

Short-Term Portfolio Construction: Separating Forecasting from Portfolio Design

Nikolaos Halidias

Actuarial-Financial Mathematics Lab

Department of Statistics and Actuarial - Financial Mathematics

March 13, 2026

Long-Term vs Short-Term Investment Decisions

- Portfolio construction depends strongly on the investment horizon.
- For **long-term investments**, the analysis should focus on structural factors such as:
 - the financial robustness of firms,
 - the quality of management,
 - business fundamentals,
 - competitive advantages,
 - long-run growth prospects.
- For **short-term investments**, market behavior is often driven by fast and unstable information flows, including:
 - fake news and misinformation,
 - sudden geopolitical developments,
 - macroeconomic announcements,
 - market sentiment and speculative reactions.
- Therefore, portfolio construction is not a one-size-fits-all problem:

investment horizon \implies different risk factors, signals, and strategies.

Focus of This Talk: Short-Term Portfolio Construction

- In this talk, we focus on **short-term investments** over a finite time horizon

$[0, T]$.

- In such a setting, portfolio construction is strongly affected by rapidly changing market information and short-lived opportunities.
- Two ingredients will play a central role:
 - a **predictive mechanism**, which provides information or signals about the future evolution of market variables,
 - **options**, which offer flexibility for hedging, leverage, and the design of structured payoff profiles.
- Therefore, our main objective is to combine prediction and derivative instruments in order to construct portfolios adapted to short-term market dynamics.

One Underlying and Its Option Market

We begin with the simplest nontrivial setting:

one stock S and the corresponding options with maturity T .

At time $t = 0$, the investor may combine:

- a position in the stock,
- a cash position,
- call options at selected strikes,
- put options at selected strikes.

If $x = S_T$, the terminal profit of the portfolio can be written in the form

$$\Pi(x) = b - V + ax + \sum_{i=1}^m \gamma_i (x - K_i^C)^+ + \sum_{j=1}^n \delta_j (K_j^P - x)^+.$$

Thus, the problem is not to choose a stock in isolation, but to design a **profit function** $\Pi(x)$ adapted to our market view.

Forecast First, Construction Second

Our philosophy is:

forecasting \neq portfolio construction.

We first produce a forecast of the form

$$S_T \in A,$$

where $A \subseteq \mathbb{R}_+$ is a predicted set of plausible terminal prices.

Then we construct a portfolio whose payoff profile reflects this view.

In other words:

- the forecasting mechanism provides information,
- the portfolio construction mechanism converts this information into a tradable decision.

This separation is essential in short-term investment decisions, where forecasts may come from highly heterogeneous information sources.

Portfolio Design Under a Prediction Set

Given a forecast set A and a loss tolerance $D > 0$, we want to construct a portfolio such that:

$$\begin{aligned}\Pi(x) &\geq \kappa, & x \in A, \\ \Pi(x) &\geq -D, & x \geq 0,\end{aligned}$$

for the largest possible value of κ .

Interpretation:

- on the **predicted scenario** A , the portfolio should be profitable,
- on the **bad scenarios**, losses should never exceed the prescribed amount D .

Hence the investor does not ask for a universally optimal portfolio.

The investor asks for a portfolio that is:

aggressive where the forecast supports it, and controlled elsewhere.

Classical Option Strategies are Included as Special Cases I

An important point is that this framework does not exclude the classical named option strategies.

On the contrary, since the portfolio is constructed from:

- cash,
- the underlying asset,
- call options at different strikes,
- put options at different strikes,

all standard payoff structures are automatically contained as special cases.

For example, by choosing appropriate positions, one recovers:

- bull and bear spreads,
- straddles and strangles,
- covered calls and protective puts,

Classical Option Strategies are Included as Special Cases II

- butterfly spreads,
- and other standard structured payoffs.

So the present approach is not a competitor to the classical strategies.

It is a general construction framework that contains them all.

Why This is More Flexible than Markowitz-Type Approaches

This framework is substantially more flexible than classical Markowitz-style portfolio selection.

In Markowitz theory, the investor typically chooses weights in the underlying assets using estimated means, variances, and covariances.

Here, instead:

- the investor may use **any forecasting technology**,
- the forecast may come from
 - technical analysis,
 - market intuition,
 - news and sentiment analysis,
 - machine learning,
 - macroeconomic or geopolitical views,
 - stress scenarios.
- this information is summarized by the set A ,
- and the portfolio is then obtained by optimization.

So the methodology is modular:

The key idea is simple but powerful:

Do not force forecasting and portfolio choice into one fragile model.

Instead:

- 1 form a prediction set A for the terminal stock price,
- 2 impose a loss budget D ,
- 3 use the stock and the option market to engineer a payoff function Π .

This leads to portfolios that are:

- scenario-driven,
- risk-controlled,
- and much closer to actual financial engineering practice.

Bid–Ask Spread, Transaction Costs, and Real Execution

A practical advantage of this framework is that market frictions can be incorporated very naturally.

In the static one-period setting, the terminal payoff is unchanged; what changes is the **initial execution cost**.

For each instrument:

- long positions are executed at the **ask** price,
- short positions are executed at the **bid** price.

Hence, transaction costs can be incorporated through the bid–ask spread:

$$V_0 = \sum_i q_i^{\text{buy}} p_i^{\text{ask}} - \sum_i q_i^{\text{sell}} p_i^{\text{bid}}, \quad q_i^{\text{buy}}, q_i^{\text{sell}} \geq 0.$$

Therefore, the optimization is based not on idealized mid-prices, but on **real executable market prices**.

This makes the portfolio construction problem immediately more realistic and closer to actual trading practice.

All-or-None Execution and Order Limits I

Another practical issue is that the portfolio should be implemented **exactly as designed**.

In practice, this may require an **all-or-none** instruction:

either the entire package of trades is executed, or none of it is executed.

So the investor does not want partial implementation of the strategy, because partial execution may destroy the intended hedge or payoff profile.

Moreover, trading platforms or execution rules often impose an upper bound on the number of simultaneous orders. For example, we may require

at most 8 orders at the same time.

This leads naturally to binary variables:

- a variable indicating whether order i is activated,

- constraints enforcing all-or-none logic,
- and a cardinality constraint such as

$$\sum_{i=1}^N z_i \leq 8, \quad z_i \in \{0, 1\}.$$

Therefore, the portfolio selection problem becomes not only a pricing or hedging problem, but also an **execution-feasible optimization problem**.

Diversification Beyond the Number of Stocks I

A key message is that diversification is *not* only cross-sectional.

Even with a single underlying stock, the investor can enlarge the range of possible risk profiles by combining:

- the stock itself,
- call options at different strikes,
- put options at different strikes,
- and cash.

Thus, diversification is achieved not only across *assets*, but also across *payoff shapes*.

one stock + option menu \implies richer exposure structure

This allows us to design portfolios with:

Diversification Beyond the Number of Stocks II

- directional exposure when our forecast is strong,
- downside protection when we want bounded losses,
- convex upside when we expect large moves,
- and intermediate profiles that are impossible in stock-only investing.

From One Asset to Many Assets I

We now move from the one-asset case to a portfolio with **many underlying assets**.

The generalization is conceptually similar:

- before, the terminal state was described by one variable S_T ,
- now, the terminal state is described by the vector

$$S_T = (S_1^T, \dots, S_d^T) \in \mathbb{R}_+^d.$$

Thus, instead of designing a profit function of one variable, we now design a profit function on a multi-dimensional state space:

$$\Pi : \mathbb{R}_+^d \rightarrow \mathbb{R}.$$

So the philosophy remains unchanged:

forecast first, construction second.

Forecast as a Prediction Set in \mathbb{R}_+^d

In the multi-asset case, the investor may produce a forecast of the form

$$(S_1^T, \dots, S_d^T) \in G, \quad G \subseteq \mathbb{R}_+^d.$$

Here, G represents the set of terminal scenarios that the investor considers plausible.

This is very flexible, because G can encode:

- bullish or bearish views on specific assets,
- relative-value views,
- sector or macro scenarios,
- dependence beliefs between assets,
- information coming from statistics, news, machine learning, or intuition.

So the forecast is not forced into a single parametric probabilistic model.

Instead, the investor specifies a **region of belief** in the state space.

Construction of the Multi-Asset Portfolio I

At time 0, the investor constructs a static portfolio consisting of:

- cash,
- positions in the d underlyings,
- call options on each asset,
- put options on each asset.

Hence, the terminal profit takes the form

$$\Pi(x_1, \dots, x_d) = b - V + \sum_{i=1}^d \phi_i(x_i),$$

where each ϕ_i is built from the corresponding stock and its options.

Therefore, the objective is again:

Construction of the Multi-Asset Portfolio II

- to make the portfolio profitable on the prediction set G ,
- while controlling the loss outside G .

So the one-asset problem is extended naturally from intervals in \mathbb{R}_+ to regions in \mathbb{R}_+^d .

Design Objective in the Multi-Asset Case I

The investor now seeks a portfolio such that:

$$\Pi(x) \geq 0 \quad \text{for } x \in G,$$

while at the same time

$$\Pi(x) \geq -D \quad \text{for all } x \in \mathbb{R}_+^d.$$

Interpretation:

- on the scenarios that the investor believes in, the portfolio should be profitable,
- on all other scenarios, the maximum loss should remain controlled.

This is exactly the same engineering principle as before, but now in higher dimension.

The difficulty is greater, but so is the expressive power:

multi-asset portfolios allow us to encode much richer market views.

Markowitz Theory with Options

The classical Markowitz framework can be extended naturally once options are included in the investment universe.

Indeed, instead of selecting positions only in the underlying assets, we may choose positions in:

- cash,
- the underlying assets,
- call and put options.

If all these instruments are viewed as random payoffs (or returns) at time T , then the portfolio becomes

$$R_\theta = \theta^\top Y,$$

where Y is the vector of payoffs of all available instruments and θ is the vector of positions.

Therefore one may formulate a mean–variance problem such as

$$\min_{\theta} \text{Var}(R_\theta),$$

subject to budget, expected return, and position constraints

Why Go Beyond Mean–Variance?

Although the Markowitz problem with options is perfectly legitimate, it does not fully capture the portfolio design philosophy developed here.

In our framework, the investor may wish to incorporate:

- a prediction set $G \subseteq \mathbb{R}_+^d$,
- profitability on the scenarios in G ,
- explicit lower bounds on losses,
- transaction costs and execution constraints.

Thus, minimizing variance is only one possible criterion.

A construction-based optimization problem is more flexible, because it allows the investor to shape the payoff directly according to the forecast and the market constraints, instead of summarizing risk only through variance.

Hence:

Markowitz with options is possible, but forecast-driven construction is richer.

- [1] H. Markowitz,
“Portfolio Selection,”
The Journal of Finance, 7(1), 77–91, 1952.
DOI: 10.1111/j.1540-6261.1952.tb01525.x.
- [2] N. Halidias,
Financial Engineering with Python,
forthcoming, World Scientific Publishing.
Draft available on ResearchGate:
https://www.researchgate.net/publication/394413705_Financial_Engineering_with_Python.
- [3] N. Halidias,
Financial-Engineering,
GitHub repository with Python code accompanying *Financial Engineering with Python*.
Available at:
<https://github.com/nikoshalidias/Financial-Engineering>.
Includes, among others, the notebooks PCUP.ipynb (one-asset portfolio construction) and Multi_Asset4.ipynb (two-asset construction with bid–ask spread).