Financial Mathematics / Financial Engineering
Assignment Topics

Topic 1: Deterministic Arbitrage (price-forming no-arbitrage)

1.1 Definition (statewise / model-free). State a precise definition of deterministic arbitrage in a frictionless market with a stock, a bank account, and European calls/puts with common maturity T. Use a terminal payoff function $\Pi(x)$ with $x = S_T$ and specify the conditions on:

$$\Pi(x) \geq 0 \quad \forall x \geq 0, \qquad \Pi(x) > 0 \quad \text{for some } x \geq 0,$$

together with the requirement on the initial cost (non-positive net investment). Explain why this notion is *price-forming*.

1.2 Example A: Put–Call Parity arbitrage. Assume no dividends, constant rate r, maturity T, strike K, and observed prices S_0, $C(K)$, $P(K)$.

(a) Write the parity relation and interpret it as an equality of two portfolios with identical payoff at T.

(b) Provide a numerical example where parity fails (choose parameters and quotes), and construct an explicit arbitrage strategy (positions in stock, bond, call, put).

(c) Compute the time-0 cashflow and the terminal payoff in each region $S_T \leq K$ and $S_T > K$.

1.3 Example B: Upper/lower bounds for calls and puts. Assume a non-dividend-paying underlying and frictionless trading.

(a) State the classic no-arbitrage bounds for $C(K)$ and $P(K)$.

(b) Construct an explicit deterministic arbitrage strategy for *each* violated bound (e.g., $C > S_0$, $P > Ke^{-rT}$, etc.).

(c) Provide a numerical table of strikes and quotes and identify any violations.

1.4 Example C: Cross-strike constraints (monotonicity and convexity). Fix a maturity T and consider call prices across strikes.

(a) State and prove (via static portfolios) that call prices are nonincreasing in K.

(b) State a convexity condition in strike (finite-difference/butterfly form) and interpret it economically.

(c) Give a numerical counterexample (fabricated or from data) that violates monotonicity or convexity and build the corresponding arbitrage portfolio.

1.5 Arbitrage-free interval and "price discipline". Explain what an *arbitrage-free interval* for an option price means in a given market, and how additional traded strikes tighten this interval. Discuss how deterministic arbitrage restrictions shape the option surface (parity, monotonicity, convexity).

1.6 Other notions of arbitrage in the literature (and why they are not "real" arbitrage). Search the literature and identify at least **three** concepts besides deterministic arbitrage, such as:

- statistical arbitrage / arbitrage in expectation,
- approximate or asymptotic arbitrage,
- "no free lunch" type conditions (vanishing risk / bounded risk variants).

For each notion:

(a) give the definition (with a citation),

(b) explain precisely what is weakened relative to deterministic arbitrage (statewise dominance, horizon, limit operations, probability model dependence),

(c) argue why it is not price-forming for a *single* option quote (i.e., why it does not produce hard model-free bounds).

Topic 2: Arbitrage-Free Bounds and Sanity Checks on Real Option Chains

Download a call/put option chain for a fixed maturity T and:

2.1 Verify numerically the bounds for calls and puts across strikes.

2.2 Test put–call parity across strikes and report the maximum deviation.

2.3 Test monotonicity and convexity in strike; visualize $C(K)$ and highlight violations (if any).

2.4 Discuss how bid–ask spreads and transaction costs affect whether an apparent violation is executable.

Topic 3: Detecting Arbitrage via Linear Programming

Formulate arbitrage detection as an optimization problem over static portfolios (stock, cash, multiple calls/puts):

3.1 Define a portfolio payoff $\Pi(x)$ and a budget constraint using observed option quotes.

3.2 Discretize the state space $x \in [0, \infty)$ on a sufficiently fine grid and implement the constraints.

3.3 Solve the linear program and interpret the optimal value (arbitrage vs. no-arbitrage within bounds).

3.4 Study the effect of position limits on whether arbitrage can be detected/exploited.

Topic 4: Model-Free Super/Sub-Hedging and the Arbitrage-Free Price Interval

Using traded strikes $K_1 < \cdots < K_n$ at the same maturity:

4.1 Define writer super-hedging and buyer sub-hedging problems for a piecewise-linear payoff $f(S_T)$.

4.2 Compute the model-free arbitrage-free interval $[Y_{\text{buyer}}, Y_{\text{writer}}]$.

4.3 Explain how adding more strikes changes (tightens) the interval.

4.4 Apply the method to a new payoff (e.g., a spread, a capped payoff, or a piecewise-linear exotic).

Topic 5: Model Prices vs No-Arbitrage Discipline

Compare a model-based price (e.g., binomial tree or Black–Scholes with historical parameters) with model-free bounds:

5.1 Compute the model price for a European call/put for chosen (S_0, K, T, r, σ).

5.2 Compute the model-free arbitrage-free interval implied by observed option quotes.

5.3 Check whether the model price lies inside the interval; if not, explain what this means.

5.4 Discuss why replication-based arguments may fail in practice (discrete trading, costs, impact) and why no-arbitrage bounds remain robust.

Topic 6: Markowitz (Mean–Variance) and Risk Management Beyond Diversification

6.1 Markowitz problem formulation and interpretation. Consider n risky assets with expected returns $m \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n\times n}$. Let $w \in \mathbb{R}^n$ be portfolio weights.

(a) Define portfolio mean and variance:

$$\mu_p(w) = w^\top m, \qquad \sigma_p^2(w) = w^\top \Sigma w.$$

(b) State the classical Markowitz optimization problem (target-return form):

$$\min_{w\in\mathbb{R}^n} \; w^\top \Sigma w \quad \text{s.t.} \quad \mathbf{1}^\top w = 1, \;\; w^\top m = \mu_0,$$

and explain the economic meaning of each constraint.

(c) Explain why covariance (co-movement) is central: diversification reduces risk through imperfect correlation.

6.2 Computational part (Python): Efficient frontier with and without short selling. Using either a small toy dataset (e.g., $n = 3$ assets) or real historical data:

(a) Estimate m and Σ from returns and compute the Global Minimum-Variance Portfolio (GMVP).

(b) Compute the efficient frontier by sweeping μ_0 over a grid and solving the Markowitz problem.

(c) Repeat under the constraint $w_i \geq 0$ (no short selling) and compare frontiers (risk/return trade-offs).

(d) Visualize and briefly comment on the geometry and the role of constraints.

6.3 Why diversification is not the only risk-management strategy. Provide a critical discussion showing that mean–variance diversification alone can be insufficient:

(a) Discuss estimation/forecast risk: m and Σ are typically inferred from historical data and may be unstable under structural breaks or regime changes.

(b) Explain why variance may not capture tail risk (downside asymmetry, fat tails).

(c) Explain that restricting the action space to underlyings (and possibly a risk-free asset) omits powerful risk-management tools, notably derivatives-based hedging and payoff engineering.

6.4 Alternative strategy 1: Derivatives-based hedging (option overlay) and payoff engineering. Choose a baseline portfolio (e.g., a Markowitz portfolio in one equity index/stock) and add an option overlay.

(a) Construct either a *protective put* or a *collar* on the main risky exposure.

(b) Write the terminal P&L as a function of S_T and show explicitly how the option overlay creates a downside floor (bounded loss).

(c) Compare the unhedged vs hedged portfolio using at least two metrics: worst-case loss over a grid of S_T, and one statistical measure (e.g., VaR or CVaR).

6.5 Alternative strategy 2: Robust (distribution-free) risk constraints beyond variance. Formulate a portfolio-design problem that enforces a statewise loss bound:

$$\Pi(x) \geq -D \quad \text{for all terminal states } x \text{ in a test set (grid/scenarios).}$$

(a) Explain how to enforce the constraint numerically (finite grid of terminal prices plus appropriate tail constraints if needed).

(b) Explain why this is more robust to model misspecification than purely distributional constraints.

(c) Discuss how this framework generalizes classical Markowitz (broader instruments + explicit downside control).

6.6 Alternative strategy 3: Tail-risk measures (VaR / CVaR) as objective or constraint.

(a) Compute historical and/or parametric VaR at two confidence levels (e.g., 95% and 99%) for two portfolios.

(b) Compute/estimate CVaR and give an example where two portfolios have similar variance but different tail risk.

(c) Discuss why "variance as risk" can be inadequate for practical risk management.

6.7 Synthesis (1–2 pages). Write a concluding section explaining clearly:

- what diversification achieves in the Markowitz framework (role of Σ),
- what it fails to control (tail events, forecast instability, regime shifts, limited instrument set),
- and how hedging with options, robust statewise constraints, and tail-risk measures complement diversification as a broader risk-management toolkit.
- **Why separate forecasting from portfolio construction? (pipeline advantage)** Explain the practical advantage of separating the *forecasting layer* (prediction) from the *construction layer* (portfolio design/optimization). Your answer must address all of the following:

(a) **Clear roles and interfaces.** State precisely what each layer produces/consumes. Give at least two examples of admissible forecasting outputs (e.g., prediction set, stress-scenario set, or an estimated distribution), and explain how the construction layer turns these inputs into a tradable portfolio via an optimization problem.

(b) **Modularity and upgradability.** Argue why improving the forecasting layer should *mechanically* improve the resulting portfolios without changing the portfolio engine. Explain what it means for the interface to be "clean".

(c) **Robustness and model-risk reduction.** Explain why a prediction-set / scenario-based interface can be more robust than committing to a single full joint distribution. Discuss how heterogeneous information (data + judgment + stress scenarios) can be incorporated without over-committing to fragile parametric assumptions.

(d) **Auditability and failure diagnosis.** Explain why entangling prediction and portfolio choice into a single monolithic model can be fragile, and why separation makes it easier to diagnose whether poor performance came from (i) bad forecasts or (ii) poor construction (constraints/objective/implementation).

(e) **Connection to risk management.** Explain how the separation supports a disciplined workflow such as enforcing explicit risk constraints (e.g., worst-case loss bounds when possible) and only then using statistical risk measures (VaR/CVaR) when deterministic control is not feasible.

Guiding Note: What "Option Pricing & Hedging" Really Means

In this course, we interpret **option pricing and hedging** as the following *engineering* problem:

> *List (or characterize) all hedging strategies that are* ***feasible*** *in the given market and compute how much initial capital each strategy requires.*

Equivalently, for a target payoff $f(S_T)$ (e.g., a European option payoff), we ask:

- **Feasibility:** Which trading strategies (dynamic and/or static) can be implemented using the available instruments?
- **Cost:** What is the minimal initial capital needed to achieve a required hedging objective?

Why would I buy or sell an option?

An option is not only a "bet"; it is a **payoff-shaping tool**. You buy/sell an option because you want to modify your terminal P&L profile:

- **Risk reduction / insurance:** cap losses (downside protection), reduce drawdowns, control worst-case outcomes.
- **Exposure engineering:** express a view on direction, volatility, or tail events, often in a convex way.
- **Budget or constraint satisfaction:** meet regulatory or internal risk limits with a targeted payoff profile.

Is the market price "convenient" for my purpose?

Let V_0^{mkt} be the observed market price of an option with payoff $f(S_T)$.

- If you want to **buy** the option, you should compare V_0^{mkt} to the **buyer side** benchmark:

$$Y_{\text{buyer}} = \sup\{\text{initial cash you can raise by a portfolio with payoff} \le f(S_T)\}.$$

  If V_0^{mkt} is *high* relative to what you can achieve by constructing a similar payoff from traded instruments, the option may be **uneconomical** for your objective.

- If you want to **sell** (write) the option, you should compare V_0^{mkt} to the **writer side** benchmark:

$$Y_{\text{writer}} = \inf\{\text{initial cash needed for a portfolio with payoff} \ge f(S_T)\}.$$

  If V_0^{mkt} is *low* relative to the cost of a safe hedge, writing the option is **too risky/underpaid** for your objective.

Student deliverable (required in every option question)

Whenever you analyze an option quote, you must answer the following:

(a) What is the **goal** (insurance, exposure, constraint satisfaction, speculation)?

(b) What are the **feasible hedges** (static/dynamic) using the available instruments?

(c) What are the corresponding **costs** (capital required) and the resulting terminal payoffs?

(d) Given the market quote V_0^{mkt}, is the option **useful** (cost-effective) for the stated goal?

Why classic models (Black–Scholes, Binomial) do *not* answer the questions above

At the end of your report, write a short discussion (about 1 page) explaining why the classical Black–Scholes (BS) and binomial-model theories do *not* directly answer any of the practical questions posed above, namely: "What hedges are actually feasible in this market?", "How much capital do they require?", and "Is the observed market quote convenient for my specific purpose?"

Your discussion must address **all** of the following points:

(a) **Model dependence vs. market dependence.** BS/binomial produce a *model price* under specific assumptions (dynamics for S_t, frictionless trading, continuous/discrete rebalancing rules, etc.). The questions above are *market questions*: they depend on the *actual set of traded instruments*, their quotes (bid–ask), and the trading constraints.

(b) **Perfect replication is an assumption, not a market fact.** In BS (and in idealized binomial settings), the option price is pinned down because the payoff is assumed to be perfectly replicable by trading the underlying (and cash). In reality, replication may fail or be only approximate due to: discrete hedging, transaction costs, liquidity constraints, jumps, stochastic volatility, and position limits. Hence BS/binomial do not *catalog* the feasible hedges in the *given* market; they postulate a replicating strategy in an idealized model.

(c) **No-arbitrage bounds and hedging cost are *one-sided* in incomplete markets.** When the market is incomplete (or trading is constrained), there is typically an *arbitrage-free interval* rather than a unique price. BS/binomial still output a single number, but the buyer/writer questions above require:

   - the **writer's** super-hedging cost (safe hedge from above),

- the **buyer's** sub-hedging benchmark (hedge from below),
- and an assessment of whether the market quote is attractive *for a specific goal*.

A single model price is not the same thing as these market-dependent benchmarks.

(d) **Purpose and constraints are external to the model.** The decision "Should I buy or sell this option?" depends on the user's objective (insurance vs. speculation), risk limits (worst-case loss, VaR/CVaR constraints), capital constraints, and horizon. BS/binomial do not encode the user's utility, constraints, or institutional risk-management requirements; therefore they cannot decide whether a quote is "convenient" for a particular purpose.

(e) **Parameter risk and calibration risk.** Model outputs depend critically on inputs (e.g., volatility in BS, transition probabilities in binomial trees). In practice these must be estimated or calibrated, and the resulting price can vary widely. This reinforces that BS/binomial provide *conditional* answers ("if the model and parameters are correct") rather than the robust, goal-driven answers required by the questions above.

Deliverable. Conclude with a clear statement of the takeaway: *BS/binomial are valuable for producing internally consistent model prices and model hedges under strong assumptions, but they do not, by themselves, determine (i) all feasible hedges in the actual market, (ii) their true implementable costs, or (iii) whether the observed market price is suitable for a specific risk-management objective.*