# Model-Free Pricing and Dynamic Trading

Nikolaos Halidias

Actuarial-Financial Mathematics Lab
Department of Statistics and Actuarial - Financial Mathematics

February 9, 2026

# Regime switches: past data is a poor guide when the process changes

*"The greatest danger in times of turbulence is not the turbulence—it is to act with yesterday's logic."*
— Peter F. Drucker

*"Risk comes from not knowing what you are doing."*
— Warren Buffett

Upward drift → switch at $t$ → downward drift

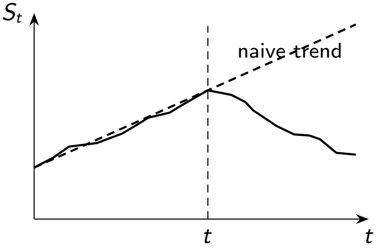


Downward drift → switch at $t$ → upward drift

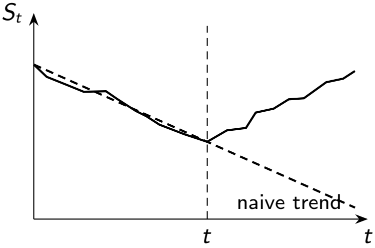


**Takeaway:** Blind extrapolation of past trends can fail after a regime shift (drift change + Brownian noise).

# Price formation: supply, demand, and the crowd

- A stock price is not an *intrinsic* number: it is the *clearing price* of the next trade.
- News, narratives, and constraints shift investors' beliefs & risk budgets ⇒ reallocations across assets ⇒ changing order flow.
- Therefore, transaction prices are shaped by **supply/demand**, **liquidity**, and bargaining pressure (who must trade, and how urgently).

**Key question:** How can we predict the movement of the masses?

- Predicting the "crowd" ≈ forecasting future **net order imbalance** and liquidity.
- Deterministic prediction is impossible; at best we build **probabilistic scenarios** and manage risk (position sizing / hedging) accordingly.

# Motto: forecasting vs. strategy design

*For now, we set aside the forecasting problem and focus on the construction problem:* given an input (forecast or data), how do we build an implementable, risk-controlled decision?

**Give me a forecast,** and I will design the **optimal static portfolio**.

**Give me data,** and I will design the **optimal dynamic strategy**.

- This may sound obvious — but it is not.
- Turning a forecast or a dataset into a *tradable* decision requires **careful calculations**, **risk constraints**, and solving **optimization problems**.
- This is exactly what we will learn next.

# Forecasts Differ — Portfolio Design Must Be Systematic

- Investors will naturally rely on *their own* forecasting style (fundamental, technical, macro, machine learning, ...).
- But once a forecast (or set of scenarios) is specified, everyone faces the *same* challenge: How do we convert information into a tradable portfolio with explicit risk control?
- In this course we use a *shared portfolio-design framework*: forecasting supplies the inputs; portfolio construction is posed as an optimization/hedging problem.

## Key principle

We decouple **prediction** from **strategy design**: forecasts may differ across investors, but the construction step must be mathematically well-posed and operationally implementable—and, crucially, it is a unique technique.

# What You Must Learn First (Before Forecasting)

1. **Master the mathematical toolkit** behind portfolio design: no-arbitrage logic, hedging/replication ideas, constraints, and optimization-based constructions.
2. **Then study forecasting methods** as modular inputs (signals, scenarios, distributions, stress tests).

   *"But these are well known among experienced investors."*

**Our answer**

Maybe. Now we place this "folk knowledge" into a clear mathematical framework — so we can **teach it**, **justify it**, and **apply it consistently** with our students.

# Will you show backtests to demonstrate that the approach works?

- **Not in the standard sense.**

- What we presented is **not a forecasting device** and **not a predictive model**.
- It is a **pricing and hedging framework** grounded in **no-arbitrage** and **explicit risk control** (e.g., bounds, super-/sub-replication, CVaR-type criteria).

- **Backtesting becomes indispensable** when the claim is **predictive power** (signals, return forecasts, regime classification, statistical alpha).

## Practical note

When the framework involves **data-driven calibration** (e.g., tuning hedge parameters), the right validation target is **robustness / sensitivity** (stress and stability), not **prediction accuracy**.

# Vanilla Options: European Call & Put

## Contract ingredients

- Underlying (spot) price process: $S_t$.
- Strike price: $K > 0$.
- Maturity: $T$.
- Premium: the price paid at $t = 0$ to acquire the option.

## Definitions

- **European Call Option:** the right (not the obligation) to *buy* the underlying at time $T$ for price $K$.
- **European Put Option:** the right (not the obligation) to *sell* the underlying at time $T$ for price $K$.

# How the payoffs arise at maturity $T$

## Call payoff

At maturity you compare $S_T$ with $K$:

$$\text{if } S_T > K \Rightarrow \text{exercise: gain } S_T - K, \qquad \text{if } S_T \leq K \Rightarrow \text{do not exercise: } 0.$$

Hence the **long call** payoff is

$$(S_T - K)^+ = \max\{S_T - K, 0\}.$$

## Put payoff

Similarly:

$$\text{if } S_T < K \Rightarrow \text{exercise: gain } K - S_T, \qquad \text{if } S_T \geq K \Rightarrow \text{do not exercise: } 0.$$

Hence the **long put** payoff is

$$(K - S_T)^+ = \max\{K - S_T, 0\}.$$

(For **short** positions, the payoff is the negative of the above.)

# Payoff diagrams (long positions)

**Long Call: $(S_T - K)^+$**

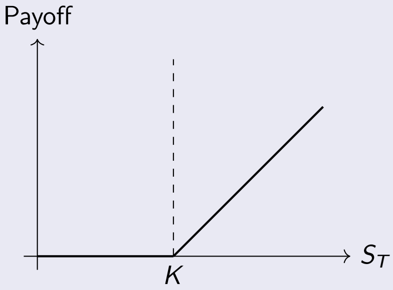


**Long Put: $(K - S_T)^+$**

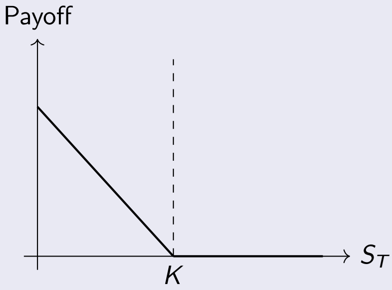

# No-Arbitrage Bounds for a European Call (and tightening with more options)

## Basic no-arbitrage bounds from the underlying and cash

Assume a frictionless market with spot price $S_0$ and no arbitrage. The time-0 price $C(K,T)$ of a European call satisfies

$$\max\{0,\ S_0 - K\,e^{-rT}\} \ \leq\ C(K,T)\ \leq\ S_0.$$

## Tightening the interval using additional traded options

The bounds above use only *static hedges* built from the underlying and cash. If, in addition, we incorporate *all other traded options* (e.g. calls/puts at multiple strikes), then the set of arbitrage free values consistent with market prices shrinks. Consequently, the *no-arbitrage price interval* for $C(K, T)$ becomes *narrower*:

$$\left[\underline{C}(K,T), \overline{C}(K,T)\right] \subseteq \left[\max\{0, S_0 - Ke^{-rT}\}, S_0\right],$$

where $\underline{C}, \overline{C}$ are obtained by enforcing consistency with the observed option surface (equivalently: via tighter super-/sub-replication constraints).

# How Additional Options Tighten Price Bounds

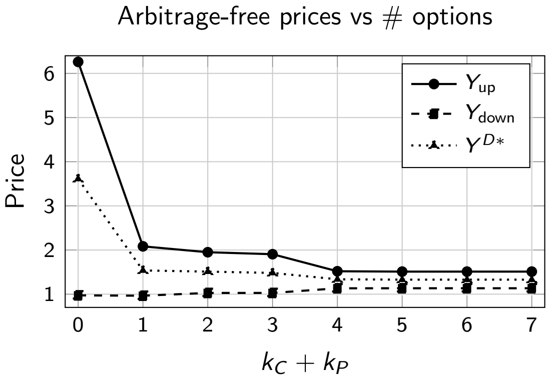


Arbitrage-free prices as a function of $(k_C + k_P)$ (B/W).

## Observation

We observe that as the set of available options expands, the **bounds tighten**. Hence, the **admissible price interval shrinks**, and we progressively move towards a **more targeted (narrower) set of arbitrage-free prices** (illustratively, clustering near $Y^{D*}$ which is the middle).

## Question

Does this type of bound-tightening hold for **all** options?

# Why Model-Free for non-path dependent options? (Motivation)

- **No strong distributional assumptions:** we do *not* want to assume lognormal returns, a specific volatility model, or a jump process.
- **What we actually observe:** today's underlying price $S_0$ and a finite set of market option quotes (calls/puts) at a few strikes and maturities.
- **What we want to guarantee:** *model-independent* price bounds and hedges (super-/sub-replication) that remain valid under **any** terminal distribution consistent with observed prices.

**Key Question**

Given only traded prices, what can we say about the fair value of a payoff $f(S_T)$ *without* specifying a probabilistic model?

# Model-free setup: static portfolios and terminal profit

Fix maturity $T$ and denote by $x := S_T \geq 0$ the terminal underlying price. Assume that, in addition to the underlying, the market trades European calls and puts $\{C(K_i), P(K_i)\}_{i=1}^n$ with strikes $K_i$.

## Static portfolio (underlying + cash + traded options)

Choose positions $(a, b, \gamma_1, \ldots, \gamma_n, \delta_1, \ldots, \delta_n)$, where $a$ is the number of shares, $b$ is the cash position (bank account), and $\gamma_i, \delta_i$ are call/put holdings at strike $K_i$.

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

## Profit for writer vs. buyer of an option with payoff $f(x)$

$$\Pi_{\text{writer}}(x) = V_T(x) - f(x), \qquad \Pi_{\text{buyer}}(x) = V_T(x) + f(x).$$

Requiring $\Pi(\cdot) \geq 0$ for all $x \geq 0$ enforces *no possible loss in any state*.

Note that the payoff $f(x)$ here should be continuous having finitely many branches, with the final branch being linear.

# Upper bound $Y_{\text{writer}}$: super-hedging (writer's problem)

The writer sells the claim $f(S_T)$ and seeks the *smallest* premium that allows a static hedge with *nonnegative terminal profit in every state*.

## Linear program (writer)

$$\begin{aligned} \text{minimize} \quad & Y \\ \text{subject to} \quad & aS_0 + b + \sum_{i=1}^{n} \left(\gamma_i C(K_i) + \delta_i P(K_i)\right) = Y, \\ & \Pi_{\text{writer}}(x) \geq 0 \quad \text{for all } x \geq 0, \\ & a, \gamma_i, \delta_i \in [-N_i, N_i]. \end{aligned}$$

## Definition

Let the optimal value be $Y_{\text{writer}}$. This is the *tightest model-free ask / super-replication cost* consistent with the available traded options.

# Lower bound $Y_{\text{buyer}}$ and the arbitrage-free interval $[Y_{\text{buyer}}, Y_{\text{writer}}]$

The buyer seeks the *largest* price they can pay while still being able to form a portfolio with *no possible loss in any state*.

## Linear program (buyer)

$$
\begin{aligned}
\text{maximize} \quad & Y \\
\text{subject to} \quad & aS_0 + b + \sum_{i=1}^{n} \big(\gamma_i C(K_i) + \delta_i P(K_i)\big) = -Y, \\
& \Pi_{\text{buyer}}(x) \geq 0 \quad \text{for all } x \geq 0, \\
& a, \gamma_i, \delta_i \in [-N_i, N_i].
\end{aligned}
$$

# Static superhedging on a finite grid (one-asset case)

## Piecewise-affine profit function and node-wise minimum check

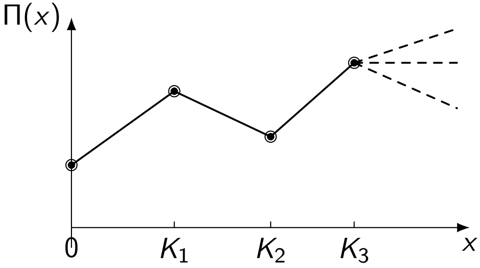


Figure 1: Illustration of a profit function $\Pi(x)$ on $x > 0$ with finitely many affine branches and grid nodes (kinks) at $S = \{0, K_1, K_2, K_3\}$. Since $\Pi$ is affine between consecutive nodes, its minimum over $(0, \infty)$ is attained by checking the node values, including the boundary point $x = 0$, provided the tail slope condition prevents $\Pi(x) \to -\infty$ as $x \to \infty$. The three dashed rays depict alternative tail slopes on the last branch: increasing, flat, and decreasing.

## Arbitrage-free price interval

Let the optimal value be $Y_{\text{buyer}}$. If $Y_{\text{buyer}} < Y_{\text{writer}}$, we define the (no-arbitrage) price interval as $(Y_{\text{buyer}}, Y_{\text{writer}})$. We often denote the admissible range as

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

Any transaction price outside this range creates an arbitrage for one side; within the range, no riskless profit is available.

# Market remark: price discovery and the arbitrage-free interval

## Comment

If an option is actively traded (high liquidity, frequent buy/sell orders), then competitive price discovery and arbitrage activity tend to push its market price into the *arbitrage-free interval* $[Y_{\text{buyer}}, Y_{\text{writer}}]$, provided that this interval is *non-empty* (i.e. $Y_{\text{buyer}} \leq Y_{\text{writer}}$).

# The center price $Y^{D^*}$: an equal-max-loss interpretation

## Arbitrage-free interval (assumed non-empty)

Assume $Y_{\text{buyer}} \leq Y_{\text{writer}}$, so the arbitrage-free interval $[Y_{\text{buyer}}, Y_{\text{writer}}]$ is non-empty.

## Definition (midpoint / center price)

We define the center of the interval by

$$Y^{D^*} := \frac{Y_{\text{buyer}} + Y_{\text{writer}}}{2}.$$

## Physical meaning

At the transaction price $Y^{D^*}$, *both sides* (buyer and writer) can construct a static portfolio such that their *worst-case loss* is the same:

$$\max_{x \geq 0} \big( -\Pi_{\text{buyer}}(x) \big) \;=\; \max_{x \geq 0} \big( -\Pi_{\text{writer}}(x) \big).$$

In other words, $Y^{D^*}$ balances the deal so that the two parties face an equal maximum possible loss under the model-free (pathwise) viewpoint.

# Negotiation remark: $Y^{D^*}$ is a reference, not a prediction

## Comment

The center price $Y^{D^*}$ does *not* imply that the market will trade exactly at $Y^{D^*}$. The actual transaction price may be anywhere inside (or even outside) the arbitrage-free interval $[Y_{\text{buyer}}, Y_{\text{writer}}]$, depending on supply/demand, liquidity, and bargaining power.

## Practical use

Nevertheless, $Y^{D^*}$ can be used as a *reference price* during negotiations: it provides a symmetric benchmark that equalizes the two parties' worst-case exposure under the model-free viewpoint.

## Perspective

We cannot (and need not) *predict* the future price of an option. What matters operationally is *risk management*: how we hedge and manage the position *after* we have bought or sold the option.

Often the *arbitrage–free* interval of prices is *empty*, i.e.

$$Y_{\text{buyer}} > Y_{\text{writer}}.$$

In this case, *for any* proposed option price $Y$ there exists a sure–profit opportunity for one of the two parties: if $Y < Y_{\text{buyer}}$ there is an arbitrage for the buyer, whereas if $Y > Y_{\text{writer}}$ there is an arbitrage for the writer.
Nevertheless, the central price

$$Y^{D^*} := \frac{Y_{\text{buyer}} + Y_{\text{writer}}}{2}$$

is always well-defined and admits a natural "balance" interpretation: at the price $Y^{D^*}$, both parties can construct static portfolios with the *same maximal possible loss* (equal max–loss viewpoint).
Moreover, whenever the corresponding *arbitrage–free* interval is non-empty ($Y_{\text{buyer}} \leq Y_{\text{writer}}$), we have

$$Y^{D^*} \in [\, Y_{\text{buyer}},\ Y_{\text{writer}} \,],$$

and hence $Y^{D^*}$ is itself *arbitrage–free*. These properties make $Y^{D^*}$ distinguishable among other candidate prices and therefore a natural reference price.

# What can we do after buying/selling an option? (Roadmap)

Trading an option is only the beginning. Given the premium paid/received, we can manage the position in (at least) three ways:

1. **Model-free static hedging:** build a portfolio that minimizes the *worst-case loss* (Paper 1).
2. **Forecast-driven static portfolio design:** use a price scenario (forecast) and require *profit in that scenario* while keeping worst-case loss below a fixed budget $D$ (Paper 1).
3. **Dynamic hedging in discrete time:** implement a *self-financing* trading rule across rebalancing dates (Paper 3).

# (1) Model-free static hedging: smallest possible loss (Paper 1)

## Goal

Given the premium (paid or received) and the set of traded instruments (stock, cash, and listed options), construct a *static* portfolio that makes the terminal profit as safe as possible *in every state* $x = S_T \geq 0$. That is, construct a portfolio with maximum possible loss in all scenarios (i.e. $S_T \geq 0$) the amount $D$.

## Interpretation

This is a **worst-case** (pathwise) viewpoint:

- We do *not* assume a probabilistic model for $S_T$.
- We search for a portfolio that minimizes the *maximum possible loss* over all $x \geq 0$, subject to using only available traded instruments.

## Outcome

The resulting construction yields the *best achievable protection* against adverse moves, consistent with no-arbitrage and the observed option prices (super-/sub-hedging via LP).

## Goal

Use the same premium to construct a *static* portfolio that is **profitable under a forecast scenario** while controlling downside risk globally.

## Two requirements

Let $A \subseteq \mathbb{R}_+$ be a forecast region for $x = S_T$ (e.g. "likely outcomes"):

- **Profit on the scenario:** $\Pi(x) \geq 0$ for all $x \in A$.
- **Worst-case loss cap:** $\Pi(x) \geq -D$ for all $x \geq 0$, for a given $D \geq 0$.

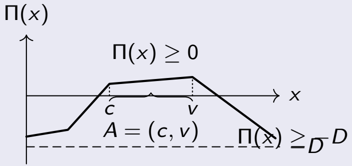

# (2) Forecast-driven design with loss budget $D$ (Paper 1) II

## Interpretation

We *separate* forecasting from portfolio construction: any forecasting method can be used to specify $A$, while the portfolio is computed by optimization (LP) under the above constraints.

## Our viewpoint: classical 2–4-leg strategies as special cases

The standard "named" option strategies built from a small number of legs (typically 2 to 4 European options, with or without stock), such as the Bull Call Spread, Bear Put Spread, Straddles/Strangles, and iron structures, can be viewed as *particular feasible portfolios* in our design space. In our framework we do not pre-select a template; instead, we optimize over a broad class of static combinations of calls/puts (and possibly the underlying), while enforcing the same global risk control (e.g. a loss budget $D$) and the desired forecast-driven profit constraints. Hence, the textbook strategies are simply special cases of our proposed design.

# (3) Dynamic strategy in discrete time (Paper 3)

## Goal

Instead of a purely static hedge, we rebalance positions through time:

$$0 = t_0 < t_1 < \cdots < t_N = T,$$

trading the stock and the bank account (and possibly options) at each $t_k$.

## Key features

- **Self-financing:** portfolio value changes only through gains/losses from holdings.
- **Path dependence:** trading decisions may depend on the observed history $(S_{t_0}, \dots, S_{t_k})$.
- **Objective:** design a rule that manages risk (e.g. profit vs. CVaR) while remaining implementable.

## Interpretation

Dynamic hedging is essential for path-dependent claims and often improves performance relative to purely static constructions, at the cost of rebalancing and model/estimation inputs.

# Path-dependent option (discrete time): payoff and hedged P&L

## Discrete-time setting

Fix trading dates $0 = t_0 < t_1 < \cdots < t_N = T$ and a nonnegative price process $\left(S_{t_k}\right)_{k=0}^{N}$. A **path-dependent** claim is specified by a payoff functional

$$H = H\big(S_{t_0}, S_{t_1}, \ldots, S_{t_N}\big).$$

Let $Y \in \mathbb{R}_+$ denote the option premium paid at inception. Here, we do not assume anything about the payoff function, that is, it can be nonlinear as well.

## Hedged profit-and-loss (P&L)

Each side $s \in \{w, b\}$ (writer $w$, buyer $b$) selects an admissible (self-financing) trading strategy from a set $\Theta_s$. Let $G_T^{\theta}$ denote the terminal trading gain from strategy $\theta \in \Theta_s$. We define the terminal hedged P&L as

$$\Pi_w(Y, \theta) = Y + G_T^{\theta} - H, \qquad \Pi_b(Y, \theta) = -Y + G_T^{\theta} + H.$$

# Downside hedging loss and CVaR

## Downside loss

For each side $s \in \{w, b\}$, define the **downside hedging loss** as the negative part of the P&L:

$$L_s(Y, \theta) \;:=\; \big(\Pi_s(Y, \theta)\big)^- \;=\; \max\{-\Pi_s(Y, \theta),\, 0\}.$$

## Conditional Value-at-Risk

Fix a confidence level $\alpha \in (0, 1)$. The **Conditional Value-at-Risk** of a nonnegative loss $L$ is

$$\mathrm{CVaR}_\alpha(L) \;=\; \inf_{u \in \mathbb{R}} \left\{ u + \frac{1}{1-\alpha}\, \mathbb{E}\big[(L-u)^+\big] \right\}.$$

# Value-at-Risk (VaR): definition and meaning

## Definition

Fix a confidence level $\alpha \in (0,1)$. The **Value-at-Risk** of a loss r.v. $L \geq 0$ is

$$\mathrm{VaR}_\alpha(L) \ := \ \inf\{\ell \in \mathbb{R} : \ \mathbb{P}(L \leq \ell) \geq \alpha\}.$$

Equivalently, if $F_L(\ell) = \mathbb{P}(L \leq \ell)$ is the CDF of $L$, then

$$\mathrm{VaR}_\alpha(L) = F_L^{-1}(\alpha).$$

## Sketch (CDF viewpoint)

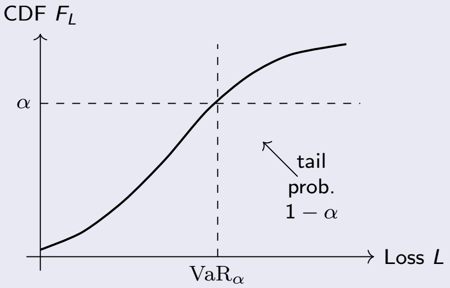


## Physical interpretation

$\mathrm{VaR}_\alpha(L)$ is the $\alpha$-quantile (a **loss threshold**): with probability at least $\alpha$, the loss will *not exceed* $\mathrm{VaR}_\alpha(L)$. The remaining probability mass $1-\alpha$ corresponds to the **tail events** where losses are larger.

# Conditional Value-at-Risk (CVaR): physical meaning

## Physical interpretation

Fix $\alpha \in (0,1)$. The quantity $\mathrm{CVaR}_\alpha(L)$ measures the **average loss in the worst** $100(1-\alpha)\%$ **cases**.

- $\mathrm{VaR}_\alpha(L)$ is the $\alpha$-quantile of the loss distribution (a cutoff level).
- $\mathrm{CVaR}_\alpha(L)$ looks *beyond* that cutoff and averages the tail losses.
- Hence, CVaR is a **tail-risk** criterion: it penalizes rare but severe outcomes.

If $L$ has a continuous distribution, one can write

$$\mathrm{CVaR}_\alpha(L) = \mathbb{E}[L \mid L \geq \mathrm{VaR}_\alpha(L)].$$

## Sketch (loss distribution tail)

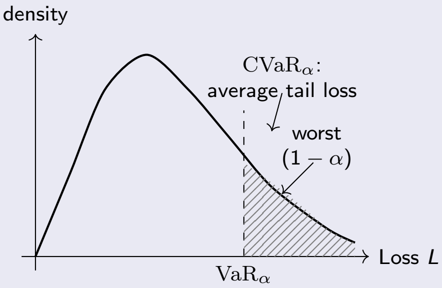

# CVaR-based fair price for a path-dependent option

## Minimal residual downside tail risk (per side)

At premium $Y$, side $s \in \{w, b\}$ chooses the strategy that minimizes tail downside risk:

$$R_s(Y) := \inf_{\theta \in \Theta_s} \mathrm{CVaR}_\alpha\big(L_s(Y, \theta)\big).$$

Here $R_s(Y)$ is the **best achievable residual downside risk** (in the $\mathrm{CVaR}$ sense) after optimal hedging by side $s$.

## Definition (fair price in the equal-downside-risk sense)

A premium $Y^\star$ is called a "**fair price**" if it equalizes the two sides' best achievable residual downside tail risk:

$$R_w(Y^\star) = R_b(Y^\star).$$

## Interpretation

At $Y^{\star}$, both the writer and the buyer can hedge optimally so that their *minimal* $\mathrm{CVaR}$-downside exposure coincides; the transaction is balanced in terms of tail-loss risk.

# Remark: CVaR-based fair price as a negotiation benchmark

## How to read $Y^\star$ (order-of-magnitude value)

- Calibrated from **historical/backtested** data *or* a **reference model + simulation**.
- Therefore, $Y^\star$ is **not** a guaranteed "true" market price.
- Use $Y^\star$ as a **benchmark**: a reference point to structure negotiation.

## Negotiation: bounds differ, but must be feasible

**Seller (writer): lower bound $Y_{\text{min}}$**

- Wants an acceptable *minimum* price.

**Buyer: upper bound $Y_{\text{max}}$**

- Wants an acceptable *maximum* price.

- Bounds may differ due to: datasets, calibration windows, backtesting protocols, model assumptions, stress scenarios.
- **Credibility requires feasibility:** a proposed interval should come with an **implementable hedge** in the chosen market model meeting the relevant risk constraints (e.g. CVaR limits).

# Beyond Option Pricing: What about Markowitz?

## Motivating question

We understand option pricing. How does this relate to Markowitz portfolio theory—and can we do better?

**Distributional portfolio design with a loss budget.** Fix an acceptable loss level $D \geq 0$. Choose portfolio coefficients (cash, underlyings, and traded options) by solving

$$\begin{aligned} \max \quad & F\Big(\mathbb{P}(\Pi(S^T) \geq 0), \mathbb{E}[\Pi(S^T)], \mathrm{Var}\,(\Pi(S^T)), \dots\Big) \\ \text{s.t.} \quad & \Pi(x) \;\geq\; -D, \qquad \forall x \in \mathbb{R}^d_+, \\ & c_i \in [-N_i, N_i] \qquad \text{for all } i. \end{aligned} \tag{1}$$

Here $\Pi(x)$ is the terminal profit (net payoff) of the constructed portfolio when terminal prices are $x \in \mathbb{R}^d_+$, and $F$ is an investor-chosen objective encoding a trade-off between performance and risk under the assumed joint distribution of $S^T$. The variables $c_i$ are portfolio decision parameters (shares, option contracts, and cash), and the bounds $N_i$ model position limits due to liquidity and operational constraints.

**Markowitz as a special case:** if option positions are disallowed (i.e., $N_i = 0$ for option coefficients), the problem reduces to the classical mean–variance portfolio selection framework.

**Remark.** With a statewise loss budget $\Pi(x) \geq -D$, VaR-style constraints become essentially non-binding, since losses cannot exceed $D$ in any scenario.

## Sanity check

Any pricing and hedging approach should be able to defend the quoted price: *Why is it reasonable to trade at this level? Does the method prescribe a feasible hedging strategy? If not, what is the point of studying it?*

# Dynamic hedging in Black–Scholes: a very specific sigmoid rule

## Black–Scholes model (reference)

Assume the stock follows

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t, \qquad r = \text{constant},$$

and consider a European call with strike $K$ and maturity $T$.

# Core Assumptions of the Black–Scholes Model

The Black–Scholes framework relies on the following assumptions:

- The (instantaneous) volatility $\sigma$ over the horizon $[0, T]$ is known.
- The investor can trade continuously in time to construct a replicating portfolio.

### How do we predict future volatility?

A common approach is to compute historical volatilities from past data and assume that the volatility over $[0, T]$ will be an appropriate average of those past values.

### Implied Volatility: What it is

*Implied volatility* $\sigma_{\text{impl}}$ is the volatility input that, when plugged into the Black–Scholes formula, reproduces the observed market price of the option.

# Is the Black-Scholes price Arbitrage-Free?

## Arbitrage-Free Value

The Black-Scholes theory states that this value is arbitrage-free. But this is true only if someone can indeed replicate the hedging portfolio continuously in time!

## Is the Black-Scholes theory a prediction method?

If you are not intended to use the proposed replication method continously in time then the Black-Scholes method can be seen only as a prediction technique. Do you really believe that?

# Are these assumptions true in the real world?

## Why Volatility Smiles?

In the real world, no one can replicate continuously in time. The computation of future volatility is, in fact, a prediction. Volatility smile means that the behavior of the asset depend on the option characteristics! That's why volatility smiles! :-)

## Despite the mismatch with practice

Although these assumptions are not strictly valid in the real world, the research literature has largely evolved toward developing increasingly sophisticated models *but under the same idealized assumptions*, rather than focusing on making those assumptions more realistic. For instance, a more constructive direction is to study hedging strategies that are genuinely implementable in practice—in *discrete time*—while explicitly accounting for transaction costs and liquidity constraints. This is exactly what we do next; and, more importantly, we will leave no open problems behind!

## Call price and delta

Under Black–Scholes,

$$C(t,S) = S\,\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

where

$$d_1 = \frac{\ln(S/K) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}, \qquad d_2 = d_1 - \sigma\sqrt{T-t},$$

and $\Phi(\cdot)$ is the standard normal CDF. The **replicating strategy** holds

$$\Delta(t,S) = \frac{\partial C}{\partial S}(t,S) = \Phi(d_1) \in (0,1).$$

# Beyond GBM: why Black–Scholes dynamics may be misspecified

## GBM is a convenient reference, not a law of nature

The Black–Scholes model assumes continuous paths and lognormal returns:

$$dS_t = \mu S_t \, dt + \sigma S_t \, dW_t.$$

In practice, market data often exhibit features that are inconsistent with pure GBM.

## Empirical deviations from GBM (typical)

- **Jumps / discontinuities:** sudden moves due to news, earnings, macro shocks.
- **Heavy tails & skewness:** return distributions with more extreme events than Gaussian.
- **Volatility clustering:** time-varying volatility (stochastic volatility / GARCH-like behavior).
- **Smile/skew:** option-implied distributions not consistent with constant $\sigma$.

# Alternative dynamics: jumps and Lévy models (and the case for model-agnostic hedging)

## Examples of richer dynamics

A common extension is a jump-diffusion (Merton-type):

$$\frac{dS_t}{S_{t-}} = \mu \, dt + \sigma \, dW_t + (J - 1) \, dN_t,$$

where $N_t$ is a Poisson process and $J$ is a jump size r.v. More generally, one may assume that $\log S_t$ is driven by a **Lévy process** (allowing jumps and heavy tails), possibly combined with stochastic volatility.

## Which model is true? What about the replicating strategy?

This raises a key question: which of these models should be regarded as the “right” one? The seller might simply adopt the specification that delivers the highest valuation, while the buyer will naturally prefer the one that generates the lowest. Bargaining then becomes unavoidable—because we no longer have a single, mutually acceptable price in hand.

And what about replication? Do these approaches actually yield feasible hedging strategies, or do they only produce a range of valuations without an implementable hedge?

## Can you replicate continuously in time? What are the consequences of this discreteness?

It is well known that *continuous-time* replication is not feasible in practice: trading occurs at discrete dates, transaction costs are present, and model parameters (e.g. volatility) are uncertain. What are the consequences of this discreteness?

In particular:

- How does the Black–Scholes $\Delta$-hedging strategy behave under *discrete* rebalancing?
- How large is the residual hedging error, and how does it depend on the rebalancing frequency?
- Ultimately, does the Black–Scholes model provide a convincing rationale for *why* one should buy or sell at the quoted price?
- The same questions holds for any other similar continuous time replication method.

## Key observation

The mapping $z \mapsto \Phi(z)$ is **sigmoid-shaped**. Hence, in Black–Scholes the stock holding is given by a *very specific* sigmoidal function of a standardized log-moneyness variable $d_1$. Changing model will drive us to a different hedging strategy. But we have the **Universal approximation property of sigmoidal functions!**. With this, we can approximate any possible hedging strategy.

**Reminder:** $\Phi(d_1)$ is the Black–Scholes *delta* (the shares held per option).

# From models to a universal feedback form: find the rule $f$ I

**Discrete-time trading.** On a grid $0 = t_0 < t_1 < \cdots < t_N = T$, trade only stock and cash (self-financing). At each time $t_k$ we choose a *predictable* stock position

$$a_k \;=\; f\big(t_k, T - t_k, S_{t_0}, \ldots, S_{t_k},\, X_{t_k}\big) \qquad \text{(Markov special case: } a_k = f(t_k, T - t_k, S_{t_k}, X_{t_k})\text{)}.$$

**What is $X_{t_k}$? (recent information / news layer)**

- $X_{t_k} \in \mathbb{R}^p$ is a vector of *exogenous signals* observable at (or just before) $t_k$, built from recent information beyond prices (news, macro releases, order-flow/liquidity indicators, etc.).
- Example (news analytics): $X_{t_k} = \mathrm{Sent}_{t_k}$, a rolling sentiment score computed from headlines and announcements over $(t_{k-1}, t_k]$.
- The role of $X_{t_k}$ is to let the rule react to newly observed conditions during rebalancing (the forecasting layer supplies $X$, the construction layer optimizes/calibrates $f$).

**Universal portfolio form.** Any such strategy necessarily has portfolio value

$$V_{t_k} = a_k S_{t_k} + b_k = f\big(t_k, T - t_k, S_{t_0}, \ldots, S_{t_k}, X_{t_k}\big)\, S_{t_k} + b_k,$$

where $b_k$ is the cash account (endogenously determined by self-financing once $a_k$ is fixed).

**Extended universal portfolio form (allowing long-only American options).** Similarly, we may allow the portfolio to hold *American* calls/puts *long only*. (Recall: for American-type contracts, model-free static pricing bounds can be obtained using portfolios of American calls and puts, but restricting positions to the long side only.)
Let $C^A_{t_k}(K_i), P^A_{t_k}(K_i)$ denote traded American call/put prices at time $t_k$, for strikes $K_1, \ldots, K_m$. At each rebalancing date $t_k$ we choose holdings

$$a_k \in \mathbb{R}, \qquad c_{k,i} \geq 0, \qquad d_{k,i} \geq 0, \qquad i = 1, \ldots, m,$$

and (optional) exercise quantities

$$e^C_{k,i} \in [0, c_{k,i}], \qquad e^P_{k,i} \in [0, d_{k,i}], \qquad i = 1, \ldots, m,$$

# From models to a universal feedback form: find the rule $f$ III

where $e^C_{k,i}$ (resp. $e^P_{k,i}$) is the number of call (resp. put) contracts exercised at $t_k$.
A universal feedback specification is then

$$(a_k,\ c_k,\ d_k,\ e^C_k,\ e^P_k) = F\big(t_k, T - t_k, S_{t_0}, \ldots, S_{t_k}, X_{t_k}\big),$$

and the (marked-to-market) portfolio value becomes

$$V_{t_k} = a_k S_{t_k} + b_k + \sum_{i=1}^{m} c_{k,i}\, C^A_{t_k}(K_i) + \sum_{i=1}^{m} d_{k,i}\, P^A_{t_k}(K_i),$$

where $b_k$ is the cash account (endogenously determined by self-financing once $(a_k, c_k, d_k, e^C_k, e^P_k)$ is fixed).

**Remark (practicality).** Now compare the strategy

$$a_k = f(T - t_k, S_{t_k})$$

# From models to a universal feedback form: find the rule $f$ IV

with the proposed universal feedback rule

$$a_k = f\big(t_k, T - t_k, S_{t_0}, \dots, S_{t_k}, X_{t_k}\big),$$

and you will see that the latter is *by far* more realistic and more efficient in practice. Notice that the Black–Scholes strategy, and essentially all of its standard (parametric) extensions, lead in discrete time to trading rules that fall into exactly this kind of simple **universal portfolio form**, that is $a_k = f(t_k, S_{t_k})$.

# From models to a universal feedback form: find the rule $f$ V

## So what is the *optimal* $f$? (example criteria)

**Parametric approximation.** We restrict $f$ to a flexible *sigmoidal* family and optimize over its parameters:

$$f(s) \approx f_\theta(s) \quad \text{with} \quad f_\theta(s) := \sum_{j=1}^{J} a_j\, \sigma\big(b_j(s - c_j)\big) + d, \qquad \sigma(x) := \frac{1}{1 + e^{-x}},$$

so choosing $f$ is equivalent to choosing $\theta = (a_j, b_j, c_j, d)_{j=1}^{J} \in \Theta_s$.

Fix a premium $Y$ and choose $f$ (equivalently, $\theta \in \Theta_s$) by solving

$$R_s(Y) \ := \ \inf_{\theta \in \Theta_s} \ \rho\big(L_s(Y, \theta)\big),$$

with a risk criterion $\rho$, e.g.

- $\rho = \mathrm{CVaR}_\alpha$ (tail-risk optimal hedging),
- mean–CVaR trade-off, or $\min \mathbb{E}[L]$ subject to $\mathrm{CVaR}_\alpha(L) \le D$,

# Remark: selecting a dynamic strategy from an infinite family (data-driven) I

## Universal approximation property of sigmoidal functions

For a given option, we may choose the most appropriate hedging rule from an *infinite class* of admissible (self-financing) discrete-time strategies. This class *contains* the Black–Scholes delta hedge as a special case, while also covering a vast range of alternative hedging rules enabled by the **universal approximation property of sigmoidal functions**.

## What are the best criteria? Examples...

- **Tail-risk:** min $\mathrm{CVaR}_\alpha(L_s)$
- **Mean–tail:** min $\mathbb{E}[L_s] + \lambda\,\mathrm{CVaR}_\alpha(L_s)$
- **Risk budget:** min $\mathbb{E}[L_s]$ s.t. $\mathrm{CVaR}_\alpha(L_s) \leq D$
- **Robustness:** min $\sup_{\omega\in\Omega} \mathrm{CVaR}_\alpha(L_s^\omega)$

# Approximation is not unique

- We model the hedge as a **bounded decision rule** in discrete time:

$$a_k = f(t_k, \cdots).$$

- Sigmoids are **one** convenient approximation family, but not the only one (e.g. **Fourier/Chebyshev**, **splines**, **wavelets**, **ReLU/hinge bases**, ...).
- **Research agenda:** determine which approximation family is best **per payoff class and regime**, using criteria such as **out-of-sample hedging P&L** (mean–CVaR / tail risk), **robustness** under stress, and **computational efficiency**.

## No need to commit to a single stochastic model

We do not need to decide *a priori* which stochastic process drives the stock price (e.g. diffusion vs. Lévy, jump models, etc.). The strategy is selected by calibration on *historical data* (and/or simulated scenarios), so jump-type behavior can be incorporated whenever it is supported by the data. Hence, no volatility smiles!

Of course, in the path-dependent case, one should not generally expect to obtain either an arbitrage-free price interval for the option or a single valuation on which all investors would agree. However, the above method (which is model-free but data driven) provides feasible hedging strategies once the price of the option is specified.

### Stress-testing with extreme scenarios

Even if jumps or extreme moves have *not* been observed in the past, we can enrich the calibration set with *stress scenarios* (tail events) to make the chosen strategy more robust to rare but impactful outcomes.

# Rolling 3-Month Scenarios from Daily Close Data

**Input:** daily close prices $S_1, \ldots, S_N$.

**Rolling-window construction (moving block):**

- Fix a 3-month horizon as $m$ **trading days** (typically $m \approx 63$).
- Slide the start date forward by **one day**.
- Each window position is one historical 3-month **scenario path**.

This generates a large set of overlapping scenarios that preserve realistic within-quarter dynamics (drawdowns, rebounds, volatility clustering).

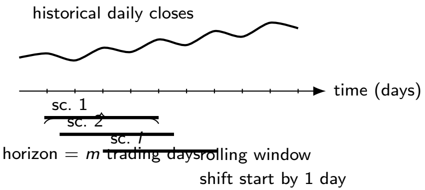

## Scenario Set: Historical Rolling Windows + Forward-Looking Stress Paths I

**Historical rolling-window scenarios.** Given daily closes $S_1, \ldots, S_N$ and a 3-month horizon of $m$ trading days, define for each start day $i$ the path

$$\omega_i := (S_i, S_{i+1}, \ldots, S_{i+m}), \qquad i = 1, \ldots, N-m.$$

**Count:** $M_{\text{hist}} = N - m$.

**Optional re-basing to today's spot $S_0$.**

$$\widetilde{S}_k^{(i)} = S_0 \cdot \frac{S_{i+k}}{S_i}, \qquad k = 0, \ldots, m.$$

**Augmenting the scenario set (forward-looking / hypothetical paths).** Beyond historical windows, we can add **non-historical** scenarios that have *not* occurred in the sample but are **plausible** under our views or fears (e.g., crash/gap moves, volatility spikes, regime shifts, event risk).

# Scenario Set: Historical Rolling Windows + Forward-Looking Stress Paths II

Let $\Omega_{\text{hist}}$ be the historical set and $\Omega_{\text{stress}}$ the added set. We work with the augmented scenario universe

$$\Omega = \Omega_{\text{hist}} \cup \Omega_{\text{stress}}.$$

**Why do this?** We can **tune model/valuation parameters** (or their priors/weights) so that risk measures and option prices are consistent with both:

- the empirical distribution from history, and
- our forward-looking expectations about tail events.

# Forecasting vs. Construction: A Clean Separation

- **Scenario selection is the forecasting layer:** it requires experience, judgment, and intuition.
- **Construction is different:** once scenarios (or data) are given, building an *optimal, implementable* dynamic strategy is a *mathematical optimization problem*.
- In practice we select a flexible class of decision rules (e.g., **sigmoidal families**) and **calibrate parameters** to minimize a risk criterion (CVaR / mean–CVaR / risk-budget).
- **Pedagogical point:** students (even MSc) should be trained *primarily* on the construction toolkit (no-arbitrage logic, constraints, replication, optimization); forecasting can come *much later* as a modular input.

# Unique construction principle – non-unique approximations

- **The construction methodology is unique:** once the input is fixed (forecast / scenarios / data), we solve a well-posed optimization/hedging problem to obtain an implementable decision rule.
- **The approximation family is not unique:** sigmoids are convenient, but one can also use Fourier/Chebyshev bases, splines, wavelets, etc.
- **But the output is conditional on the scenarios:** change the scenario set (even by adding a few "wrong" scenarios) and the calibrated *optimal* dynamic strategy can change.
- **Takeaway:** stick to *one* approximation family you know well and refine it systematically.

### Adaptation during rebalancing

Finally, the framework naturally allows us to incorporate *recent market information* as we rebalance: parameters can be updated on a rolling window and the trading rule can react to newly observed conditions, such as news analytics etc.

### Multi-asset options

Similar arguments and methods extend to multi-asset options, whether European/American-style (non path-dependent) or path-dependent.

# Pricing Bounds and Feasibility

- An investor can determine the lowest (resp. highest) option price they would accept (resp. be willing to pay) while still being able to implement a chosen dynamic hedging strategy. In this context, the "fair" price $R(Y^\star)$ provides a natural benchmark.
- For **non–path-dependent** options, these amounts (coming from dynamic hedging strategies) need not lie within the arbitrage-free interval. A given hedging strategy may produce prices inside the arbitrage-free price interval, yet still be unattainable for a capital-constrained investor—especially when the hedging portfolio must accommodate multiple options. Note, however, that the arbitrage-free interval is typically pinned down by well-capitalized (i.e., "strong") market participants.
- In general, we can not define an arbitrage - free interval to **path-dependent** options.

# Market Pricing vs. Negotiated Pricing

- **If the market price is already set:** the investor must decide whether that price is compatible with the intended hedging strategy, either static or dynamic; if not, they should avoid entering the contract.
- **If the market price is not set:** the investor should compute the minimum/maximum amount they can receive/pay to implement the desired hedge; these bounds are useful during negotiation.

## When do we need Black-Scholes?

**So do we ever really need Black–Scholes?** In principle, *no*. Oops—there is one case: when you *don't know (or don't use) the present methodology*; then BS is used merely as a *prediction crutch*, not as an implementable hedging prescription.

*Some may ask: "I want to know why this is the option price." But the price changes every 1–2 minutes, and so do the forces shaping it—what are you going to do, re-derive the reasons every two minutes?*

# Methodological Implications

- Any option pricing method that does not yield a *practically implementable* hedging strategy is of no use to the investor (except when used purely for prediction).
- Any technique that **does** propose an implementable hedging strategy is encompassed within our methodology and results via the *universal approximation property of sigmoidal functions*.

# Conclusion-Part I I

## European-Type Options

Consider an option with payoff $f(x)$ that is piecewise linear with finitely many linear segments. For such payoffs, one can determine a deterministic no-arbitrage price interval by constructing static portfolios of European calls and puts, allowing both long and short positions.

## American-Type Options

Consider an option with payoff $f(x)$ that is piecewise linear with finitely many linear segments. For these contracts, a deterministic no-arbitrage price interval can also be obtained using portfolios of American calls and puts, but restricting positions to the long side only.

# Conclusion-Part I II

## Forecast-driven portfolio design and dynamic overlay

Beyond pricing bounds, our approach can be used to *construct an optimal portfolio* that is **profitable on a forecast scenario** (i.e., on a predicted set of outcomes), while respecting a prescribed downside constraint (e.g. a loss budget). Moreover, one may **overlay a dynamic trading strategy**—as described above—to adapt exposures over time, enhancing robustness and performance under evolving market conditions.

## From pricing bounds to negotiation and hedging

The *center* (midpoint) of the arbitrage-free price interval can be viewed as a natural **reference quote** for the negotiation phase, providing an objective benchmark between the admissible bid–ask extremes. Hence, our methodology is not only about **pricing** (via $Y^{D*}$), but also about **hedging**: we can construct *static* hedging portfolios (built from calls/puts and possibly the underlying) and, when desired, implement a *dynamic* trading overlay as described above.

# Conclusion-Part II I

## Path-Dependent Options

For path-dependent options, a deterministic no-arbitrage price interval cannot, in general, be defined or computed. Hedging may be approached via dynamic trading strategies, but such strategies do not yield deterministic arbitrage-free price bounds in the same sense as the static constructions above.

## Non Piecewise-Linear Payoffs

If the payoff function $f(x)$ is not piecewise linear with finitely many linear segments, one may still employ a dynamic trading strategy of choice for hedging purposes. In general, purely static portfolios of traded options are not sufficient to hedge the payoff or to derive deterministic no-arbitrage price bounds.

# Conclusion-Part II II

## Pricing via $R(Y^*)$ and discrete-time dynamic hedging

In this setting, our output $R(Y^*)$ can still be interpreted as a **pricing method**: it provides a principled reference value (or range) consistent with the admissible market model and the imposed risk constraints. At the same time, the same framework naturally supports **hedging** through **discrete-time dynamic strategies**, i.e. by rebalancing a self-financing portfolio along time on a chosen trading grid, as described above.

# On the notions of Arbitrage (I)

## Static market setup and profit function

Let $x = S_T$. At $t = 0$ we can form a *static* portfolio using: (i) $a$ shares of the stock, (ii) $b$ units in the bank account (so $b \mapsto b\,e^{rT}$ at $T$), and (iii) a finite set of European calls/puts with strikes $K_i$ and prices $C(K_i), P(K_i)$.
The resulting terminal profit (payoff) function is

$$h(x) = a\,x + b\,e^{rT} + \sum_i \gamma_i (x - K_i)^+ + \sum_i \delta_i (K_i - x)^+,$$

with initial cost

$$V_0 = aS_0 + b + \sum_i \gamma_i C(K_i) + \sum_i \delta_i P(K_i).$$

# On the notions of Arbitrage (II)

## Deterministic vs Statistical Arbitrage

**Deterministic (model-free / static) arbitrage.** A deterministic arbitrage exists if we can choose positions so that

$$V_0 \leq 0, \qquad \Pi(x) := h(x) \geq 0 \;\; \forall x \geq 0, \qquad \Pi(x) > 0 \;\; \text{for some } x \geq 0.$$

This is a *statewise dominance* condition: no downside in any scenario and some upside, with no net investment (possibly an initial cash inflow). Because it is riskless and scalable, this notion imposes *hard* no-arbitrage relations (e.g. put–call parity, monotonicity/convexity) and is therefore the notion that *shapes market prices*.

**Statistical arbitrage.** By contrast, statistical arbitrage refers to strategies that are profitable only *in expectation* or *with high probability*. Losses can occur over finite horizons and performance depends on modeling/estimation and rebalancing. Hence, it does *not* generate deterministic no-arbitrage price bounds and is not the mechanism that enforces price relations in the market.

# Deterministic arbitrage is what enforces prices (and relations) I

## Pricing discipline

- **Deterministic (statewise) arbitrage** is a dominance condition: no downside in *any* scenario, some upside, with non-positive initial cost.
- Because it is **riskless and scalable**, it generates **hard constraints** that *must* hold in market prices (monotonicity, convexity, **put–call parity**, etc.).
- Hence, **this** is the notion that shapes option prices (not statistical arbitrage).

# Deterministic arbitrage is what enforces prices (and relations) II

## Example: put–call parity is a deterministic no-arbitrage relation

Consider European options with strike $K$ and maturity $T$ (no dividends). At maturity,

$$(S_T - K)^+ + K \;=\; (K - S_T)^+ + S_T.$$

Thus the portfolios *(Call + bond)* and *(Put + stock)* have identical payoff for every $S_T$. No deterministic arbitrage implies equal time-0 costs:

$$C_0 + Ke^{-rT} \;=\; P_0 + S_0.$$

Any violation yields a *riskless* profit via a static trade.

*Remark.* Throughout we use **deterministic** no-arbitrage arguments to derive such exact relations. However, much of the literature shifts to **statistical arbitrage** (profitability "on average" or with high probability), which is a different notion and can be a source of substantial confusion if the transition is not made explicit.

# Model-based vs. model-free (data-driven) hedging

**Classical model-based**

Historical data $S_t$

↓

Assume model (e.g. SDE)

↓

Estimate parameters (e.g. $\sigma, \mu$)

↓

$\Delta$-hedge at discrete $t_k$

Model risk: the SDE assumption may be wrong (*misspecification*).

**Paper: model-free, data-driven**

Paths: historical + stress/scenarios

↓

Choose criterion (e.g. mean–CVaR)

↓

Optimize discrete-time hedge (sigmoidal family)

↓

Implement at $t_k$ without an SDE

Bypass modelling and compute the “best” strategy directly from paths.

*And for those who seek "pure" mathematical beauty: the literature still offers difficult, deep problems!*

### Simplicity Wins (in Practice)

Most of the time, the right answer is the *least complicated* one that you can actually execute. Sigmoidal-based optimal strategies may sacrifice some Black–Scholes elegance, but they win on practicality and implementable hedging.

**Takeaway**

In option pricing and hedging, every question has been settled *except those that require genuine prediction*.

# Indeed, we leave no open problems beyond prediction

**Static portfolios.** If we restrict ourselves to *static* positions, the only additional degree of freedom is to expand the set of traded contracts—for instance by including forwards, futures, or similar instruments. From a mathematical standpoint, this alters very little: these payoffs are typically still *piecewise linear*, with finitely many kinks, and therefore fit naturally within the same framework.

**Dynamic portfolios.** For *dynamic* trading, we push the model as far as it can reasonably go: given a concrete dataset, we approximate the best *admissible* rebalancing rule—namely, the best dynamic strategy *conditional on the information available*.

*In this sense, the only genuinely open questions that remain are predictive in nature.*

# In a Nutshell: What Do We Mean by Financial Mathematics? I

## Core idea

**Financial Mathematics** is the **mathematics that comes after prediction**: *given a forecast (or data)*, we **formulate** and **solve** a well-posed decision problem (portfolio / hedge / strategy) using **no-arbitrage principles** and **optimization**.

1. **Deterministic arbitrage:** solid understanding and correct application of *deterministic (statewise) no-arbitrage* as a discipline of pricing relations.
2. **Portfolio construction with options:** *given a prediction*, modeling the construction of a portfolio (including options) as an **optimization problem** under constraints (risk/positions/budget).
3. **Option pricing (non path-independent):** set up the super-/sub-replication problems that yield the **arbitrage-free price band** for **non path-independent** payoffs. Given a predictive view, construct the **optimal static hedge** that delivers positive payoff under the predicted scenario.
4. **Dynamic trading (given a prediction):** formulating an optimization problem to find the **optimal dynamic strategy** for both **non path-dependent** and **path-dependent** options.

# In a Nutshell: What Do We Mean by Financial Mathematics? II

## Remark (on prediction)

Next, one may *carefully* mention a few forecasting techniques. For example, if the stock price follows a GBM, one can find a set $A \subseteq \mathbb{R}_+$ such that $P(S_T \in A) = p$. However, it should be made explicit that forecasts of this type **do not incorporate recent events** (or news-driven information) and are primarily model-based distributional statements.

## Forecasting is not a purely mathematical object.

The forecasting stage is intrinsically chaotic: it is driven by noisy data, regime shifts, and reflexive crowd dynamics, so deterministic prediction is out of reach. Of course, forecasting *contains* mathematics—arguably the mother of all sciences—through probability, statistics, and dynamical-systems ideas. Yet forecasting itself is not an exclusively mathematical object: it is an empirical, information-driven activity shaped by institutions, incentives, and human behavior, and its effective "rules" can change as the market adapts.

# Finance $\leftrightarrow$ Actuarial: layer payoffs as a common language

## Finance (Options / Hedging)

- Underlying terminal value: $S_T$.
- Layers via vanilla options:

$$(S_T - K)^+, \quad (K - S_T)^+.$$

- Static portfolios (cash + stock + traded options) deliver super-/sub-hedging bounds.
- Risk criteria (tail risk): $\mathrm{VaR}_\alpha$, $\mathrm{CVaR}_\alpha$ of hedging loss.

## Insurance (Loss / Reinsurance)

- Aggregate loss over horizon: $L \geq 0$.
- Stop-loss / XoL indemnity (a tail layer):

$$I_K(L) = (L - K)^+.$$

- Retained loss (retention):

$$L^{net} = L - I_K(L) = L \wedge K = K - (K - L)^+.$$

- Tail transfer reduces $\mathrm{VaR}_\alpha$ and $\mathrm{CVaR}_\alpha$ of net losses.

### One-line analogy

$$\textbf{Option layers on } S_T \quad \longleftrightarrow \quad \textbf{Reinsurance layers on } L.$$

# Why sigmoids? Bounded decision rules in finance & insurance

## (A) Finance: Black–Scholes delta is already a sigmoid

For a European call,

$$\Delta(t,S) = \Phi(d_1) \in (0,1), \qquad d_1 = \frac{\ln(S/K) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}.$$

Hence $\Delta$ is a *sigmoidal* function of log-moneyness $\ln(S/K)$.

## (B) Generalisation: replace the fixed $\Phi$ by a calibratable sigmoid family

A flexible bounded hedge ratio:

$$\Delta(t,S) \approx a + (b-a)\,\sigma\Big(k\big(\ln(S/K) - m(t)\big)\Big),$$

or more generally a shallow sigmoidal network (discrete-time implementable rule)

$$f_\theta(t,s) = \sum_{j=1}^{m} \alpha_j\, \sigma(\beta_{j,0} t + \beta_{j,1} s + \gamma_j)\,, \qquad a_k = f_\theta(t_k, S_{t_k}).$$

## (C) Actuarial: the same sigmoid logic for probabilities, reserves, tail risk

$$p(x) \approx \sigma(ax+b) \in (0,1), \qquad R(z) \approx a + (b-a)\sigma(g(z)),$$

so outputs are *bounded/stable* and can approximate rich actuarial maps on compact domains.

# From Distribution Fitting to Optimization

- The goal is not merely to identify the “best” distribution:
  - distributional models can be refined to achieve arbitrarily good fits.
- The fundamental problem is to **translate actuarial decisions into optimization**:
  - determine the **premium** and select the **best available reinsurance arrangement**
  - by minimizing/maximizing an objective function subject to constraints
  - according to **chosen criteria** (risk measures, capital requirements, utility, etc.).

# References

Below you can find the relevant references. Item [6] contains the Python code used primarily in [2].


[1] N. Halidias. *Option Pricing of Non Path Dependent Options*. *Preprint (ResearchGate)*, 2026.

[2] N. Halidias. *An overview of financial mathematics with Python codes*. *Monte Carlo Methods and Applications*, 30(3):249–267, 2024. doi:10.1515/mcma-2024-2009.

[3] N. Halidias. *Dynamic Trading and Path Dependent Options*. Preprint (ResearchGate), 2025.

[4] N. Halidias. Financial Engineering with Python. Preprint (ResearchGate), 2025.

[5] N. Halidias. Sigmoidal, Data-Driven Approximations for Core Actuarial Problems. Preprint (ResearchGate), 2026.

[6] N. Halidias. *Financial Engineering: Python codes (GitHub repository)*. *GitHub repository*, accessed February 7, 2026.