Designing Dynamic Trading Strategies

Nikolaos Halidias

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

April 21, 2026

Why the Classical Starting Point is Too Narrow

The limitation

Classical option-pricing frameworks are built on two powerful but restrictive idealisations:

- the finite-state paradigm of tree-based models,
- and the continuous-trading illusion behind continuous-time replication.

Both are mathematically elegant, but neither describes how hedging is actually implemented in real markets.

The practical issue

If the hedging rule is not feasible, then the associated pricing logic is not directly implementable either. Its conclusions may therefore be materially disconnected from market reality. The real questions are: how large is this gap, and what is the economic cost of ignoring it?

A Different Starting Point

Central idea

We remove both restrictions and work directly with all feasible discrete-time hedging strategies.

- We do not assume that the stock price evolves on a finite tree.
- We do not rely on idealised continuous rebalancing.
- We start from the class of implementable trading rules in discrete time, including static hedging positions.
- The only economically natural state restriction is

$$S_T \geq 0.$$

What Determines the Hedging Rule?

Selection principle

Once the arena is opened to all feasible discrete-time hedging strategies, we let optimisation, explicit risk criteria, and market data determine which rule survives.

- We do not declare in advance that one particular rule must be optimal.
- The framework tests candidate strategies inside a larger feasible class.
- What survives is the strategy that performs best under the chosen implementation and risk criteria as well as the investor's expectations about the market's future behavior.

Where Does Delta Hedging Stand?

Key message

Delta hedging naturally belongs to this candidate class.

- If Delta hedging is truly optimal under real market conditions, the optimisation procedure will reveal it.
- If it is not optimal, then a more robust feasible discrete-time rule will emerge.

Takeaway

The objective is not to reject classical hedging rules a priori, but to test them inside a broader implementable framework and retain only what remains optimal in practice.

Two Formal Definitions

Definition (Fair Selling Price of a Contract)

Let H be the payoff of a financial contract at maturity. A number $p \in \mathbb{R}$ is called a *fair selling price* of the contract if it is determined through hedging techniques that are *feasible*, that is, implementable in the actual market under the admissible trading constraints.

More precisely, a fair selling price is obtained by considering feasible hedging strategies for the two parties and selecting the price p so that a prescribed quantity, such as a residual downside risk, a hedging-loss functional, or another risk criterion, is equalized between seller and buyer after optimization over the corresponding admissible strategy classes.

If, in addition, this fair price p is arbitrage-free, then this notion becomes stronger, since the price is not only balanced with respect to feasible hedging considerations, but is also consistent with the no-arbitrage restrictions of the market.

Therefore, a fair selling price should be interpreted not necessarily as the actual transaction price, but rather as a rational benchmark for negotiation between the two sides.

Predictive Valuation Map I

Definition (Predictive Valuation Map)

Let $\mathcal{X}$ be a set of market-state variables and predictive features, such as prices of related assets, moneyness, time to maturity, trading volume, volatility indicators, liquidity variables, order-flow information, news analytics, and other relevant market signals. A *predictive valuation map* is a function

$$V : \mathcal{X} \to \mathbb{R},$$

such that, for each $x \in \mathcal{X}$, the quantity $V(x)$ represents a prediction of the market value of a given asset or financial contract under the market conditions encoded by x.

Predictive Valuation Map II

Interpretation

The role of V is predictive rather than normative: it is designed to forecast the values of traded assets and standard derivatives, such as vanilla calls and puts, and its outputs may be used in portfolio construction, dynamic hedging, or other trading procedures.

In this sense, Black–Scholes-type methods, binomial methods, and their extensions may also be viewed only as *valuation predictive maps*. They provide parametric valuation rules based on selected state variables and modelling assumptions, but they should not be interpreted as literal fair-value mechanisms whenever their associated hedging logic relies on idealized continuous trading, finite-state structures, or other assumptions that are not fully feasible in actual markets.

For exotic options, the direct construction of such a predictive valuation map is often of limited practical relevance, since sufficiently rich and reliable market data for systematic training are usually unavailable.

Fair Price Is Not Market Price

- If you want to construct a method for computing some kind of a fair price, do not try to fit parameters so that the output matches the observed market prices.
- The world is not fair. Moreover, the market does not know your method, nor does it think according to it.
- Therefore, any fair price notion should be understood as a reference benchmark, not necessarily as the transaction price.
- If, instead, the goal is to build a predictive mechanism, then one should develop a valuation-predictive map and calibrate it appropriately to observed market prices, selecting the parameters that shape option values. Nevertheless, a volatility-smile-type effect will typically remain, because the predictive map is not trained on all the factors that influence prices. Since the Black–Scholes-type models should be viewed merely as valuation-predictive maps, the volatility smile signals will remain indicating that the model does not capture the full set of relevant pricing variables.

Historical Background I

The mathematical study of option pricing is usually traced back to Louis Bachelier (1900), who gave the first systematic mathematical treatment of a speculation / option-pricing problem. The basic idea was that a possible price of the option is the $\mathbb{E}(g(S_T))$ where $g(\cdot)$ is the payoff of the option.

A decisive breakthrough came in 1973 with the Black–Scholes–Merton theory. In that framework, option valuation became directly linked to dynamic hedging in continuous time.

Its most notable implication was the idealized possibility of perfect replication in a frictionless market with continuous rebalancing. A recent paper by Cont and Vuletić explores this idea by seeking the dynamic strategy in discrete time that comes as close as possible to a given hedge.

For model-based methods built on the principle of perfect replication, calibration alone cannot yield results that are genuinely implementable in practice; the volatility smile is one indication of this limitation. On the one hand, the corresponding hedging strategy is not feasible in real

Historical Background II

markets. On the other hand, there exist more effective predictive pricing approaches—such as the valuation maps in [9]—which incorporate a much richer set of inputs, including, for example, news analytics, in order to forecast future prices.

Later, in 1979, the Cox–Ross–Rubinstein binomial model reformulated the pricing and hedging problem in discrete time. This made the logic of replication especially transparent, but within a finite-state tree for the underlying asset price. See chapter 9 of *Financial Engineering with Python* for the limitation of the Binomial model.

Historical Background III

Historical lesson

The classical development moved from (see [1, 2, 3, 4]):

- the first mathematical option-pricing attempt,
- to continuous-time replication,
- and then to discrete-time replication.

The present viewpoint preserves the central idea of hedging, but works directly in discrete time and avoids the restriction that the stock price can take only finitely many terminal values.

Instead of assuming a finite tree for the terminal stock price, we only impose the minimal condition

$$S_T \geq 0.$$

Within this framework, one can construct:

- static hedging portfolios (see [6], [7])
- dynamic hedging strategies (see [8], [9])
- and corresponding notions of a fair option price.

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 γ_i, δ_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_{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)

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

Definition

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

Lower bound Y_{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

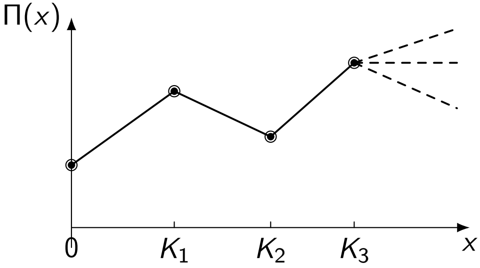

Figure 1: Illustration of a profit function $\Pi(x)$ on $x > 0$ with finitely many affine branches and grid nodes (kinks) at $S = \{0, K_1, K_2, K_3\}$. Since Π is affine between consecutive nodes, its minimum over $(0, \infty)$ is attained by checking the node values, including the boundary point $x = 0$, 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.

How Additional Options Tighten Price Bounds

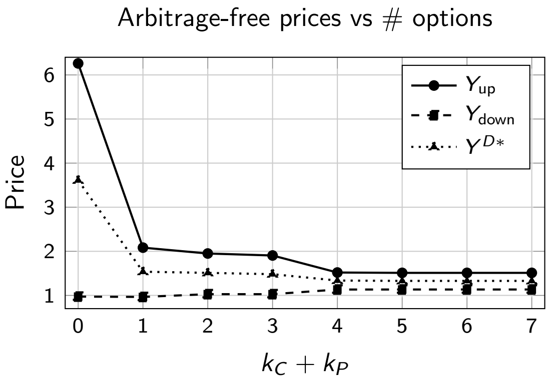

Arbitrage-free prices as a function of $(k_C + k_P)$ (B/W).

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).

Pricing

This suggests that the arbitrage-free interval may be regarded as a pricing technique.

Classical Model-Based Option Pricing I

Classical option pricing approaches, such as the Black–Scholes framework, start from a *specific model* for the dynamics of the underlying asset.

For instance, in the Black–Scholes model one assumes that the stock price follows a geometric Brownian motion:

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

The option valuation rule is then derived from this assumed price dynamics.

Main observation

When one uses a model-based option pricing formula, one also commits to a specific asset pricing model and to the structural assumptions behind it.

Thus, model-based pricing formulas are tied to assumptions such as:

Classical Model-Based Option Pricing II

- a prescribed stochastic law for the underlying asset,
- a specific parametric structure,
- and an idealized market environment supporting valuation.

Moreover, their hedging interpretation is based on replication arguments that rely on

- continuous-time rebalancing,
- frictionless trading,
- unlimited liquidity,
- and absence of transaction costs.

Therefore, the associated hedging portfolios they do not directly provide feasible hedging rules in realistic discrete-time markets.

Why We Need a Data-Driven Valuation Mechanism I

The previous discussion leads to a natural conclusion.

If the goal is to construct a practically useful dynamic trading strategy, then option valuation should not rely exclusively on

- a fixed stochastic model for the underlying asset,
- or on a historical calibration of the stock price process alone.

Instead, the pricing mechanism should be able to use information that is directly relevant to the option market while the hedging strategy should be either static or in discrete time.

Main principle

To obtain logically consistent option valuations, one should take into account, in some form, the information contained in the available traded options.

Dynamic Portfolio Construction I

We consider a portfolio consisting of

- one risky asset,
- one bank account,
- call and put options with strikes $K_1 < \cdots < K_n$,

all having the same expiration date T.

At the initial time $t_0 = 0$, we allocate the initial capital V according to

$$a_0 S_0 + b_0 + \sum_{i=1}^{n} \gamma_i \, C(K_i) + \sum_{i=1}^{n} \delta_i \, P(K_i) = V.$$

The key question is the following:

Dynamic Portfolio Construction II

Main problem

How should we update the portfolio composition at the rebalancing dates

$$t_1 < t_2 < \cdots < t_{N-1}\,?$$

In other words, we seek a dynamic rule that determines the new holdings in the stock, the bank account, and the options at each trading date.

Dynamic Portfolio Construction III

Bank account or Bonds?

In a discrete-time rebalancing framework, the non-risky asset is modeled by the money-market account. A fixed-maturity bond is tied to a specific terminal date and therefore cannot play the role of a common risk-free asset over all rebalancing periods.
This also justifies the use of the same interest rate for both investment and borrowing, since both are represented by the same bank account process: positive positions correspond to lending, whereas negative positions correspond to borrowing.

Scenario Selection for the Underlying Asset I

The methodology that we propose is based on the selection of future price scenarios for the underlying asset in its simplest form (see [9]), namely paths of the form

$$S_{t_1}(\omega),\ S_{t_2}(\omega),\ \ldots,\ S_{t_N}(\omega).$$

These scenarios represent possible future evolutions of the asset price over the trading dates

$$t_1 < t_2 < \cdots < t_N.$$

Main point

The choice of these paths is of a purely predictive character. Therefore, it is necessarily subjective.

Scenario Selection for the Underlying Asset II

Indeed, the methodology does not assume that there exists a unique “correct” future dynamics for the underlying asset.

Instead, one selects a family of plausible scenarios, according to

- the investor’s market view,
- the modeller’s predictive methodology,
- or statistical and econometric evidence.

Thus, subjectivity enters the framework through the scenario generation step, and not through the imposition of a universal pricing formula.

Possible Constructions of Scenario Paths I

There are many possible ways to construct the scenario paths

$$S_{t_1}(\omega),\ S_{t_2}(\omega),\ \ldots,\ S_{t_N}(\omega).$$

For example, one may generate them from a specific asset-price model, such as a geometric Brownian motion.

In that case, the scenarios are produced according to

$$S_{t_{k+1}}(\omega) = S_{t_k}(\omega)\exp\Big(\mu\Delta t + \sigma\sqrt{\Delta t}\, Z_{k+1}(\omega)\Big),$$

where (Z_k) is a sequence of standard normal shocks.

More generally, one may combine scenarios coming from several different models.

For instance, one may decide that:

Possible Constructions of Scenario Paths II

- $x\%$ of the generated paths come from a first model,
- $y\%$ of the paths come from a second model,
- and the remaining scenarios come from other predictive mechanisms.

Interpretation

Hence, the scenario family may be generated from a mixture of predictive models, rather than from a single universally imposed dynamics.

This makes the framework flexible: it allows the investor to incorporate different market views, stress scenarios, and alternative predictive structures into the dynamic trading procedure.

When Dynamic Trading in Options is Allowed I

Suppose that, at the rebalancing dates

$$t_1, \dots, t_{N-1},$$

the investor is allowed not only to rebalance the stock and the bank account, but also to buy or sell options.

In that case, the portfolio construction requires information about the possible future values of the traded options at the rebalancing dates.

Main implication

If options can be bought or sold dynamically, then one needs a predictive mechanism for their future market values.

When Dynamic Trading in Options is Allowed II

Indeed, at time t_k, the investor must decide the new positions in calls and puts on the basis of their expected values at that date.

Therefore, the dynamic strategy must include a valuation component capable of producing approximations such as

$$C_k(K_i) \approx \widehat{F}^{\mathrm{call}}(X_{k,i}), \qquad P_k(K_i) \approx \widehat{F}^{\mathrm{put}}(X_{k,i}).$$

Without such a predictive valuation mechanism, dynamic trading in options cannot be implemented in a coherent way.

There is, however, an alternative framework (see [8]).

Instead of allowing repeated purchases and sales of options at the rebalancing dates, one may equip the portfolio initially with options that are only bought and then held.

If these contracts are of American type, then at each rebalancing date the investor may decide whether to exercise some of them.

Alternative viewpoint

In this case, the dynamic decision is no longer how many options to buy or sell, but rather how many and which options should be exercised at each rebalancing date.

Thus, at time t_k, the control variables may include:

- the stock position,

An Alternative: Buy-Only Option Positions with Exercise Decisions II

- the bank account position,
- and the number of American-style option contracts to be exercised.

This formulation avoids the need to predict resale prices of options at every date, but replaces it with an exercise policy problem.

Hence, the methodology may proceed in two different ways:

- either through predictive valuation and dynamic trading of options (see [9]),
- or through initial option positions together with dynamic exercise decisions (see [8]).

Below, we describe (see [9]) the first methodology given that we have construct suitable valuation maps for the call and put options.

Portfolio Value at Time t_k I

At time t_k, the portfolio value is

$$V_k = a_k S_{t_k} + b_k + \sum_{i=1}^{n} \gamma_i^k \, C_k(K_i) + \sum_{i=1}^{n} \delta_i^k \, P_k(K_i),$$

where $C_k(K_i)$ and $P_k(K_i)$ denote the time-t_k values of the call and put options with strike K_i.

Valuation Maps for the Options I

The portfolio value at time t_k depends on the current values of the calls and puts.

Since these values are not known in advance, we estimate them by means of valuation maps.

More precisely,

$$C_k(K_i) \approx \widehat{F}^{\mathrm{call}}(X_{k,i}), \qquad P_k(K_i) \approx \widehat{F}^{\mathrm{put}}(X_{k,i}).$$

Therefore,

$$V_k \approx a_k S_{t_k} + b_k + \sum_{i=1}^{n} \gamma_i^k \, \widehat{F}^{\mathrm{call}}(X_{k,i}) + \sum_{i=1}^{n} \delta_i^k \, \widehat{F}^{\mathrm{put}}(X_{k,i}).$$

Learning the Rebalancing Rules from Selected Paths I

We now turn to the portfolio positions

$$a_k, \qquad \gamma_i^k, \qquad \delta_i^k.$$

These quantities are not chosen arbitrarily. We approximate them by sigmoidal functions with unknown parameters.

The parameters of these functions are estimated from a set of selected price paths.

In this way, the dynamic strategy is learned from data:

- we select representative paths,
- we fit the parameters of the sigmoidal rules,
- we use the fitted rules to rebalance the portfolio over time.

Parametric Approximation of the Portfolio Rules I

More precisely, we write

$$a_k \approx a_k^{\theta}, \qquad \gamma_i^k \approx \gamma_i^{\theta,k}, \qquad \delta_i^k \approx \delta_i^{\theta,k},$$

where θ is a vector of unknown parameters.

The parameter vector θ is computed using the selected paths and using criteria like in [8].

Mean–CVaR Optimisation

$$L_\theta := -\Pi_\theta(S_T), \qquad J(\theta) = \mathbb{E}[\Pi_\theta(S_T)] - \lambda\,\mathrm{CVaR}_\alpha(L_\theta), \qquad \lambda > 0.$$

- Maximise expected hedging performance
- Penalise tail risk of residual losses
- Optimise over admissible parameters:

$$\sup_{\theta\in\Theta} J(\theta)$$

Theorem 15

If Θ is nonempty and compact, then there exists $\theta^* \in \Theta$ such that

$$J(\theta^*) = \sup_{\theta\in\Theta} J(\theta).$$

Dynamic Hedging of Options I

By inserting the payoff function of the claim into the terminal performance equation, we obtain a dynamic hedging scheme for that derivative.

Therefore, the same framework can be used for:

- standard options,
- exotic options,
- and, in general, any contingent claim.

This gives a flexible, implementable, and data-driven approach to dynamic hedging.

Main implication

A fair price for the option can be defined via the equality of the residual downside risks of the two sides. If the arbitrage-free interval is nonempty, then, under suitable assumptions, this fair price lies inside that interval.

Dynamic Hedging of Options II

Main contribution

The proposed framework determines a fair price that is coherent with the prices of the other traded options. At the same time, it delivers a feasible dynamic hedging strategy, and therefore combines valuation with practical implementation.

Unified modelling viewpoint

By combining scenario selection with sigmoidal approximations, the proposed methodology can incorporate diffusion models, jump models, stochastic-volatility models, Lévy-type dynamics, and mixtures of predictive mechanisms. In this sense, it provides a unified discrete-time framework for approximating pricing and hedging rules generated by a broad class of stochastic models.

Dynamic Hedging of Options III

Interpretation

The proposed viewpoint may be regarded as a natural extension of classical model-based option pricing frameworks to the setting of feasible discrete-time hedging.

Related literature

This viewpoint is closely related to recent work (see [5]) on data-driven hedging with generative models, where forward-looking scenarios are combined with optimisation in order to construct dynamic self-financing hedging strategies. More generally, such scenario-based methodologies may also be used for the design of trading strategies under alternative criteria. However, the authors do not address the issue of defining a fair price for a general path-dependent option, despite the practical importance of such a notion at the negotiation stage.

Toy Example: A European Call with Stock–Bank Hedge I

We illustrate the methodology on a simple non path-dependent claim, namely a European call option with payoff

$$H = (S_T - K)^+.$$

We consider trading dates

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

and a self-financing hedging portfolio consisting only of

- the underlying asset,
- and the bank account.

Toy Example: A European Call with Stock–Bank Hedge II

Hence, at time t_k the portfolio value is

$$V_k = a_k S_{t_k} + b_k,$$

where a_k is the number of shares and b_k is the amount invested in the bank account.

Key simplification

In this toy example no intermediate option valuation map is needed, since the hedging portfolio contains only the stock and the bank account.

A Data-Driven Rebalancing Rule I

We assume that the stock position is determined only by the current stock price:

$$a_k = a_k^{\theta}(S_{t_k}), \qquad k = 0, 1, \ldots, N-1,$$

where θ is a vector of parameters to be learned from selected scenarios.
A convenient choice is a sigmoidal rule of the form

$$a_k^{\theta}(S) = \alpha_k + \beta_k \frac{1}{1 + e^{-\gamma_k (S - c_k)}}.$$

Thus, the input of the rebalancing rule is one-dimensional:

$$\text{input} = S_{t_k}.$$

If the call is sold at price p, then the initial portfolio is

$$V_0 = p, \qquad b_0 = p - a_0^{\theta}(S_0) S_0.$$

A Data-Driven Rebalancing Rule II

Interpretation

The dynamic hedging strategy is entirely determined by the current value of the underlying asset, without using path-dependent state variables or additional option positions.

Self-Financing Dynamics Along the Selected Paths I

Let

$$R_k = \frac{B_{t_{k+1}}}{B_{t_k}}$$

denote the bank account growth factor from t_k to t_{k+1}.
For each selected scenario

$$S_{t_0}^{(m)}, S_{t_1}^{(m)}, \ldots, S_{t_N}^{(m)}, \qquad m = 1, \ldots, M,$$

the portfolio evolves according to

$$V_{k+1}^{(m)} = a_k^{\theta}\left(S_{t_k}^{(m)}\right) S_{t_{k+1}}^{(m)} + R_k\, b_k^{(m)},$$

and after rebalancing we set

$$b_{k+1}^{(m)} = V_{k+1}^{(m)} - a_{k+1}^{\theta}\left(S_{t_{k+1}}^{(m)}\right) S_{t_{k+1}}^{(m)}.$$

Self-Financing Dynamics Along the Selected Paths II

At maturity, the hedging error is

$$\varepsilon^{(m)}(p,\theta) = V_N^{(m)} - (S_T^{(m)} - K)^+.$$

The parameters θ are estimated from the selected paths by optimising a risk-sensitive criterion based on the terminal hedging error.
For example, for a fixed initial price p, one may solve

$$\sup_{\theta \in \Theta} \left[\frac{1}{M} \sum_{m=1}^{M} \varepsilon^{(m)}(p, \theta) - \lambda \, \mathrm{CVaR}_{\alpha}\Big(\big((S_T - K)^+ - V_N \big)^+ \Big) \right].$$

This produces:

- a feasible dynamic hedging rule $a_k^{\theta}(S_{t_k})$,
- and a residual risk level associated with the candidate price p.

Valuation viewpoint

By varying p, one may define a fair price through the equality of the optimised residual downside risks of the two sides.

A Fair Price Within the Classical No-Arbitrage Bounds I

For the European call toy example, the proposed methodology yields a fair price Y^* defined through equality of the minimal residual downside risks of the two sides.
Under suitable assumptions, this fair price lies in the arbitrage-free interval. Since here the market consists only of the underlying asset and the bank account, the relevant no-arbitrage bounds are the classical ones:

$$\max\{0, S_0 - Ke^{-rT}\} \leq Y^* \leq S_0.$$

Interpretation

Hence the method produces a valuation which is not only data-driven and compatible with feasible discrete-time hedging, but also consistent with the standard no-arbitrage restrictions for a European call.

Implicit Dependence on the Option Parameters I

In the toy example, the stock position is written in the simple form

$$a_k^* = A_k^*(S_{t_k}).$$

Thus, the current stock price is the only explicit input of the trading rule.

However, the fitted map A_k^* depends implicitly on

- the specific option payoff,
- the strike K,
- the maturity T,
- the trading dates,
- and the scenario-generation mechanism used in the training step.

Interpretation

Hence the rule is one-dimensional in its observable input, but option-specific through its calibration.

Relation with the Black–Scholes Viewpoint I

The scenario paths may be generated by a geometric Brownian motion.
Therefore, the toy example remains close to the Black–Scholes framework at the level of predictive scenario generation.

The essential difference appears at the valuation level: the proposed fair price is associated with an optimal feasible hedge in discrete time, whereas the Black–Scholes price is derived from idealised continuous-time replication.

Main point

Thus, the proposed valuation is more closely connected with implementable hedging decisions in realistic discrete-time trading. **The fair price $Y^\star$ is thus a fairer price than the Black-Scholes fair price!**

Relation with the Black–Scholes Viewpoint II

Choice of approximation class and dependence on scenario selection

It may be preferable not to approximate the optimal trading rule by a single sigmoidal function only, but rather by a richer class of sigmoidal-type functions, possibly including a Delta-type hedging rule among them. In that case, the optimal strategy may naturally be represented as a suitable linear combination of such candidate functions.

This also makes clear that the resulting optimal strategy depends directly on the paths used in the training procedure. Since the selection of these paths is itself predictive in nature, the final hedging rule necessarily reflects the underlying forecasting choice.

Why we do not compare separately with Delta hedging

It is not particularly meaningful to compare the proposed method separately with Delta hedging. Indeed, the Delta-hedging rule is itself of sigmoidal type and therefore belongs to the class of functions used to approximate the optimal strategy. Consequently, if Delta hedging is in fact optimal, this will be revealed by the optimisation problem itself.

Relation with the Black–Scholes Viewpoint III

Further implication

The learned rebalancing rule a_k^* can also be used to compute data-driven Greeks and an implied volatility. These quantities are more relevant for practical hedging than the corresponding Black–Scholes ones, since they are generated by a feasible discrete-time hedging portfolio.

Why This is a Toy Example

- **Only GBM scenarios:** the future stock-price paths are generated only by a geometric Brownian motion.
- **Restricted information set:** the rule depends only on basic variables such as T, σ, r, K and the current stock price, while richer inputs are ignored, such as past path information or news analytics.
- **No option-market information:** available call and put prices are not used.

Interpretation

Hence, this example stays very close to the Black–Scholes setting. Even there, the proposed methodology is more relevant in practice, because it is based on feasible discrete-time hedging strategies rather than idealised continuous-time replication.

Bounded Sigmoidal Approximation of the Hedging Rule I

In the toy example, the stock position is written in the form

$$a_k^* = A_k^*(S_{t_k}),$$

where the dependence on the strike, the maturity, and the scenario-generation mechanism is implicit through calibration.

To make this dependence explicit, we introduce the parametrised hedging map

$$A_k(S;\sigma,m,K,T) := A(t_k,S;\sigma,m,K,T), \qquad \tau_k := T - t_k.$$

Bounded Sigmoidal Approximation of the Hedging Rule II

Standing assumption

For each trading date t_k, the map

$$A_k : \mathcal{D}_k \subset \mathbb{R}^5 \to [0,1], \qquad (S, \sigma, m, K, \tau_k) \mapsto A_k(S; \sigma, m, K, T),$$

is assumed to be continuous on a compact domain $\mathcal{D}_k$.

We then approximate A_k by bounded sigmoidal functions.

Sigmoidal approximation

Let

$$\Lambda(x) := \frac{1}{1+e^{-x}}, \qquad x \in \mathbb{R}.$$

For every $\varepsilon > 0$, there exists an integer $q \geq 1$ and coefficients

$$c_0,\ c_1, \dots, c_q, \qquad b_1, \dots, b_q \in \mathbb{R},$$

together with vectors

$$u_j = (u_{j,1}, u_{j,2}, u_{j,3}, u_{j,4}, u_{j,5}) \in \mathbb{R}^5, \qquad j = 1, \dots, q,$$

such that the function

$$A_k^{(q)}(S; \sigma, m, K, T) = \Lambda\left(c_0 + \sum_{j=1}^{q} c_j\, \Lambda\big(u_{j,1}S + u_{j,2}\sigma + u_{j,3}m + u_{j,4}K + u_{j,5}\tau_k + b_j\big) \right)$$

Symbolic Data-Driven Greeks I

Once the hedging rule is represented by a smooth bounded sigmoidal approximation

$$a_k \approx A_k^{(q)}(S; \sigma, m, K, T),$$

its sensitivities can be computed symbolically.

Indeed, since

$$\Lambda'(x) = \Lambda(x)\big(1 - \Lambda(x)\big),$$

the map $A_k^{(q)}$ is infinitely differentiable with respect to all its variables. Therefore one may define the data-driven hedge sensitivities

$$\frac{\partial A_k^{(q)}}{\partial S}, \qquad \frac{\partial A_k^{(q)}}{\partial \sigma}, \qquad \frac{\partial A_k^{(q)}}{\partial m}, \qquad \frac{\partial A_k^{(q)}}{\partial K}, \qquad \frac{\partial A_k^{(q)}}{\partial t}.$$

Symbolic Data-Driven Greeks II

If, moreover, the data-driven price surface is denoted by

$$Y^\star = Y^\star(t, S; \sigma, m, K, T),$$

and is sufficiently smooth, then the corresponding Greeks are defined by

$$\Delta^{DD} = \frac{\partial Y^\star}{\partial S}, \qquad \Gamma^{DD} = \frac{\partial^2 Y^\star}{\partial S^2}, \qquad \mathcal{V}^{DD} = \frac{\partial Y^\star}{\partial \sigma}, \qquad \Theta^{DD} = \frac{\partial Y^\star}{\partial t}.$$

A Price Surface Associated with the Hedging Rule I

Assume that the learned hedging rule is represented by

$$A = A(t, S; \sigma, m, K, T).$$

A price surface associated with this hedge map may be defined by

$$Y^\star(t, S; \sigma, m, K, T) = \int_{S_{\mathrm{ref}}}^{S} A(t, u; \sigma, m, K, T)\, du + c(t, \sigma, m, K, T),$$

where S_{ref} is a fixed reference level of the underlying price, and $c(t, \sigma, m, K, T)$ is determined by a normalisation or boundary condition.
Then, by construction,

$$\frac{\partial Y^\star}{\partial S}(t, S; \sigma, m, K, T) = A(t, S; \sigma, m, K, T).$$

A Price Surface Associated with the Hedging Rule II

Therefore, if $Y^\star$ is sufficiently smooth, the corresponding Greeks are

$$\Delta^{DD} = \frac{\partial Y^\star}{\partial S} = A, \qquad \Gamma^{DD} = \frac{\partial^2 Y^\star}{\partial S^2}, \qquad \mathcal{V}^{DD} = \frac{\partial Y^\star}{\partial \sigma}, \qquad \Theta^{DD} = \frac{\partial Y^\star}{\partial t}.$$

Interpretation

The learned hedge map A plays the role of a data-driven Delta, while $Y^\star$ is an associated price surface obtained as an S-primitive of A.

Toy Example: Fair Price and Market Valuation

- The fair price $Y^\star$ may still display a behaviour analogous to the classical volatility smile.
- This is because $Y^\star$ is a fair-value benchmark, not necessarily the market selling price of the option.
- Hence, fair value and observed market quote need not coincide.
- If one seeks a valuation map more directly aligned with market prices, one may use the data-driven approach of [9].
- There, the valuation map is learned from real option data, so pricing becomes predictive and data-driven.

Fair Value as a Negotiation Benchmark I

- In our framework, the fair value is linked to feasible hedging strategies, not to idealized continuous-time replication.
- More precisely, it is defined through the equality of the minimal residual downside risks of the writer and the buyer.
- Therefore, this fair value should be interpreted as a bilateral reference point for negotiation.
- **It is not necessarily the actual market transaction price**.
- For this reason, even in a toy example close in spirit to Black–Scholes, the fair-value benchmark may still exhibit a smile-like effect.

Fair Value as a Negotiation Benchmark II

Why a fair price is still useful

The seller naturally aims to sell at a high price, and the buyer naturally aims to buy at a low price. However, without some notion of a fair or reference value, neither side can meaningfully assess what "high" or "low" actually means.

In this sense, the fair price is not the transaction price itself. Its role is to provide each side with a scenario-based benchmark relative to which an expensive selling price or a cheap buying price can be judged. The crucial point is that any notion of fair price must be grounded in feasibility.

Valuation Maps: A Different Objective

- A valuation map has a different role: it aims to predict market option values rather than define a fair bilateral benchmark grounded on feasible hedging strategies.
- It is trained on real option data and should incorporate as many relevant factors as possible:
  - spot and moneyness,
  - time to maturity,
  - **realised** volatility,
  - liquidity variables,
  - order-flow or bid–ask information,
  - news and other market signals.
- In this case, the smile is not a structural defect of the method.
- Instead, the valuation map can absorb and reflect the information already embedded in market prices.

Sensitivities and Greeks

- If the objective is to compute sensitivities, Greeks, or implied volatilities, then the natural object is a valuation map or an associated smooth price surface.
- Such a map is built as a function of relevant state variables and can therefore be differentiated with respect to spot, volatility, strike, or time.
- By contrast, the fair value in our framework is mainly a bilateral benchmark linked to feasible hedging strategies of the two sides.
- Hence, fair value is more suitable for negotiation and reference pricing, whereas valuation maps are more suitable for sensitivity analysis and prediction.

Black-Scholes model: A valuation map

In this sense, Black–Scholes is used in practice less as a literal fair-value mechanism and more as a benchmark valuation map, precisely because its original hedging interpretation relies on continuous rebalancing, which is not feasible in actual markets.

Fair Value Method vs. Valuation Map

- There is an important difference between the fair value method and the valuation map in the way their parameters are trained.
- In the case of the valuation map, the training procedure is based only on **historical market data**.
- By contrast, in the fair value method, the parameter estimation is based on **possible future scenarios** for the evolution of the underlying asset.
- These scenarios may come from:
  - historical data,
  - stress scenarios that we fear may occur,
  - favourable scenarios that we hope may occur,
  - stochastic models such as **GBM**, **Heston**, etc.
- Hence, the valuation map is essentially **backward-looking**, while the fair value method is **forward-looking** and **scenario-based**.

Valuation Map vs. Fair Value

- Valuation maps are primarily needed for the vanilla options that are traded dynamically inside the hedging portfolio.
- Their role is to provide intermediate option values along the trading dates, so that rebalancing remains feasible.
- For a complex contract, however, a direct valuation map is not always the natural object to estimate.
- Such an approach would require historical observations of market prices for that specific contract, and for non-standard derivatives these data may be scarce or unreliable.
- In this case, a fair value derived from feasible hedging strategies is often more useful before trading, because it provides a practical benchmark for negotiation.

References I

Below you can find the relevant references. Item [13] contains the Python code used primarily in [6].

[1] L. Bachelier.
Théorie de la spéculation.
Annales Scientifiques de l'École Normale Supérieure, 17:21–86, 1900.

[2] F. Black and M. Scholes.
The pricing of options and corporate liabilities.
Journal of Political Economy, 81(3):637–654, 1973.

[3] R. C. Merton.
Theory of rational option pricing.
Bell Journal of Economics and Management Science, 4(1):141–183, 1973.

[4] J. C. Cox, S. A. Ross, and M. Rubinstein.
Option pricing: A simplified approach.
Journal of Financial Economics, 7(3):229–263, 1979.

[5] R. Cont and M. Vuletić.
Data-driven hedging with generative models.
Annals of Operations Research, 2025.
doi:10.1007/s10479-025-06867-3.

[6] 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.

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

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

[9] N. Halidias. *Feasible Dynamic Hedging and Data-Driven Valuation of Options in Discrete Time.*
Preprint (ResearchGate), 2026.

[10] N. Halidias. *Calibrated Predictive Regions for Static Hedging of Piecewise-Affine Claims with Option Portfolios.*
Preprint (ResearchGate), 2026.

[11] N. Halidias.
Financial Engineering with Python.
World Scientific Publishing, to appear.

References II

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

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