

Financial Engineering: The Big Picture

Nikolaos Halidias

Actuarial-Financial Mathematics Lab

Department of Statistics and Actuarial - Financial Mathematics

May 15, 2026

Deterministic Arbitrage: Market Setup

We consider a one-period market with terminal stock price

$$x = S_T \in [0, \infty).$$

The traded instruments are:

$$1, \quad S, \quad (S_T - K_i)^+, \quad (K_i - S_T)^+, \quad i = 1, \dots, n.$$

A static portfolio is

$$\theta = (a, b, \gamma_1, \dots, \gamma_n, \delta_1, \dots, \delta_n),$$

where

a : stock position, b : bank account position,

γ_i : call position, δ_i : put position.

Portfolio Payoff and Initial Cost

The terminal payoff of the portfolio is

$$G_{\theta}(x) = ax + be^{rT} + \sum_{i=1}^n \gamma_i (x - K_i)^+ + \sum_{i=1}^n \delta_i (K_i - x)^+.$$

The initial setup cost is

$$Y_{\theta} = aS_0 + b + \sum_{i=1}^n (\gamma_i C_i + \delta_i P_i).$$

The financed terminal profit is

$$\Pi_{\theta}(x) = G_{\theta}(x) - Y_{\theta}e^{rT}.$$

Definition: Deterministic Arbitrage

A static portfolio θ is a **deterministic arbitrage** if

$$\Pi_{\theta}(x) \geq 0 \quad \text{for all } x \geq 0,$$

and

$$\Pi_{\theta}(x) > 0 \quad \text{for at least one } x \geq 0.$$

Equivalently, for zero-cost portfolios,

$$Y_{\theta} = 0,$$

we require

$$G_{\theta}(x) \geq 0 \quad \text{for all } x \geq 0,$$

and

$$G_{\theta}(x) > 0 \quad \text{for some } x \geq 0.$$

This is a **model-free** and **pathwise** notion of arbitrage.

Deterministic vs Model-Internal Arbitrage

Deterministic arbitrage

A deterministic arbitrage is a real static trading certificate. Using observable market prices, its financed terminal profit satisfies

$$\Pi_{\theta}(x) \geq 0 \quad \text{for every } x \geq 0,$$

and is strictly positive for at least one terminal price. Thus the statement is model-free, pathwise, and directly testable from quoted prices.

Black–Scholes and binomial arbitrage

In Black–Scholes or in the binomial model, arbitrage is defined inside an idealized model. Its absence is mainly a mathematical consistency condition that gives pricing equations or risk-neutral probabilities. The alleged arbitrage usually relies on assumptions such as continuous rebalancing, frictionless trading, exact model dynamics, or a finite prescribed price tree.

What should be clear to students

Model-internal arbitrages should not be presented as riskless profits that can be extracted in actual markets. In practice they usually cannot be implemented as guaranteed gains, because the assumptions needed for the argument are not available in the real trading environment.

Correct intuition

Deterministic arbitrage is an observable inconsistency of traded prices. By contrast, Black–Scholes and binomial no-arbitrage are consistency principles inside a chosen model. Confusing these two notions gives students the wrong intuition about what option-pricing theory can and cannot deliver in practice.

Why the Problem Becomes Finite

The payoff $G_\theta(x)$ is continuous and piecewise affine.

Its only possible breakpoints are

$$0, K_1, \dots, K_n.$$

Therefore, the infinite constraint

$$G_\theta(x) \geq 0 \quad \text{for all } x \geq 0$$

is equivalent to finitely many inequalities:

$$G_\theta(0) \geq 0, \quad G_\theta(K_i) \geq 0, \quad i = 1, \dots, n,$$

together with the tail-slope condition

$$a + \sum_{i=1}^n \gamma_i \geq 0.$$

Linear Program for Detecting Arbitrage I

We search for a zero-cost portfolio with nonnegative payoff and positive gain at some certificate point.

$$\begin{aligned} \max \quad & D_0 + \sum_{i=1}^n D_i + \eta M \\ \text{s.t.} \quad & aS_0 + b + \sum_{i=1}^n (\gamma_i C_i + \delta_i P_i) = 0, \\ & G_\theta(0) \geq D_0, \\ & G_\theta(K_i) \geq D_i, \quad i = 1, \dots, n, \\ & a + \sum_{i=1}^n \gamma_i \geq M, \\ & D_0, D_i, M \geq 0, \\ & |a| \leq N_0, \quad |\gamma_i| \leq N_i, \quad |\delta_i| \leq N_i. \end{aligned}$$

Linear Program for Detecting Arbitrage II

If the optimal value is strictly positive, the optimizer gives a deterministic arbitrage. For this topic see [5] and [2].

Jupyter notebook

Download the notebook for the arbitrage detection example.

▶ [Detecting Arbitrage](#)

Multiperiod Arbitrage

The deterministic arbitrage framework follows the viewpoint of the paper *Deterministic Arbitrage, Localized Arbitrage Portfolios, and Arbitrage-Consistent Projection* and the references therein. The key idea is that arbitrage is treated pathwise and model-free: no probability model for S_T is required. Instead, absence of arbitrage is imposed directly through payoff inequalities holding for every terminal stock price $x \geq 0$. A natural extension is to consider a multiperiod static portfolio, where positions are fixed at time 0, while the no-arbitrage condition is imposed pathwise over all admissible price paths across the whole time grid.

Role of Deterministic Arbitrage

Deterministic arbitrage should be interpreted as a market-consistency test on observable prices, not as a Black–Scholes-style replication argument. It leads to strict model-free restrictions on option prices, such as bounds, monotonicity, convexity, and put–call parity.

The key point is practical: if such a static certificate exists, the quoted prices are inconsistent in a pathwise sense. By contrast, arbitrage in Black–Scholes or binomial models is internal to the modelling assumptions and should not be confused with an implementable real-world profit opportunity. This distinction is also the basis for the Arbitrage-Consistent Projection (ACP) method (see [4]), which corrects predicted option prices so that they become arbitrage-consistent.

Separation of Forecasting and Construction

For short-term investment, the forecasting problem should be separated from the portfolio-construction problem.

Long-term investment depends on fundamental characteristics of the company, such as earnings, growth, management, debt, and strategic positioning. In contrast, short-term investment is driven by different signals: direction, volatility, liquidity, market regime, and risk tolerance.

In this short-term framework, the portfolio should not be restricted to the underlying asset. It should also include call and put options, so that we can construct nonlinear payoff profiles adapted to the forecast.

Thus, forecasting produces a view about the future, while portfolio construction transforms this view into an implementable payoff function. This separation allows us to use advanced forecasting tools, including machine learning methods, without changing the construction mechanism.

Prediction Set and Optimal Portfolio Construction I

Assume that the investor wants to allocate an initial budget Y across d stocks and their traded options.

The forecasting step produces a prediction set

$$G \subset \mathbb{R}_+^d,$$

where the terminal vector of stock prices

$$S_T = (S_T^1, \dots, S_T^d)$$

is expected to lie.

The construction step then solves a linear program in order to build a static portfolio whose profit function $\Pi(x)$ satisfies

$$\Pi(x) \geq 0, \quad x \in G,$$

while the loss in the adverse scenario is controlled by

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

Thus, the forecast defines the favourable region, while linear programming constructs the optimal option-equipped portfolio under a prescribed worst-case loss bound. For more on this topic see [2].

Markowitz as a special case

This viewpoint contains the Markowitz model as a special case. Suppose the forecasting module is a mean–variance mechanism producing expected returns $\mu \in \mathbb{R}^d$ and covariance matrix Σ , and suppose that the admissible portfolio contains only the d underlying assets, with no options. If w denotes portfolio weights, the construction step becomes the classical Markowitz problem, for example

$$\min_{w \in \mathbb{R}^d} w^\top \Sigma w \quad \text{s.t.} \quad \mu^\top w \geq m, \quad \mathbf{1}^\top w = 1,$$

or, equivalently, the maximization of expected return penalized by variance.

Prediction Set and Optimal Portfolio Construction IV

Generalization beyond Markowitz

The proposed framework keeps the same separation between forecasting and construction, but removes two restrictions. First, the forecasting module need not be mean–variance: it may be based on scenarios, prediction sets, distributional forecasts, regime models, or machine learning. Second, the admissible portfolio need not contain only the underlying assets: calls and puts can be included, allowing nonlinear payoff profiles adapted to the forecast.

Python notebooks

Download the Jupyter notebooks for the portfolio examples:

▶ [One-asset portfolio](#)

▶ [Two-asset portfolio](#)

Option Pricing and Hedging: Classical Limitations I

Classical option pricing starts from an idealized replication principle.

In the Black–Scholes framework, exact replication requires continuous trading and instantaneous rebalancing of the hedging portfolio. In real markets, however, trading is discrete and affected by liquidity, transaction costs, bid–ask spreads, and operational constraints.

Binomial and trinomial models avoid continuous time, but they impose a different idealization: at each step the stock price is restricted to a finite number of possible values.

Key point

Classical models may be arbitrage-free internally, but this does **not** mean that their prices are **deterministic-arbitrage-free** in the actual market.

Therefore, model-based prices **are not** necessarily market-consistent: they may violate the pathwise no-arbitrage restrictions encoded in real market prices.

Option Pricing and Hedging: Classical Limitations II

Therefore, option pricing and hedging should move toward realistic methods based on feasible portfolios, explicit constraints, and pathwise performance across historical, simulated, or stressed scenarios. For more on this see [2].

Non-Path Dependent Options via Static Portfolios I

Consider a non-path dependent option with payoff

$$f(S_T),$$

where f is piecewise linear with finitely many branches.

Using static portfolios in the stock, the bank account, and traded calls and puts, we can compute the arbitrage-free interval

$$[Y^{\text{buyer}}, Y^{\text{writer}}].$$

Inside this interval, a natural model-free fair price is the value Y^{D^*} that balances the worst possible losses of the two parties:

$$D_{\text{writer}}(Y^{D^*}) = D_{\text{buyer}}(Y^{D^*}).$$

Non-Path Dependent Options via Static Portfolios II

Key point

The proposed fair price is also deterministic-arbitrage-free.

Hence, it is not merely a model output; it is a market-consistent reference value that respects the pathwise no-arbitrage restrictions of the actual market.

After the contract is priced (the actual price will not be necessarily fair in some particular sense or even deterministic arbitrage free), an investor may include it in a larger static portfolio and solve a linear program so that the portfolio is profitable on a forecast set A , while the loss in the adverse scenario is bounded by

$$\Pi(x) \geq -D.$$

Thus, option pricing and hedging are directly connected with the previous portfolio-construction problems under prediction sets. See the book [2].

Non-Path Dependent Options via Static Portfolios III

Python code

Notebook for the option pricing and hedging example:

▶ [Option Pricing and Hedging](#)

Path-Dependent Options: Feasible Dynamic Hedging I

For path-dependent options, static portfolios are generally not sufficient. Instead, we work with feasible dynamic strategies on a discrete time grid.

The forecasting step now produces possible future paths

$$(S_{t_0}, S_{t_1}, \dots, S_{t_N}),$$

coming from historical data, simulated models, stress scenarios, or extreme hypothetical trajectories.

Given this scenario set, we construct a self-financing dynamic strategy according to a chosen risk criterion, such as CVaR or mean-CVaR.

Market-consistent fair price

Within this framework, a fair price notion is obtained by balancing the residual downside risks of the writer and the buyer.

Under suitable assumptions, this fair price lies in the deterministic arbitrage-free interval. Therefore, it is not merely a model output, but a model-free negotiation benchmark and a market-consistent reference value respecting the pathwise no-arbitrage restrictions of the actual market.

The actual price will not necessarily be any kind of fair or even arbitrage free. Given this price the investor will choose the best hedging strategy.

For more on this see [2].

Predictive Option Values inside Dynamic Portfolios

If traded options are included in a dynamic portfolio, then we must be able to assign pathwise values to the calls and puts at each rebalancing date.

In this context, Black–Scholes can be useful as a predictive valuation map for future prices of call and put options. Unfortunately, this is the only possible usage of these kind of models.

More flexible predictive maps may be constructed from data, for example using machine learning methods, in order to estimate future call and put prices along simulated, historical, or stressed paths.

If the predicted option prices violate no-arbitrage restrictions, they can be post-processed through an ACP-type correction so that the resulting call–put system becomes arbitrage-consistent. For more on this see [1] and [4].

This methodology is quite advanced and requires modern forecasting tools.

Moreover, the student should also study effective machine-learning methods for constructing predictive valuation maps.

A natural question is whether it is worthwhile to devote an entire semester to Black–Scholes, if its practical role is mainly limited to serving as an input inside a more advanced framework, where better data-driven tools may also be used.

Conclusion: Option Pricing and Hedging I

The investor needs a reference price before negotiation.

This price should be derived from feasible hedging strategies that can actually be implemented in the market.

The investor must first choose the admissible hedging class:

static portfolio or discrete-time dynamic strategy.

Conclusion: Option Pricing and Hedging II

A common confusion

Predicting the future value of an option is not the same problem as option pricing and hedging. Prediction aims to forecast a future market price. It should therefore rely on a predictive valuation map using all relevant information: asset dynamics, volatility, liquidity, macroeconomic variables, news analytics, sentiment, and other market signals.

Option pricing and hedging ask a different question: how much initial capital is needed to construct a feasible hedge that performs satisfactorily under the chosen risk criterion and market constraints.

These two problems are often conflated in the literature, but they are conceptually different.

Conclusion: Option Pricing and Hedging III

Ask yourself before doing anything

- ▶ Are you trying to compute a good price for an option from your own point of view? If so, that price should be justified by a feasible hedging strategy of your choice.
- ▶ Are you trying to propose a hedging strategy for an option whose price is already given? If so, the strategy should be feasible; even better, it should be the best among all feasible strategies.
- ▶ Are you trying to predict the future value of an option? If so, you should carefully choose the parameters that determine its price. Remember that predicting the price of an exotic option is usually not the goal; price prediction is mainly useful for standard call and put options.
- ▶ In any case, do not confuse these notions.

Conclusion: Option Pricing and Hedging IV

A tongue-in-cheek remark. In light of the discussion above, it is not entirely clear what the practical significance is of pricing techniques for exotic options that produce neither feasible hedging strategies nor efficient price-prediction methods. Perhaps their main contribution is the mathematical pleasure they provide.

One conceptual issue

The problem is not that classical arbitrage concepts are mathematically wrong. The problem is that they are often taught without clearly separating model-internal arbitrage from executable real-market arbitrage.

This can create a misleading intuition: students may come to believe that instantaneous hedging, continuous rebalancing, or finitely many future asset prices are realistic features of actual markets.

Another common confusion

Option pricing and hedging are often confused with the prediction of future option prices.

Prediction is a different problem: it aims to forecast a future market value and should therefore use a predictive valuation map incorporating relevant market information, such as volatility, liquidity, macroeconomic variables, news analytics, sentiment, and other market signals.

By contrast, option pricing and hedging ask how much initial capital is required to implement a satisfactory hedge under a chosen risk criterion and explicit market constraints.

We also tend to confine portfolio construction to the Markowitz mean–variance paradigm. Markowitz theory is important, but it is only a special case.

The alternative

The proposed viewpoint generalizes this framework: it allows arbitrary forecasting mechanisms, realistic constraints, and portfolios that may include both the underlying assets and options.

References

- [1] S. N. Cohen, C. Reisinger, and S. Wang. *Detecting and repairing arbitrage in traded option prices*. Applied Mathematical Finance, 27(5):345–373, 2020.
- [2] N. Halidias. *Financial Engineering with Python*. World Scientific Publishing Company (WSPC), forthcoming.
- [3] N. Halidias. *Feasible Dynamic Hedging and Data-Driven Valuation of Options in Discrete Time*. Available on ResearchGate, 2026.
- [4] N. Halidias. *Deterministic Arbitrage, Localized Arbitrage Portfolios, and Arbitrage-Consistent Projection*. Available on ResearchGate, 2026.
- [5] S. Herzel. *Arbitrage opportunities on derivatives: a linear programming approach*. Dynamics of Continuous, Discrete and Impulsive Systems. Series B, 12(4):589–606, 2005.