interpolation is then employed to estimate values between these discrete points. A critical parameter in this numerical implementation is the number of terminal portfolio values in the grid. This parameter directly influences the granularity and precision of the computations. It is important to note that the range of potential portfolio values in a single period is infinite. This is due to the fact that the future portfolio value is contingent on the variable $u$, which is not inherently constrained.

To illustrate this numerical approach, consider a one-period setting, and assume that $Y$ follows a distribution given by

$$Y = \begin{cases} r_u \text{ with probability } p \in (0,1), \\ r_d \text{ with probability } 1-p, \end{cases} \tag{3.15}$$

where $r_u$ and $r_d$ are the upward and downward return rates, respectively. The parameter $p$ denotes the probability of an upward movement in value, while $1-p$ represents the probability of a downward movement.

The binomial model is particularly suited for our mean-semivariance analysis due to its inherent property of generating non-symmetric distributions of portfolio returns. In such contexts, the mean-semivariance framework becomes especially relevant, as it focuses on downside risk. Additionally, its discrete distribution nature facilitates the numerical implementation of Theorem 2.6. The binomial model's extensive application in discrete time finance, notably in derivatives pricing, further highlights its relevance to our analysis, although we do not need to consider risk-neutral probabilities [7]. Nonetheless, it is important to note that any non-symmetric discrete distribution, such as the Poisson or geometric distribution, could potentially serve in similar analytical capacities.

In this context, Figure 2 serves as a visual representation of how the game-theoretic approach is applied within the one-period binomial model and its precision. Specifically, it showcases the influence of different numbers of terminal portfolio values on the derived optimal control. This visual analysis is conducted with $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$ and $\gamma = 2$. The figure not only includes the numerical and theoretical mean-semivariance optimal controls but also showcases the numerical and theoretical mean-variance optimal controls for comparison. Our recursive expression for the reward functional, detailed in Lemma 2.5, can easily be adapted to the mean-variance case by setting the indicator function to one at all times.

To establish the range for the terminal portfolio values, we consider the potential outcomes at the end of the investment period. The lower bound we have chosen corresponds to the scenario where the entire portfolio is invested in the risky asset, and the asset experiences a downward movement, resulting in $x(1 + r_d)$. Conversely, the upper bound represents the situation where the entire portfolio is allocated to the risky asset, and the asset undergoes an upward movement, leading to $x(1 + r_u)$. These bounds define the range of possible terminal portfolio values. In our numerical implementation, we discretize this range by considering portfolio values that are equally spaced between these two extremes.

*Remark.* There is no automatic assurance that the optimal control will fall within the interval $[0, x]$, a constraint subtly imposed by the permissible range of terminal portfolio values. However, rigorous testing and verification ensure that, for our selected parameters, the optimal control indeed lies within this range. Furthermore, extending the range to accommodate a broader spectrum of scenarios, if necessary, is straightforward and maintains the integrity of our analysis.

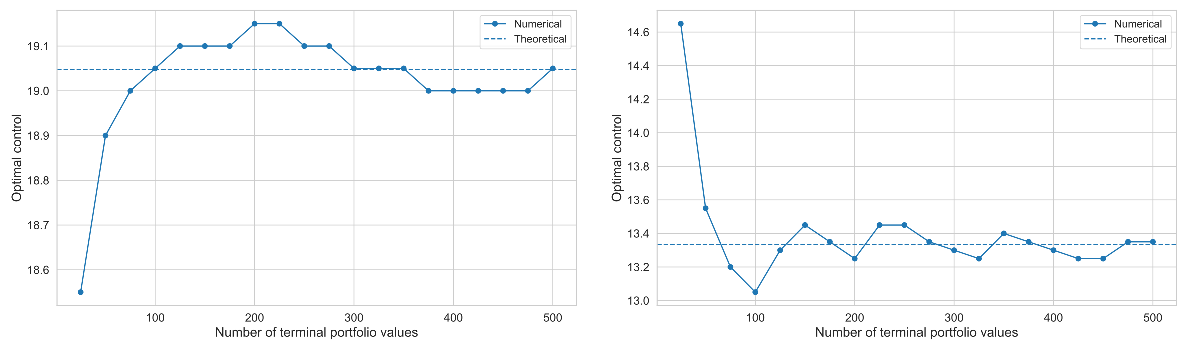


(a) Mean-semivariance optimal control

(b) Mean-variance optimal control

Figure 2: Optimal controls in a one-period binomial setting, focusing on the impact of varying the number of terminal portfolio values. Both numerical and theoretical results for mean-semivariance and mean-variance optimal controls are presented, utilizing parameters $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 2$.

The analysis depicted in Figure 2 indicates that our recursive approach, when applied to both the mean-semivariance and mean-variance models, demonstrates a tendency to converge towards the true value as the number of terminal portfolio values increases. It is important to note that in our numerical implementation, the control variable $u$ is incremented in steps of 0.05.

This discretization granularity significantly influences the rate of convergence, with smaller step sizes potentially leading to more precise approximations of the optimal control but requiring a larger computational effort.

It is also insightful to explore how the reward functional $J$ is influenced by the control variable $u$. This relationship is captured in Figure 3, which delineates the dependence of $J$ on the control strategies $u$ for both the mean-semivariance and mean-variance reward functionals, as defined in (1.5) and (1.2), respectively.

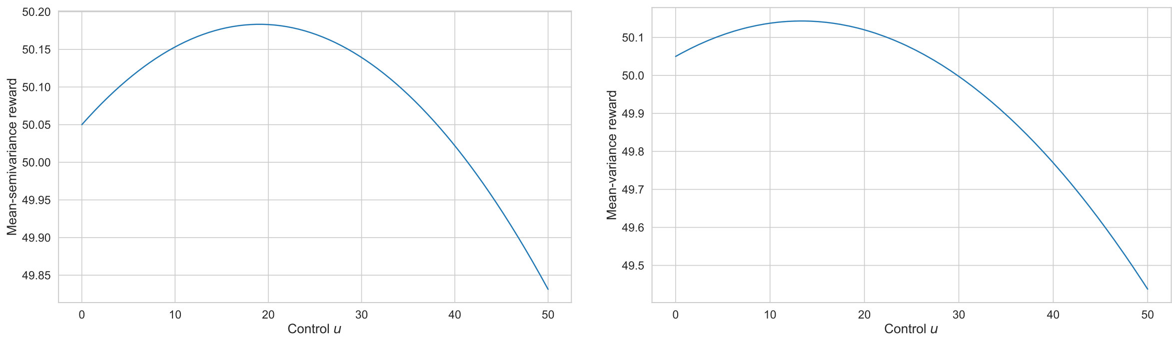


(a) Mean-semivariance reward

(b) Mean-variance reward

Figure 3: Rewards in a one-period binomial setting for different control values $u$, utilizing parameters $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 2$.

The analysis indicates that the reward functionals for both mean-semivariance and mean-variance exhibit concavity with respect to the control variable $u$. This characteristic is substantiated by the formulation presented in (3.9), along with its readily available counterpart for the mean-variance scenario. Notably, the analysis reveals that the incremental reward associated with variations in $u$ remains marginal across both frameworks in our setup.

# 4 Multi-period setting

In this section we consider a generalized version of dynamics of the one-period model. In particular, with a slight abuse of notation, we assume that the controlled portfolio dynamics can be written as

$$X_{n+1}^{u_n} = X_n^{u_{n-1}}(1 + r_f) + u_n(Y_{n+1} - r_f), \tag{4.1}$$

where the sequence $(Y_n)$ comprises i.i.d. non-degenerate random variables. Based on this setup, we obtain the following proposition.

*Proposition* 4.1. *The equilibrium control given by Theorem 2.6 for the mean-semivariance reward functional* (1.5) *in the multi-period setting with the portfolio dynamics* (4.1) *is independent of the portfolio value.*

*Proof.* We have already established the validity of the statement in the one-period setting (see Proposition 3.1); hence it holds for $n = T-1$. To extend the statement to the multi-period scenario we use an induction argument. Consider an arbitrary value for $(n, x)$. Under the induction assumption that the control values from time $n+1$ and onwards, i.e., $u_{n+1}, u_{n+2} \ldots, u_{T-1}$, are constants independent of the portfolio values, we have, with an obvious use of notation,

$$
\begin{aligned}
X_{n+1}^{u_n} =& x(1+r_f) + u_n(Y_{n+1} - r_f), && (4.2)\\
X_{n+2}^{u_n,u_{n+1}} =& X_{n+1}^{u_n}(1+r_f) + u_{n+1}(Y_{n+2} - r_f)\\
=& \big(x(1+r_f) + u_n(Y_{n+1} - r_f)\big)(1+r_f) + u_{n+1}(Y_{n+2} - r_f) && (4.3)\\
X_{n+3}^{u_n,u_{n+1},u_{n+2}} =& X_{n+2}^{u_n,u_{n+1}}(1+r_f) + u_{n+2}(Y_{n+3} - r_f)\\
=& \Big( \big(x(1+r_f) + u_n(Y_{n+1} - r_f)\big)(1+r_f) + u_{n+1}(Y_{n+2} - r_f)\Big)(1+r_f)\\
& + u_{n+2}(Y_{n+3} - r_f) && (4.4)\\
& \vdots\\
X_T^{\boldsymbol{u}} =& x(1+r_f)^{T-n} + u_n(Y_{n+1} - r_f)(1+r_f)^{T-n-1} + u_{n+1}(Y_{n+2} - r_f)(1+r_f)^{T-n-2}\\
& + \cdots + u_{T-1}(Y_T - r_f)\\
=& x(1+r_f)^{T-n} + \sum_{i=0}^{T-n-1} u_{n+i}(Y_{n+i+1} - r_f)(1+r_f)^{T-n-(1+i)}. && (4.5)
\end{aligned}
$$

Hence

$$
\mathbb{E}_{n,x}[X_T^{\boldsymbol{u}}] = x(1+r_f)^{T-n} + (\mathbb{E}[Y] - r_f) \sum_{i=0}^{T-n-1} u_{n+i}(1+r_f)^{T-n-(1+i)}, \tag{4.6}
$$

where $\mathbb{E}[Y]$ denotes the expectation of $Y_i$ and

$$
\begin{aligned}
\mathrm{SV}_{n,x}(X_T^{\boldsymbol{u}}) = \mathbb{E}\Bigg[ \Bigg( & \sum_{i=0}^{T-n-1} u_{n+i}(Y_{n+i+1} - r_f)(1+r_f)^{T-n-(1+i)}\\
& -(\mathbb{E}[Y] - r_f) \sum_{i=0}^{T-n-1} u_{n+i}(1+r_f)^{T-n-(1+i)} \Bigg)^2\\
& I_{\sum_{i=0}^{T-n-1} u_{n+i}(Y_{n+i+1}-r_f)(1+r_f)^{T-n-(1+i)} < (\mathbb{E}[Y]-r_f)\mathbb{E}\left[\sum_{i=0}^{T-n-1} u_{n+i}(1+r_f)^{T-n-(1+i)}\right]} \Bigg].
\end{aligned} \tag{4.7}
$$

Given the linear interactions between $x$ and $u_n$ in (4.6) and the independence of (4.7) from $x$, along with the induction assumption, we find that the optimal $u_n$ for maximizing the mean-semivariance reward (1.5) is independent of $x$. The result thus follows from the case when $n = T-1$, which has already been established. □

*Remark.* While the equilibrium control remains independent of the portfolio value at each distinct point in time, deriving an explicit formula for the mean-semivariance equilibrium control in the multi-period setting presents significantly greater challenges than for the mean-variance scenario. This is due to the inclusion of the indicator function within the semivariance expression, coupled with the inability to eliminate its dependence on the current control, as was possible in the one-period case. Consequently, our game-theoretic framework and the extended Bellman equation presented in Theorem 2.6, along with its numerical implementation, are useful for the analysis and determination of the equilibrium control across multiple periods.

## 4.1 Results based on binomial returns

We will now implement our game-theoretic method for $T = 5$, representative of one week of returns, where an investor pre-commits to a one-week portfolio strategy at the beginning of each week. Once again, we employ a binomial framework. In particular, we assume that the sequence $(Y_n)$ determining the return of the portfolio strategy defined in (4.1) comprises i.i.d. random variables corresponding to

$$Y_i = \begin{cases} r_u \text{ with probability } p \in (0,1), \\ r_d \text{ with probability } 1-p. \end{cases} \tag{4.8}$$

The parameters we consider are as follows. The initial portfolio value is $x = 50$, the possible returns are $r_u = 0.03$ for and $r_d = -0.02$, and the probability of a positive return is $p = 0.7$. Diverging from the one-period model, we explore the risk-free rate $r_f$ at two distinct levels, 0 or 0.001, to evaluate the impact of $r_f$ across multiple steps. This exploration is of interest as the equilibrium control in a mean-variance framework, according to [5], is obtained by discounting (3.14) by the risk-free rate

$$\hat{u} = \frac{\mathbb{E}[Y - r_f]}{\gamma \operatorname{Var}(Y)}(1 + r_f)^{-(T-n-1)}. \tag{4.9}$$

Hence, when the risk-free rate is zero, the mean-variance equilibrium control remains constant. Furthermore, we incorporate an examination of investor behavior under different degrees of risk aversion by exploring two specific levels of the risk aversion coefficient, $\gamma = 2$ and $\gamma = 4$.

Recalling our earlier discussion, constructing a grid of potential portfolio values across different time points is essential for implementing our game-theoretic framework. This grid is determined by the extreme scenarios at each time point, analogous to the one-period setting, and corresponds to the portfolio being fully invested in the risky asset from the outset. The lower bound at each node is derived by assuming the risky asset decreases in value at every step up to that time point. Conversely, the upper bound is determined by assuming an increase in the risky asset's value at every step up to that time point. By incorporating 200 equally spaced portfolio values at each temporal junction, we ensure computational precision and robustness. These values facilitate the calculation of the sequences $(g_n^{\boldsymbol{u}})$ and $(h_n^{\boldsymbol{u}})$ at designated nodes, employing linear interpolation to estimate values between these nodes, thereby reinforcing the robustness of our method's implementation.

*Remark.* Similar to the one-period framework, there is no inherent guarantee that the equilibrium control $u_k$ will be between 0 and $X_k$. However, through meticulous testing and verification, we confirm that, for our chosen parameters, the optimal control indeed always falls within these bounds. Furthermore, extending the range to accommodate a broader spectrum of scenarios, if necessary, is straightforward and does not compromise the integrity of our analysis.

Figure 4 provides a comparative visualization of the equilibrium control strategies when $r_f$ is set to zero and $\gamma = 2$. This figure showcases the equilibrium controls under both mean-semivariance and mean-variance frameworks within our binomial setting, following Theorem 2.6. It is important to note that the horizontal width of the lines depicting the equilibrium controls at various points in time represents the extreme scenarios: one where the investor is fully invested in the risky asset and the asset trends upwards or downwards until that point in time. This visual aid is intended merely for illustration purposes, to underscore the potential range of portfolio values encountered in these extreme conditions.

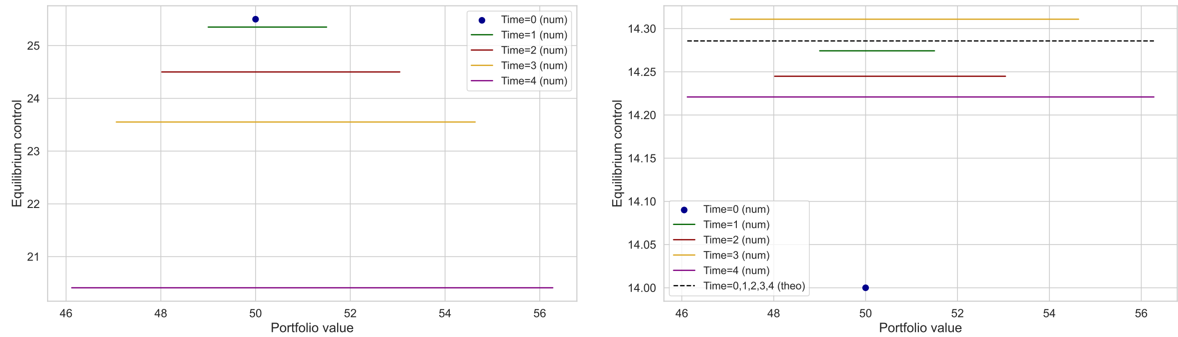


(a) Mean-semivariance equilibrium control

(b) Mean-variance equilibrium control

Figure 4: Comparison of equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0$, and $\gamma = 2$.

An interesting observation from Figure 4 is the changing behavior of the mean-semivariance equilibrium control across different time points. This is in stark contrast to the mean-variance equilibrium control, which remains almost constant. Note that the constancy in the mean-variance control aligns with the theoretical results in (4.9), with only minor deviations observable due to numerical approximations. It is also worth mentioning that, following the discussion in the one-period setting, the mean-semivariance equilibrium control significantly diverges from twice the mean-variance equilibrium control.

Continuing our exploration into the impact of varying the risk-free rate $r_f$ on portfolio strategies, while keeping $\gamma = 2$, we also have results corresponding to the other $r_f$ value, 0.001, as depicted in Figure 5

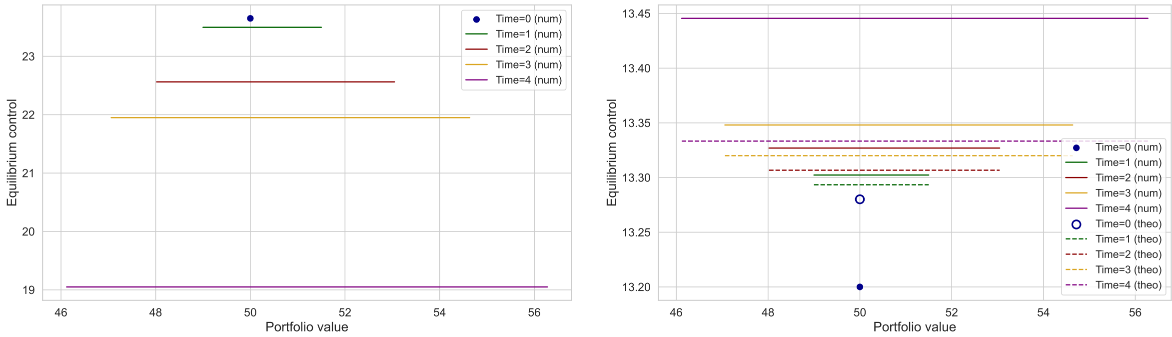


(a) Mean-semivariance equilibrium control

(b) Mean-variance equilibrium control

Figure 5: Comparison of equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 2$.

Upon analyzing the equilibrium control strategies in Figure 5, an interesting pattern emerges, further elucidating the distinctions between mean-semivariance and mean-variance allocations over time. Notably, the mean-semivariance equilibrium control exhibits a decreasing trend as the timeline advances towards the terminal period, suggesting a more conservative investment approach as the end of the period approaches. This behavior contrasts sharply with the mean-variance equilibrium control, which demonstrates an increasing trend under the same conditions, indicative of a strategy that becomes progressively more aggressive, allocating more to the risky asset as time progresses. Such divergence in control strategies highlights a fundamental difference in risk considerations between the two frameworks. Specifically, the cautious approach of the mean-semivariance model, decreasing risk exposure as the investment period nears its end, closely aligns with traditional investment wisdom. This wisdom advocates for adjusting risk exposure dynamically, becoming more conservative as less time remains to recover from potential losses, thereby providing a mathematical foundation to a well-accepted investment principle.

To conclude our analysis of the equilibrium control based on binomial returns, we extend our investigation to consider the scenario where $r_f = 0.001$ and $\gamma = 4$, effectively doubling $\gamma$ from its previous value of 2. Although we do not assume a symmetric distribution, this exploration is compelling because, in a symmetric distribution, semivariance is typically half of the variance. Therefore, it becomes interesting to observe the relationship between the mean-variance control at $\gamma = 2$, as illustrated in Figure 5, and the mean-semivariance control at $\gamma = 4$, as illustrated in Figure 6.

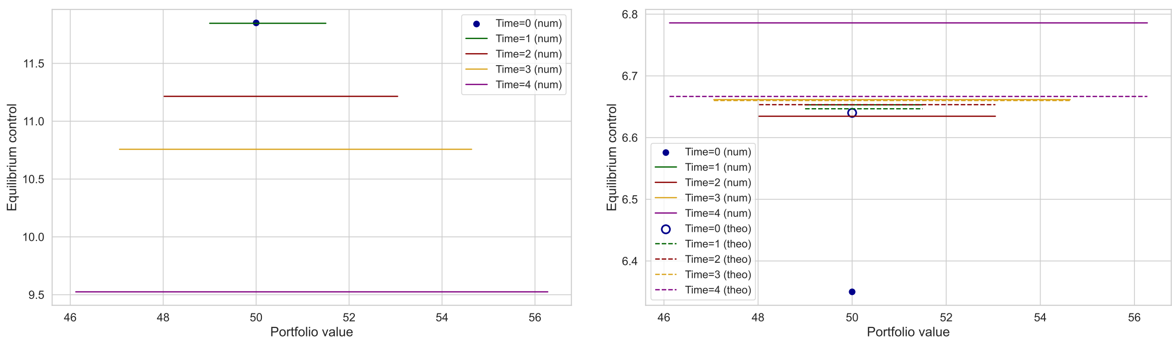


(a) Mean-semivariance equilibrium control

(b) Mean-variance equilibrium control

Figure 6: Comparison of equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 4$.

Upon reflection, we observe the same type of temporal pattern as previously identified for the equilibrium control strategies. Moreover, the adjustment to $\gamma = 4$ for the mean-semivariance equilibrium control does not align with the mean-variance control with $\gamma = 2$.

The differences observed between the numerical and theoretical outcomes for the mean-variance controls, as illustrated in Figures 4-6, underscore the inherent challenges associated with numerical approximations. Notably, the disparity is particularly pronounced at $t = 0$, where our numerical implementation considers only a single portfolio value; which contrasts with subsequent times, for which we compute the equilibrium control across a spectrum of portfolio values. Recognizing that the equilibrium control is, theoretically, independent of the portfolio value at each time point, we derive a more reliable estimate by averaging the control across multiple portfolio values. This methodology significantly narrows the gap between numerical and theoretical outcomes for times beyond the initial point, thereby enhancing the robustness of our numerical approach. The consistency of these refined numerical results with established theoretical results not only underscores the reliability of our proposed method but also lends substantial credence to the mean-semivariance findings, for which theoretical benchmarks are currently unavailable.

### 4.1.1 One week simulation results

Additionally, we have conducted an extensive analysis based on 100,000 simulations of one week's returns to further understand the differences between mean-semivariance and mean-variance equilibrium controls. The outcomes of these simulations are concisely summarized in Table 1. This table presents the average realized values of the reward functionals (1.5) and (1.2) over all simulations, given the corresponding equilibrium control investment strategies.

Table 1: Comparison of portfolio performance based on different equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, and incorporating two distinct levels of $r_f$ and $\gamma$.

| Control type | $r_f$ | $\gamma$ | Average control | Mean-variance reward | Mean-semivariance reward |
|---|---|---|---|---|---|
| Mean-semivariance equilibrium control (num) | 0.0 | 2 | 23.35 | 50.30 | 50.95 |
| Mean-variance equilibrium control (num) | 0.0 | 2 | 14.23 | 50.53 | 50.77 |
| Mean-variance equilibrium control (theo) | 0.0 | 2 | 14.29 | 50.53 | 50.77 |
| Mean-semivariance equilibrium control (num) | 0.001 | 2 | 21.68 | 50.52 | 51.08 |
| Mean-variance equilibrium control (num) | 0.001 | 2 | 13.35 | 50.71 | 50.92 |
| Mean-variance equilibrium control (theo) | 0.001 | 2 | 13.31 | 50.71 | 50.92 |
| Mean-semivariance equilibrium control (num) | 0.001 | 4 | 10.81 | 50.39 | 50.66 |
| Mean-variance equilibrium control (num) | 0.001 | 4 | 6.66 | 50.48 | 50.59 |
| Mean-variance equilibrium control (theo) | 0.001 | 4 | 6.66 | 50.48 | 50.59 |

The performance metrics suggest that both allocation strategies tend to perform better than the other in terms of what they are expected to maximize, effectively embodying their intended risk-return profiles. Notably, the numerical and theoretical mean-variance equilibrium controls show very similar performance, underscoring the precision of our numerical approach. While the distinctions in the trade-offs between mean-semivariance and mean-variance allocations are subtly nuanced, it is important to recognize that the relative differences depend on the initial portfolio value. Additionally, this simulation, based on just one week's data, may not fully capture the potential cumulative effects over time, which could be significant.

Exploring the distribution of terminal portfolio values from our simulations offers further insights into the practical implications of the mean-semivariance and mean-variance equilibrium controls. We focus on equilibrium controls based on numerical results, noting that the numerical mean-variance equilibrium control is practically indistinguishable from the theoretical mean-variance equilibrium control. Figures 7-9 delineate these distributions across our different scenarios.

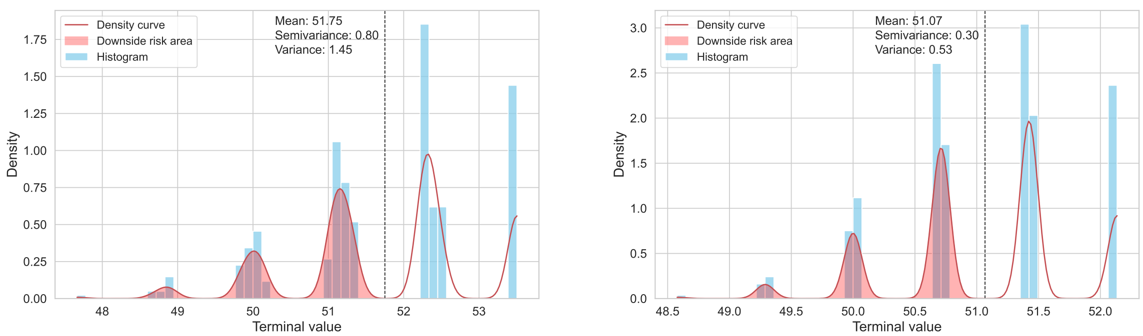


(a) Mean-semivariance equilibrium control (b) Mean-variance equilibrium control

Figure 7: Comparison of terminal portfolio values based on different equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0$, and $\gamma = 2$.

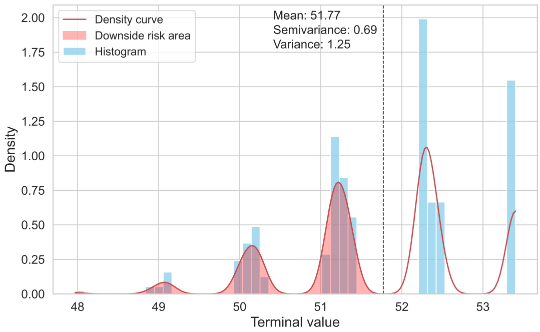


(a) Mean-semivariance equilibrium control

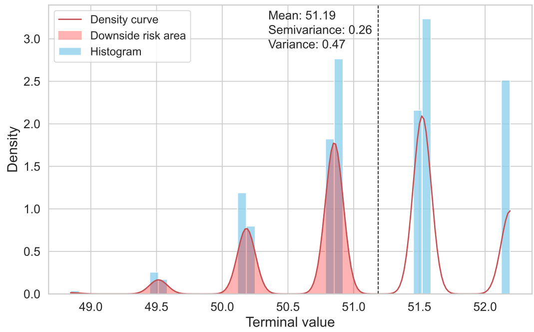


(b) Mean-variance equilibrium control

Figure 8: Comparison of terminal portfolio values based on different equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 2$.

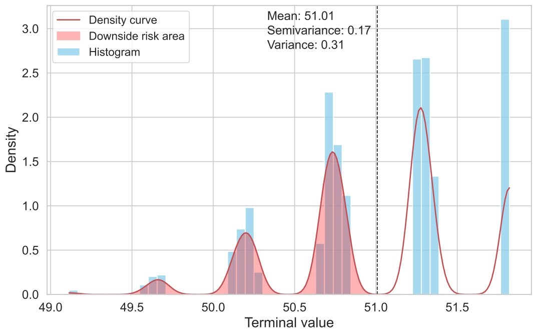


(a) Mean-semivariance equilibrium control

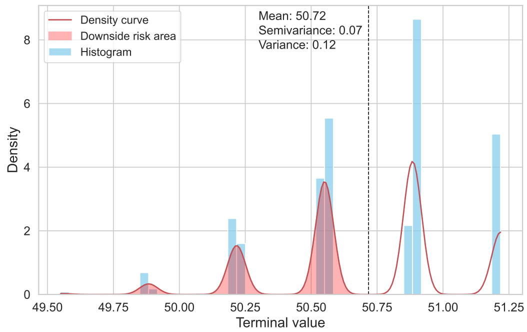


(b) Mean-variance equilibrium control

Figure 9: Comparison of terminal portfolio values based on different equilibrium controls in a multi-period binomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $p = 0.7$, $r_f = 0.001$, and $\gamma = 4$.

The figures presented illustrate how different equilibrium controls balance risk versus return. Specifically, for a given $\gamma$, the mean-semivariance equilibrium control tends to yield a higher expected return but also incurs higher variance and semivariance compared to the mean-variance equilibrium control. This outcome stems from the fact that semivariance, being smaller than variance, places a relatively lower emphasis on risk for the same $\gamma$ value. However, this observation should not be misconstrued as a flaw in the mean-semivariance approach. Rather, it highlights how the impact of $\gamma$ diverges between the mean-semivariance and mean-variance frameworks, reflecting distinct risk-return trade-offs.

## 4.2 Results based on trinomial returns

We now extend our exploration to trinomial returns on $Y$. This advancement introduces the possibility of a third potential outcome in addition to the previously analyzed upward and downward movements. The trinomial model has been successfully applied for derivatives pricing to model stock returns, as evidenced by [6] and [16] among other. In our particular setup, the third outcome represents shocks, thereby adding another dimension to more accurately reflect market dynamics. Specifically, the returns correspond to i.i.d. random variables $Y_i$ given by

$$Y_i = \begin{cases} r_s \text{ with probability } p_s, \\ r_u \text{ with probability } p, \\ r_d \text{ with probability } 1 - p_s - p, \end{cases} \tag{4.10}$$

where $r_s$ represents the return associated with a shock, $p_s$ denotes the corresponding probability, and $r_u$, $r_d$, and $p$ are consistent with their interpretations within the binomial framework.

For numerical analysis, we examine two types of shocks, positive ($r_s = 0.1$) and negative ($r_s = -0.1$), each occurring with a probability $p_s = 0.02$. In alignment with our binomial model, parameters are set to $r_u = 0.03$, $r_d = -0.02$, and the risk-free rate $r_f = 0.001$. The probability of an upward move is adjusted to $p = 0.7 - p_s = 0.68$. Additionally, the risk aversion coefficient $\gamma = 2$ is retained.

Illustrations of the equilibrium control plots under scenarios with positive and negative shocks are presented in Figures 10 and 11, respectively.

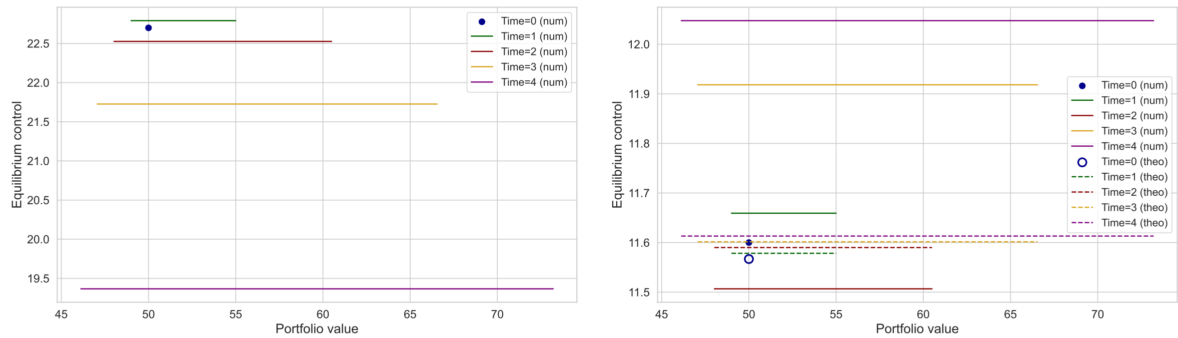


(a) Mean-semivariance equilibrium control

(b) Mean-variance equilibrium control

Figure 10: Comparison of equilibrium controls in a multi-period trinomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $r_s = 0.1$, $p = 0.68$, $p_s = 0.02$, $r_f = 0.001$, and $\gamma = 2$.

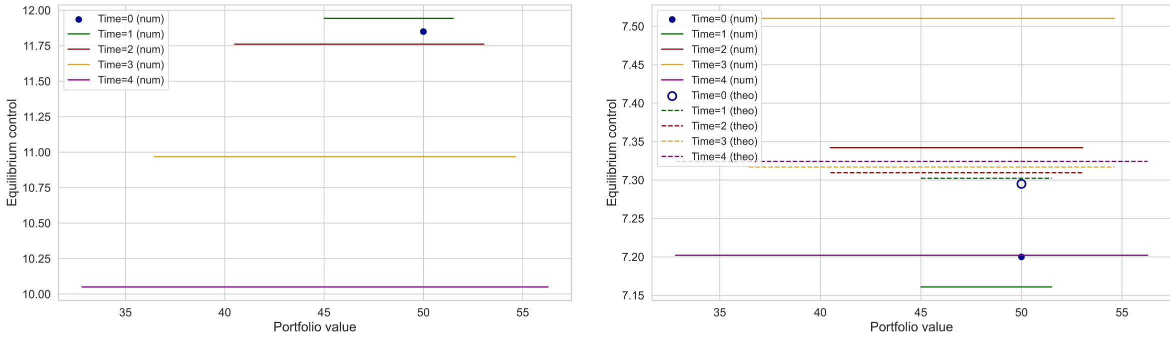


(a) Mean-semivariance equilibrium control  (b) Mean-variance equilibrium control

Figure 11: Comparison of equilibrium controls in a multi-period trinomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $r_s = -0.1$, $p = 0.68$, $p_s = 0.02$ $r_f = 0.001$, and $\gamma = 2$.

Upon reviewing the optimal control plots for scenarios with positive and negative shocks, our numerical method appears to be effective, showing a reasonable alignment with the theoretical equilibrium control in the mean-variance case. Notably, we observe that the equilibrium control is, on average, smaller when introducing negative shocks compared to scenarios with positive shocks, aligning with our expectations. However, we note a somewhat larger difference between numerical and theoretical controls in the mean-variance framework. While these differences might seem significant, they are relatively minor given the scale of the y-axis and considering that we use an optimal control step size of 0.05 during optimization. To further reduce this discrepancy and enhance the accuracy of our numerical results, adjustments such as increasing the number of portfolio values at each time step and decreasing the step size of the optimal control during optimization could be implemented.

### 4.2.1 One week simulation results

Similar to the analysis conducted in the binomial return scenario, we have executed 100,000 simulations to evaluate the performance of various equilibrium controls in terms of their realized mean-semivariance and mean-variance rewards. The outcomes of these simulations are detailed in Table 2.

Table 2: Comparison of portfolio performance based on different equilibrium controls in a multi-period trinomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $p = 0.68$, $p_s = 0.02$, $r_f = 0.001$, $\gamma = 2$ and incorporating two distinct levels of $r_s$.

| Control type | $r_s$ | Average control | Mean-variance reward | Mean-semivariance reward |
|---|---|---|---|---|
| Mean-semivariance equilibrium control (num) | 0.1 | 21.49 | 50.37 | 51.12 |
| Mean-variance equilibrium control (num) | 0.1 | 11.79 | 50.70 | 50.92 |
| Mean-variance equilibrium control (theo) | 0.1 | 11.59 | 50.70 | 50.92 |
| Mean-semivariance equilibrium control (num) | −0.1 | 11.13 | 50.40 | 50.60 |
| Mean-variance equilibrium control (num) | −0.1 | 7.28 | 50.46 | 50.54 |
| Mean-variance equilibrium control (theo) | −0.1 | 7.31 | 50.46 | 50.54 |

Our simulation results indicate that the controls effectively maximize their intended objectives when assessed relative to one another. Moreover, the effectiveness is further evidenced by a significant congruence between the numerical and theoretical mean-variance equilibrium controls across various scenarios, which underscores the accuracy of our numerical methods.

Continuing our analysis, mirroring the approach used for binomial returns, we focus on the simulated distribution of terminal portfolio values. Given the negligible difference between numerical and theoretical mean-variance controls, we limit our presentation to the numerical results. Figures 12 and 13 display these terminal values, highlighting the outcomes of our simulations.

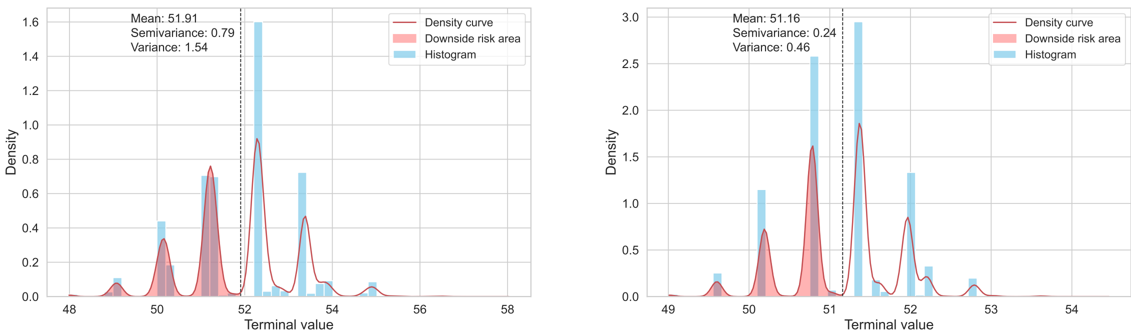


(a) Mean-semivariance equilibrium control (b) Mean-variance equilibrium control

Figure 12: Comparison of terminal portfolio values based on different equilibrium controls in a multi-period trinomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $r_s = 0.1$, $p = 0.68$, $p_s = 0.02$, $r_f = 0.001$, and $\gamma = 2$.

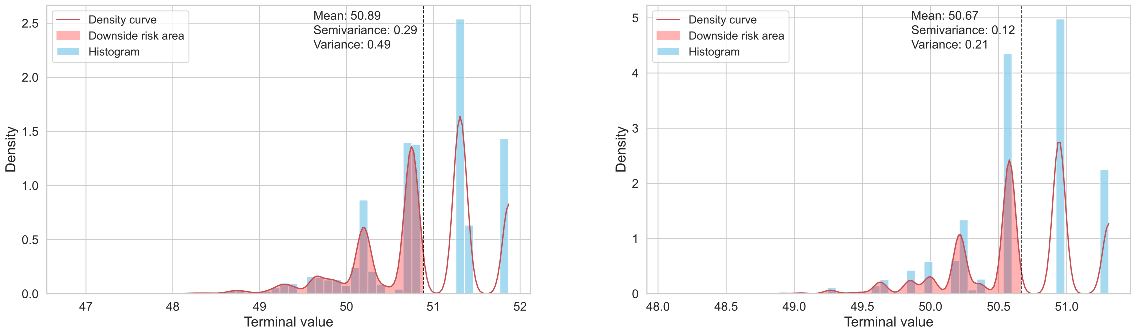


(a) Mean-semivariance equilibrium control (b) Mean-variance equilibrium control

Figure 13: Comparison of terminal portfolio values based on different equilibrium controls in a multi-period trinomial setting, utilizing parameters $T = 5$, $x = 50$, $r_u = 0.03$, $r_d = -0.02$, $r_s = -0.1$, $p = 0.68$, $p_s = 0.02$, $r_f = 0.001$, and $\gamma = 2$.

Upon examining the terminal portfolio values depicted in Figures 12 and 13, we observe a greater spread compared to the analogous binomial setup, attributable to the inclusion of shocks. Despite this increased dispersion, the patterns in terms of mean, variance, and semivariance closely align with those observed in the binomial case.

# 5 Conclusion

This study has proposed a novel recursive scheme for asset allocation in a mean-semivariance framework, catering to risk-averse investors. By extending the Bellman equation, we have facilitated the determination of Nash equilibrium strategies that aim to address time-inconsistency in multi-period financial decision making within a game-theoretic context. The proposed approach has been validated through a comprehensive numerical study, demonstrating robust performance in achieving desired investment outcomes.

Future research could explore several promising avenues. Testing the model on other types of return distributions, both discrete and continuous, could enhance the understanding of its robustness and applicability across different market conditions. Extending the framework to continuous time would also be a significant development. Additionally, applying the proposed model to empirical data could validate its practical relevance and effectiveness in real-world applications.

# 6 Acknowledgments

The authors are grateful to Mariana Khapko at the University of Toronto for suggesting considering semivariance as a measure of risk, and for fruitful discussions.

# References


[1] Suleyman Basak and Georgy Chabakauri. “Dynamic mean-variance asset allocation”. In: *The Review of Financial Studies* 23.8 (2010), pp. 2970–3016.

[2] Richard Bellman. *Dynamic programming*. Princeton University Press, 1957.

[3] Tomasz R. Bielecki, Tao Chen, and Igor Cialenco. “Time-inconsistent Markovian control problems under model uncertainty with application to the mean-variance portfolio selection”. In: *International Journal of Theoretical and Applied Finance* 24.01 (2021), p. 2150003.

[4] Tomas Björk, Mariana Khapko, and Agatha Murgoci. *Time-inconsistent control theory with finance applications*. Springer, 2021.

[5] Tomas Björk and Agatha Murgoci. “A theory of Markovian time-inconsistent stochastic control in discrete time”. In: *Finance and Stochastics* 18.3 (2014), pp. 545–592.

[6] Phelim P. Boyle. “A lattice framework for option pricing with two state variables”. In: *The Journal of Financial and Quantitative Analysis* 23.1 (1988), pp. 1–12.

[7] John C. Cox, Stephen A. Ross, and Mark Rubinstein. “Option pricing: A simplified approach”. In: *Journal of Financial Economics* 7.3 (1979), pp. 229–263.

[8] Javier Estrada. “Mean-semivariance behavior: Downside risk and capital asset pricing”. In: *International Review of Economics & Finance* 16.2 (2007), pp. 169–185.

[9] Javier Estrada. “Mean-semivariance optimization: A heuristic approach”. In: *Journal of Applied Finance* 18.1 (2008).

[10] William W. Hogan and James M. Warren. “Toward the development of an equilibrium capital-market model based on semivariance”. In: *Journal of Financial and Quantitative Analysis* 9.1 (1974), pp. 1–11.

[11] Hanqing Jin, Harry M. Markowitz, and Xun Yu Zhou. “A note on semivariance”. In: *Mathematical Finance* 16.1 (2006), pp. 53–61.

[12] Kristoffer Lindensjö. “A regular equilibrium solves the extended HJB system”. In: *Operations Research Letters* 47.5 (2019), pp. 427–432.

[13] Harry M. Markowitz. *Portfolio selection: Efficient diversification of investments*. Yale University Press, 1959.

[14] Bezalel Peleg and Menahem E. Yaari. “On the existence of a consistent course of action when tastes are changing”. In: *The Review of Economic Studies* 40.3 (1973), pp. 391–401.

[15] Robert H. Strotz. “Myopia and inconsistency in dynamic utility maximization”. In: *The Review of Economic Studies* 23.3 (1955), pp. 165–180.

[16] Fei Lung Yuen and Hailiang Yang. “Option pricing with regime switching by trinomial tree method”. In: *Journal of Computational and Applied Mathematics* 233.8 (2010), pp. 1821–1833.