Notes on Financial Risk Management*

Krish Nigam

1 Introduction

1.1 Stylized facts of asset returns

- *Little autocorrelation in returns:* daily returns are essentially uncorrelated across time, so the conditional mean is approximately constant (often set to zero).
- *Non-normality (fat tails):* return distributions are more peaked and have heavier tails than Normal, so Gaussian models underestimate extreme losses.
- *Asymmetry (negative skewness):* large negative returns occur more often than large positive ones, implying downside risk is more severe than upside gains.
- *Mean close to zero (short horizons):* at daily frequency, volatility dominates the mean, so risk modeling focuses on variance rather than expected return.
- *Volatility clustering (variance persistence):* high volatility follows high volatility (and low follows low), motivating time-varying volatility models such as GARCH.
- *Leverage effect:* volatility tends to rise after negative returns, implying asymmetric volatility response to shocks.
- *Time-varying correlation:* correlations increase in stressed markets, so constant-correlation assumptions can understate portfolio risk in downturns.
- *Conditional non-normality:* even standardized returns $z_t = R_t/\sigma_t$ remain heavy-tailed, motivating filtered historical simulation and fat-tailed shock distributions.
- *Long-horizon normality:* aggregated returns become closer to Normal at longer horizons, though volatility dynamics can still matter for long-horizon risk.

Asset returns have weak mean dynamics but strong, persistent, asymmetric, and non-Gaussian volatility and dependence, motivating GARCH/FHS, dynamic correlation, and stress testing frameworks.

1.2 Value at risk

Value-at-Risk (VaR) is a risk measure that answers the question: "given the distribution of portfolio returns, what is the loss that would be exceeded with probability p?".

For a loss random variable Loss_{t+1}, the dollar VaR ($\$VaR$) is implicitly defined by:

$$\Pr(\text{Loss}_{t+1} > \$VaR_p) = p$$

*These notes are informed by ideas and discussions shared by *Prof. Caio Almeida* at *Princeton University*.

In terms of portfolio returns (R_{PF}), this is equivalent to:

$$\Pr(R_{PF,t+1} < -VaR_p) = p$$

For example, if the 1% VaR is 2%, it means there is a 1% probability that the portfolio will lose more than 2% of its value.

A key drawback of VaR is that it ignores extreme losses; it tells you the threshold for the worst p% of outcomes but says nothing about how bad those outcomes could be.

1.3 Expected shortfall

Expected Shortfall (ES) measures the expected loss *given* that the loss exceeds the VaR.

$$\text{ES}_{p,t+1} = -E_t[R_{PF,t+1} | R_{PF,t+1} < -\text{VaR}_{p,t+1}]$$

For any continuous distribution, ES is the average of all VaR values for quantiles more extreme than p, which guarantees that $\text{ES}_p \geq \text{VaR}_p$.

ES is sensitive to the fatness of the distribution's tail; the fatter the tail (higher kurtosis), the larger the gap between ES and VaR.

For HS, ES is calculated by taking the average of all returns that are smaller than the VaR value.

1.4 Coherence of risk measures

- *Monotonicity:* If a position produces larger losses in all states, it is riskier. Risk measures must preserve this ordering.
- *Translation Invariance:* Adding a sure loss l increases risk by exactly l. Risk shifts one-for-one with deterministic cash flows.
- *Subadditivity:*
$$\rho(L_1 + L_2) \leq \rho(L_1) + \rho(L_2).$$
Diversification should not increase risk. Expected Shortfall satisfies this; VaR may fail.
- *Positive Homogeneity:*
$$\rho(\lambda L) = \lambda \rho(L).$$
Scaling a position scales its risk proportionally. This allows separation of volatility and shocks, e.g.
$$\text{VaR}^p_{R_{t+1}} = -\sigma_{t+1} F_z^{-1}(p).$$

ES satisfies all four coherence axioms, whereas VaR is not coherent since it fails subadditivity (while passing the other three).

1.5 Filtered historical simulation example

We give an example of FHS with GARCH(1,1) to estimate a 1-day ahead VaR.

We observe a time series of portfolio returns $\{R_\tau\}_{\tau=1}^m$ generated by $R_t = \sigma_t z_t$, where z_t are i.i.d. shocks with unknown distribution D_z, and σ_t follows a GARCH(1,1) volatility process.

The key idea of Filtered Historical Simulation (FHS) is to:

- use a parametric model (GARCH) to capture time-varying volatility σ_t,
- use historical data to capture the distribution of shocks z_t, without imposing normality.

The objective is to estimate VaR^p_{t+1} (*e.g.* $p = 1\%$) at the end of day t.

First, estimate the volatility model and "filter" the returns. Then, fit a GARCH(1,1) model to the historical returns $\{R_\tau\}_{\tau=1}^m$ and obtain the fitted conditional volatilities $\{\hat{\sigma}_\tau\}$.

Compute the standardized residuals (filtered shocks): $\hat{z}_\tau = \frac{R_\tau}{\hat{\sigma}_\tau}$, $\tau = 1, \ldots, m$.

Intuition: Dividing returns by their conditional volatility removes volatility clustering, so the series $\{\hat{z}_\tau\}$ is closer to i.i.d., without assuming a parametric form for D_z.

Next step is to forecast tomorrow's volatility. At the end of day t, both R_t and $\hat{\sigma}_t$ are known.

The one-step-ahead volatility forecast is obtained from the GARCH recursion:

$$\hat{\sigma}^2_{t+1} = \hat{\omega} + \hat{\alpha} R_t^2 + \hat{\beta}\hat{\sigma}_t^2.$$

Then, we simulate tomorrow's return using resampled shocks. Perform Historical Simulation on the shocks (not on raw returns):

1. Draw H values with replacement from the empirical database $\{\hat{z}_\tau\}_{\tau=1}^m$. Call them by $\hat{z}^{(h)}$, $h = 1, \ldots, H$.
2. Construct H one-day-ahead return scenarios: $R^{(h)}_{t+1} = \hat{\sigma}_{t+1}\,\hat{z}^{(h)}$, $h = 1, \ldots, H$.

Next step is to compute the 1-day VaR from simulated returns. Since losses correspond to negative returns, the one-day-ahead VaR is computed as:

$$\text{VaR}^p_{t+1} = -\text{Quantile}_p\left(\{R^{(h)}_{t+1}\}_{h=1}^H\right).$$

For $p = 1\%$, this corresponds to the negative of the 1^{st} percentile of the simulated return distribution.

2 Correlation

2.1 Portfolio correlation

While volatility captures individual asset risk, correlation ρ captures the *joint* movement of assets. It is the primary driver of portfolio diversification benefits.

For an N-asset portfolio, we must estimate N volatilities but $N(N-1)/2$ correlations. As N grows, the correlation terms dominate the portfolio variance. Portfolio return, $r_{PF,t+1} = \sum_{i=1}^n w_{i,t} r_{i,t+1}$. Even if you can model each stock's volatility, portfolio risk also depends on how they move together (covariance/correlation).

$$\sigma^2_{PF,t+1} = \sum_{i=1}^n \sum_{j=1}^n w_{i,t} w_{j,t} \sigma_{ij,t+1} = \sum_{i=1}^n \sum_{j=1}^n w_{i,t} w_{j,t} \sigma_{i,t+1} \sigma_{j,t+1} \rho_{ij,t+1}$$

Correlation can be time-varying (stylized fact) and increases during crises/financial turmoil (diversification reduces), which increases portfolio risk further. Correlation is treated as risk factor, *e.g.* $w_1 = w_2 = 0.5$, $\sigma_1 = \sigma_2 = \sigma$, if correlation $\rho = 0$ then $\sigma_{PF} = \sqrt{0.5}\sigma$; if $\rho = 0.8$ then $\sigma_{PF} = \sqrt{0.9}\sigma$.

2.2 Rolling window covariance

The simplest estimator. Calculates covariance using a fixed window of size m (*e.g.*, 250 days).

$$\sigma_{ij,t+1} = \frac{1}{m}\sum_{\tau=1}^{m} R_{i,t+1-\tau}R_{j,t+1-\tau}$$

Flaws: Equal weighting (slow reaction to recent news), ghost effects (abrupt changes when a large shock drops out of the window).

2.3 RiskMetrics covariance (EWMA)

Applies Exponentially Weighted Moving Average to covariance, similar to volatility (typically $\lambda = 0.94$).

$$\sigma_{ij,t+1} = (1-\lambda)R_{i,t}R_{j,t} + \lambda\sigma_{ij,t}$$

Flaws: Implies persistence= 1; shocks to covariance persist forever and no mean reversion to a long-run average correlation.

2.4 GARCH covariance

Introduces mean reversion to a long-run unconditional covariance ω_{ij}.

$$\sigma_{ij,t+1} = \omega_{ij} + \alpha R_{i,t}R_{j,t} + \beta\sigma_{ij,t}$$

Constraint: To ensure the matrix is positive semi-definite, we often must assume α and β are the same for all pairs, which is too restrictive (different assets have different correlation dynamics).

2.5 Dynamic conditional correlation (DCC)

The industry standard for modeling large covariance matrices. Covariance matrix is $\Sigma_{t+1} = D_{t+1}\gamma_{t+1}D_{t+1}$, $D_{t+1} = \text{diag}(\sigma_{1,t+1}, \ldots, \sigma_{n,t+1})$, γ_{t+1} is correlation matrix with entries $\rho_{ij,t+1}$.

After estimating conditional volatilities for each asset i, $\sigma_{i,t}$, (*e.g.* using GARCH or EWMA), define $z_{i,t} := \frac{R_{i,t}}{\sigma_{i,t}}$.

Then these standardized residuals $z_{i,t}$ have variance ≈ 1 and conditional correlation is $E_t[z_{i,t+1}z_{j,t+1}] = \rho_{ij,t+1}$ so correlation dynamics modelled using cross product terms.

Can also instead apply exponentially smoothing on z's giving

$$\rho_{ij,t+1} = \frac{q_{ij,t+1}}{\sqrt{q_{ii,t+1}q_{jj,t+1}}}, \quad q_{ij,t+1} = (1-\lambda)z_{i,t}z_{j,t} + \lambda q_{ij,t}$$

and we reconstruct Σ_{t+1} at the end.

2.6 Term structure of VaR example

A simple illustration from lectures is the *VaR term structure*: under GARCH/NGARCH: the 1% VaR as a function of the horizon can look very different depending on the current volatility relative to its long-run level.

- When current volatility is low (below long-run), the VaR term structure can be initially upward sloping because volatility is expected to rise toward its long-run level.
- When current volatility is high (above long-run), the VaR term structure can also be initially upward sloping, but the *shape* differs because volatility is expected to mean-revert downward. Additionally, higher moments (kurtosis) can affect short horizons.

This example shows why long-horizon VaR is not just "today's 1-day VaR times $\sqrt{K}$": the entire *future volatility path* matters, and the long-horizon distribution must be constructed using simulation methods rather than a simple analytical scaling.

3 Long Horizon Risk

Let daily portfolio returns follow $R_t = \sigma_t z_t$, where $z_t \overset{\text{i.i.d.}}{\sim} D_z$ and σ_t is time-varying (*e.g.* GARCH). For a 1-day horizon, risk depends primarily on the current conditional variance σ^2_{t+1}. For a K-day horizon, risk depends on the entire future path $\{\sigma_{t+1}, \sigma_{t+2}, \dots, \sigma_{t+K}\}$.

Monte Carlo and FHS are required to capture time-varying volatility dynamics.

3.1 The square root of time rule

If returns are i.i.d. normal with constant variance σ^2, variance over K days scales linearly $\text{Var}(R_{t:t+K}) = K\sigma^2$, $\text{VaR}_K = \sqrt{K}\,\text{VaR}_1$.

This fails under GARCH; volatility is mean-reverting and stochastic, not constant.

Under GARCH(1,1), $\sigma^2_t = \omega + \alpha R^2_{t-1} + \beta\sigma^2_{t-1}$, $\alpha + \beta < 1$. The conditional variance forecast satisfies $E_t[\sigma^2_{t+k}] = \bar{\sigma}^2 + (\alpha+\beta)^{k-1}(\sigma^2_{t+1} - \bar{\sigma}^2)$, where $\bar{\sigma}^2 = \omega/(1-\alpha-\beta)$.

The K-day variance is therefore $\text{Var}_t(R_{t:t+K}) = \sum_{k=1}^{K} E_t[\sigma^2_{t+k}]$, which is *not* equal to $K\sigma^2_{t+1}$.

Thus, if $\sigma^2_{t+1} > \bar{\sigma}^2$, the $\sqrt{K}$ rule *overestimates* long-horizon risk. If $\sigma^2_{t+1} < \bar{\sigma}^2$, the $\sqrt{K}$ rule *underestimates* long-horizon risk.

Long-horizon risk is harder because it requires forecasting volatility dynamics. Risk is path-dependent: $R_{t:t+K} = \sum_{k=1}^{K} \sigma_{t+k} z_{t+k}$.

Thus, risk depends on persistence $(\alpha + \beta)$; volatility shocks today affect risk many days ahead; closed-form VaR expressions generally do not exist.

3.2 Monte-Carlo simulation (MCS)

Monte Carlo simulation is used to compute long-horizon VaR/ES when analytical scaling fails.

1. Estimate GARCH parameters (ω, α, β) and current σ_{t+1}.

2. Draw K i.i.d. shocks $\hat{z}_1, \ldots, \hat{z}_K \sim N(0,1)$.
3. Generate returns recursively (for $k = 1, \ldots K$): $R_{t+k} = \sigma_{t+k}\hat{z}_k$, $\sigma^2_{t+k+1} = \omega + \alpha R^2_{t+k} + \beta\sigma^2_{t+k}$.
4. Aggregate returns $R_{t:t+K} = \sum_{k=1}^{K} R_{t+k}$.
5. Repeat for many paths and compute VaR as the desired percentile.

3.3 Filtered historical simulation (FHS)

FHS combines GARCH volatility dynamics with the empirical distribution of shocks - assumes distributions of shocks is constant, but volatility changes.

First, fit GARCH model to historical returns $R_{t-\tau}$. Get parameters $\hat{\omega}$, $\hat{\alpha}$, $\hat{\beta}$. Then compute standardized residuals, $\hat{z}_{t-\tau} = \frac{R_{t-\tau}}{\hat{\sigma}_{t-\tau}}$.

Then, to forecast K days ahead, draw z^* from $\{\hat{z}_{t-\tau}\}$ (with replacement, *i.e.* bootstrap).

Compute tomorrow's return, $R^*_{t+1} = \hat{\sigma}_{t+1}\hat{z}^*$.

Update volatility for $t+2$, $\hat{\sigma}^{*2}_{t+2} = \hat{\omega} + \hat{\alpha}(R^*_{t+1})^2 + \hat{\beta}\hat{\sigma}^2_{t+1}$.

Draw another z^*, calculate R^*_{t+2} *etc.* up to day K.

Then sum returns for the path, $R^*_{path} = \sum R^*$. Repeat this for many paths, then VaR is the percentile of the simulated R^*_{path} distribution.

FHS captures fat tails (from historical z) and volatility clustering (from GARCH).

4 Fixed Income Risk

Zero-coupon bond (ZCB): pays \$1 at $t+T$. Its time-t price is $P(t,T) = e^{-TR(t,T)}$, where $R(t,T)$ is the continuously compounded yield. Can rearrange for $R(t,T)$ in terms of price.

Term structure / yield curve: the function $T \mapsto R(t,T)$, mapping each maturity T to the discount rate used to value cashflows at $t+T$ (typically only *partially observed* on a discrete grid of maturities).

Coupon bond pricing (given a TS): for cashflows $\{C_k\}_{k=1}^K$ paid at maturities $\{T_k\}$, $B(t) = \sum_{k=1}^{K} C_k e^{-T_k R(t,T_k)}$.

4.1 Factor structure of yields

From $P(t,T) = e^{-TR(t,T)}$, the *primitive risk factor* for a ZCB is the yield $R(t,T)$ (and for a coupon bond, the collection $\{R(t,T_k)\}$ that discounts each cashflow). We build a *low-dimensional factor representation* for the entire curve because yields co-move strongly across maturities.

A general factor representation used in the notes is

$$\widehat{R}(t,T;\beta_t) = \sum_{j=1}^{J} \beta_{t,j}\, f_j(T),$$

where $\beta_t = (\beta_{t,1}, \ldots, \beta_{t,J})^\top$ is a vector of risk factors that changes over time (daily), $f_j(T)$ are "loading functions" (deterministic functions of maturity), and the curve at each date t is fit cross-sectionally using observed yields at a finite set of maturities.

4.2 Polynomial term structure model

A key model is the 3-factor power-polynomial approximation:

$$\widehat{R}(t, T; \beta_t) = \beta_{t,1} + \beta_{t,2}T + \beta_{t,3}T^2.$$

Factor interpretation via loadings $f_j(T)$:

- Level: $f_1(T) = 1$. A change in $\beta_{t,1}$ shifts all maturities by the same amount (parallel shift).
- Slope: $f_2(T) = T$. A change in $\beta_{t,2}$ affects long maturities more than short ones (rotation).
- Curvature: $f_3(T) = T^2$. A change in $\beta_{t,3}$ induces hump/twist type deformations.

Cross-sectional estimation at fixed t: given observed ZCB yields $\{R(t, T_i)\}_{i=1}^N$ on $\{T_1, \ldots T_n\}$, estimate β_t by least squares:

$$\beta_t = \arg\min_{\beta \in \mathbb{R}^3} \sum_{i=1}^{N} \left(R(t, T_i) - (\beta_1 + \beta_2 T_i + \beta_3 T_i^2)\right)^2.$$

This produces a time series $\{\beta_t\}$ when repeated for $t = 1, \ldots, m$.

4.3 Nelson–Siegel (NS) Model

An example of a parametric model with *exponential* loadings, of the form

$$R(t, T) = \beta_{t,1} + \beta_{t,2}\left(\frac{1 - e^{-\lambda T}}{\lambda T}\right) + \beta_{t,3}\left(\frac{1 - e^{-\lambda T}}{\lambda T} - e^{-\lambda T}\right),$$

where $\lambda > 0$ controls the decay of the loadings. Typical interpretation is: $\beta_{t,1}$ is long-run *level*, $\beta_{t,2}$ is *slope* (short-end effect), and $\beta_{t,3}$ is *curvature* (medium-term hump, governed by λ).

4.4 Dynamic two-step approach (DTSA)

We can manage risk with these models in two steps:

1. Cross-sectional step: for each day t, fit the TS model (*e.g.* polynomial) to observed yields on $\{T_1, \ldots T_n\}$ to estimate β_t (risk-factor levels).
2. Time-series step: model or simulate the factor changes $\Delta\beta_t$. A main approach is i.i.d. Historical Simulation on first differences: $\Delta\beta_{t-u} = \beta_{t-u} - \beta_{t-u-1}$, $u = 0, \ldots, m-1$, treating the vector series $\{\Delta\beta_{t-u}\}$ as i.i.d. and creating scenarios, $\beta_{t+1}^{(i)} = \beta_t + \Delta\beta^{(i)}$.Compare among different values of $\Delta\beta$ for each day to determine VaR.

4.5 Pricing under a shocked curve

Given a factor scenario $\beta_{t+1}^{(i)}$, you build a *scenario yield curve*

$$R^{(i)}(t+1,T) = \widehat{R}(t+1,T;\beta_{t+1}^{(i)}),$$

and then *reprice* the bond/portfolio by discounting each cashflow using the shocked yields:

$$P_c^{(i)}(t+1) = \sum_{k=1}^{K} C_k^{(t+1)} \exp\big(-(T_k^{(t+1)})\, R^{(i)}(t+1,T_k^{(t+1)})\big).$$

For 1-day horizon, maturities shrink by 1/252, *i.e.* $T_k^{(t+1)} = T_k - \frac{1}{252}$, and coupons are aligned to the new times. Thus, the 1-day revaluation is $P_c^{(i)}(t+1) =$

$$\sum_{j=1}^{k} c\left(T_j - \frac{1}{252}\right) \exp\left(-\left(T_j - \frac{1}{252}\right) R^{(i)}\left(t+1, T_j - \frac{1}{252}\right)\right).$$

4.6 Model risk vs market risk

The notes stress that for fixed income we often work with *model-implied prices* $P_c(t)$ rather than observed prices $P_o(t)$ to avoid contaminating the risk measure with model misspecification.

$$P_c^{(i)}(t+1) - P_o(t) = \underbrace{\big(P_c^{(i)}(t+1) - P_c(t)\big)}_{\text{market risk from TS moves}} + \underbrace{\big(P_c(t) - P_o(t)\big)}_{\text{model pricing error}}.$$

Since VaR is intended to measure risk from *movements of the term structure*, the clean object is the model-to-model difference $P_c^{(i)}(t+1) - P_c(t)$, not a difference that includes the static pricing error.

5 Fixed Income Hedging

5.1 Duration and convexity for parallel shifts

For a (small) parallel change dR in yields, the bond price change is approximated by,

$$\frac{dB}{B} \approx -D\, dR + \frac{1}{2}C\,(dR)^2$$

, where $D = -\frac{1}{B}\frac{dB}{dR}$, $C = \frac{1}{B}\frac{d^2B}{dR^2}$.

Key special cases are:

- ZCB $Z(t,T) = e^{-TR(t,T)}$: $D_Z = T$ and $C_Z = T^2$.
- Coupon bond viewed as a portfolio of zeros: its duration/convexity are cashflow-PV weighted moments:

$$D = \sum_{i=1}^{n} w_i T_i, \qquad C = \sum_{i=1}^{n} w_i T_i^2,$$

  where $w_i = \frac{c(T_i)Z(0,T_i)}{B(0,T)}$ are PV weights (same weights in duration and convexity derivations).

5.2 Factor duration and factor convexity

When yields are written as $R(t,T;\beta_t) = \sum_{j=1}^{J} \beta_{t,j} f_j(T)$, factor convexity is defined with respect to factor j as: $C_j = \frac{1}{P}\frac{d^2P}{d\beta_j^2}$.

For a ZCB $Z(t,T) = e^{-TR(t,T)}$ and using $dR(T)/d\beta_j \approx f_j(T)$, this gives: $C_{Z,j} = T^2\left(f_j(T)\right)^2$.

Similarly, factor duration measures first-order exposure to each factor (used to hedge small movements in level/slope/curvature), and factor convexity extends hedging to large moves (second-order effects).

5.3 Hedging yield-curve movements

5.3.1 Hedging a parallel shift using one ZCB

This is an example of a duration hedge. Let B_c be the current value of the bond (or bond portfolio) you want to hedge, and use one ZCB $Z(0,T)$ with duration D_z.

Define the hedged portfolio $V = B_c + k\,Z(0,T)$. Using $dB_c \approx -D\,B_c\,dR$ and $dZ \approx -D_z\,Z\,dR$, the first-order neutrality condition $dV \approx 0$ implies $k = -\frac{D\,B_c}{D_z\,Z(0,T)}$.

5.3.2 Hedging larger parallel shifts using two ZCBs

This is a duration + convexity hedge. To immunize against *both* the first- and second-order terms in dR, use two ZCBs with prices P_1, P_2, durations D_1, D_2, and convexities C_1, C_2.

Let k_1, k_2 be positions and $V = B_c + k_1P_1 + k_2P_2$. Expanding to second order:

$$dV = -(D\,B_c + k_1D_1P_1 + k_2D_2P_2)\,dR + \frac{1}{2}(C\,B_c + k_1C_1P_1 + k_2C_2P_2)\,dR^2.$$

Setting the coefficients of dR and dR^2 to zero gives the system

$$k_1D_1P_1 + k_2D_2P_2 = -D\,B_c, \qquad k_1C_1P_1 + k_2C_2P_2 = -C\,B_c,$$

with closed-form solution:

$$k_1 = -\frac{B_c}{P_1}\Big(\frac{D\,C_2 - C\,D_2}{D_1C_2 - C_1D_2}\Big), \quad k_2 = -\frac{B_c}{P_2}\Big(\frac{D\,C_1 - C\,D_1}{D_2C_1 - C_2D_1}\Big).$$

5.3.3 Hedging non-parallel yield-curve movements

We use *factor durations* for this. Instead of assuming a parallel shift, we model the yield curve with a small number of factors:

$$R(t,T,\beta_t) = \sum_{j=1}^{3} \beta_j(t)\,f_j(T) = \beta_1(t)\cdot 1 + \beta_2(t)\cdot T + \beta_3(t)\cdot T^2,$$

where $f_1(T) = 1$ (level loading), $f_2(T) = T$ (slope loading), $f_3(T) = T^2$ (curvature loading).

For asset price P, the *factor duration* w.r.t. β_j is $D_j = -\frac{1}{P}\frac{dP}{d\beta_j}$.

Since $\frac{dR(t,T,\beta_t)}{d\beta_j} = f_j(T)$, we have

$$D_j = -\frac{1}{P}\frac{dP}{dR}\frac{dR}{d\beta_j} = D \cdot f_j(T),$$

where D is the traditional duration (this is the chain rule result).

Because $D = T$ for a ZCB: $D_{\text{level}} = T \cdot 1 = T$, $D_{\text{slope}} = T \cdot T = T^2$, $D_{\text{curv}} = T \cdot T^2 = T^3$.

5.4 Hedging multiple factors

Suppose your target portfolio has price P and you want to neutralize exposure to two factors (*e.g.* level and slope, or level and curvature). Pick two ZCB hedging instruments: a short one with price P_z^S and a long one with price P_z^L. Hold positions k_S, k_L so the new portfolio is $V = P + k_S P_z^S + k_L P_z^L$.

For two factors (call them factor 1 and factor 2), first-order neutrality requires: (just differentiating)

$$k_S D_{z,1}^S P_z^S + k_L D_{z,1}^L P_z^L = -D_1 P, \quad k_S D_{z,2}^S P_z^S + k_L D_{z,2}^L P_z^L = -D_2 P.$$

The closed-form solution provided for the *level and slope* example is:

$$k_S = -\frac{P}{P_z^S}\Big(\frac{D_1 D_{z,2}^L - D_2 D_{z,1}^L}{D_{z,1}^S D_{z,2}^L - D_{z,2}^S D_{z,1}^L}\Big), k_L = -\frac{P}{P_z^L}\Big(\frac{D_1 D_{z,2}^S - D_2 D_{z,1}^S}{D_{z,1}^L D_{z,2}^S - D_{z,2}^L D_{z,1}^S}\Big),$$

and the interpretation is that investing (k_S, k_L) hedges the portfolio against *any* small combination of those two factor moves.

If we define value weights $w_1 = \frac{k_S P_z^S}{P}$, $w_2 = \frac{k_L P_z^L}{P}$, the system can also be written more simply as:

$$w_1 D_{Lev}^S + w_2 D_{Lev}^L = -D_{Lev,PF}$$
$$w_1 D_{Slope}^S + w_2 D_{Slope}^L = -D_{Slope,PF}$$

where in this particular example factor 1 is *level* and factor 2 is *slope*.

5.5 Factor convexity

To handle large moves in *multiple* factors, define *factor convexity* with respect to factor j: $C_j = \frac{1}{P}\frac{d^2P}{d\beta_j^2}$. This provides second-order sensitivity to factor moves.

For a ZCB $Z(t, T_i) = e^{-R(t,T_i)T_i}$, the slides show:

$$C_{z,j} = \frac{1}{Z(t,T_i)}\frac{d^2 Z(t,T_i)}{d\beta_j^2} = T_i^2\, f_j(T_i)^2.$$

Hedging large movements in multiple factors typically requires more instruments because each factor adds first- and second-order conditions.)

5.6 Bond risk vs. stock risk

For stocks, the natural risk factor is the return (and correlations across stocks). Pricing is trivial: $P_t = S_t$.

For bonds, the key risk is the term structure (yield curve). Even if a ZCB yield is not directly observed, we model $R(t, T)$ via a low-dimensional factor structure and manage risk via factor durations/convexities and hedging with ZCBs.

6 Option Pricing Models

Let S_T denote the underlying asset price at maturity T. The European Call Payoff is $C_T = \max(S_T - K, 0)$; the European Put Payoff is $P_T = \max(K - S_T, 0)$.European options can only be exercised at maturity T, and American options can be exercised at any time prior to T.

For European options on a non-dividend-paying stock, we have *put-call parity*: $c_t + Ke^{-rT} = p_t + S_t$. This identity follows from no-arbitrage arguments and must hold for any $t < T$. Consequences of this are:

- Knowing the call price implies the put price and vice versa.
- European calls and puts have symmetric volatility exposure.

6.1 Risk-neutral valuation principle

The fundamental pricing rule for derivatives is:

$$P_t = e^{-rT}\mathbb{E}^Q[g(S_T)]$$

where Q is the *risk-neutral probability measure*, $g(S_T)$ is the payoff function, and r is the expected return of the asset under Q.

Extra: Investors do not price risk directly; all risk premia are embedded in the change of measure from the physical measure P to Q.

6.2 Black–Scholes (BS) model

Assumptions are: continuous trading; no arbitrage; constant volatility σ; constant risk-free rate r; lognormal stock prices.

Dynamics under Q follow, $S_T = S_t e^{x^*}$, $x^* \sim \mathcal{N}\big((r - \frac{1}{2}\sigma^2)T, \sigma^2 T\big)$

BS prices of European call/put are:

$$c_t = S_t N(d_1) - Ke^{-rT}N(d_2)$$
$$p_t = Ke^{-rT}N(-d_2) - S_t N(-d_1)$$

where

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

6.2.1 Integral representation of option prices

For any payoff $g(S_T)$:

$$p_t = e^{-rT}\int_{-\infty}^{\infty} g(S_t e^{x^*}) f(x^*) dx^*$$

where $f(x^*)$ is the Gaussian density of x^*.

Useful Gaussian identities (from slides):

$$\int_a^{\infty} e^{\beta x} f_{\mu,\sigma^2}(x)\, dx = \exp\big(\beta\mu + \tfrac{1}{2}\beta^2\sigma^2\big)\, \Phi\left(\frac{\mu + \beta\sigma^2 - a}{\sigma}\right)$$

$$\mathbb{E}[e^{\beta X}] = \exp\big(\beta\mu + \tfrac{1}{2}\beta^2\sigma^2\big)$$

6.3 Implied volatility

Implied volatility σ_{imp} solves the equation: $\text{BS}(S_t, K, r, T, \sigma_{\text{imp}}) = \text{Market Price}$.

BS price is increasing in σ, so $\frac{\partial c}{\partial \sigma} > 0$.

Due to volatility skew, OTM puts typically have higher implied vol. This violates constant-σ assumption of BS, and reflects crash risk and fat tails.

6.4 Binomial Tree Model

This is a discrete-time approximation of BS. For time step Δt: $u = e^{\sigma\sqrt{\Delta t}}$, $d = e^{-\sigma\sqrt{\Delta t}} = \frac{1}{u}$.

The risk-neutral probability is $p = \frac{R_f - d}{u - d} = \frac{e^{r\Delta t} - d}{u - d}$.

Even though the tree is discrete, the interest rate is continuously compounded in the course, so the gross risk-free return over Δt is $e^{r\Delta t}$.

6.4.1 Binomial pricing algorithm (backward induction)

1. Construct the stock price tree. Start at S_0 and next prices are $S_{1,1} = u \cdot S_0$, $S_{1,0} = d \cdot S_0$, and then $S_{2,2} = u^2 S_0$, $S_{2,1} = udS_0$, $S_{2,0} = d^2 S_0$ *etc.*
2. Compute the terminal option payoffs; at time T, $V_T = \text{payoff}(S_T)$.
3. Step backwards using:
$$V_t = e^{-r\Delta t}\big(pV_u + (1-p)V_d\big)$$

For American options, use $V_t = \max(\text{exercise value}, V_{cont})$.

For example, from the practice final, for the European call at maturity T (previous step), the call payoff is $C_{4,j} = \max(S_{4,j} - K,\, 0)$, $j = 0, \ldots, 4,$. Working backwards from maturity, the option price by risk-neutral valuation at each node is:
$C_{i,j} = e^{-r\Delta t}\,[p\,C_{i+1,j+1} + (1-p)\,C_{i+1,j}]$.

For at-the-money options ($K = S_0$), the model becomes simpler. The tree becomes symmetric; many terminal nodes have zero payoff, and backward induction is simplified considerably.

6.5 American call option in the binomial model

The binomial pricing procedure for an American option is identical to that of a European option, except that at each node the holder has the right to exercise the option early. Therefore, at every node we must compare the value of continuing to hold the option with the value obtained from immediate exercise.

Consider a 6-month call option with maturity $T = 0.5$ years, strike $K = S_0$, written on a non-dividend-paying stock. We use a 4-period binomial tree ($n = 4$), so that $\Delta t = \frac{T}{n} = 0.125$.

The gross risk-free return per step (with continuously compounded rate r) is $R_f = e^{r\Delta t}$.

Using the Cox–Ross–Rubinstein (CRR) parameterization, the up and down factors are

$$u = e^{\sigma\sqrt{\Delta t}}, \quad d = e^{-\sigma\sqrt{\Delta t}} = \frac{1}{u}.$$

The risk-neutral probability of an up move is

$$p = \pi_u^Q = \frac{R_f - d}{u - d} = \frac{e^{r\Delta t} - d}{u - d},$$

with $1 - p$ the probability of a down move.

Let $S_{i,j}$ denote the stock price at time step i after j up moves. Then

$$S_{i,j} = S_0 u^j d^{i-j}, \quad i = 0, 1, 2, 3, 4, \; j = 0, \ldots, i.$$

At maturity ($i = 4$, terminal payoff) , the American call payoff equals the European payoff:

$$C_{4,j} = \max(S_{4,j} - K, \, 0), \qquad j = 0, \ldots, 4.$$

At each node (i, j) for $i = 3, 2, 1, 0$, compute:

- Continuation value: $C_{i,j}^{\text{cont}} = e^{-r\Delta t} \left[p \, C_{i+1,j+1} + (1-p) \, C_{i+1,j}\right]$.
- Immediate exercise value: $C_{i,j}^{\text{ex}} = \max(S_{i,j} - K, \, 0)$.

The American option value at node (i, j) is then $C_{i,j} = \max\left(C_{i,j}^{\text{ex}}, \, C_{i,j}^{\text{cont}}\right)$.

This recursion is applied backward through the tree until reaching the initial node $(0, 0)$, whose value $C_{0,0}$ is the American call price.

For a non-dividend-paying stock, it is never optimal to exercise an American call option early. Exercising early forfeits the remaining time value of the option while providing no benefit from early stock ownership.

Therefore, at every node, $C_{i,j}^{\text{cont}} \geq C_{i,j}^{\text{ex}}$, and the American call price coincides with the European call price: $C_{0,0}^{\text{American}} = C_{0,0}^{\text{European}}$.

7 Options Risk Management

7.1 Risk Factors for Options

The price of an option is a function of three fundamental *risk factors*: the underlying asset price S_t; the volatility σ; the risk-free interest rate r_f.

7.2 Option delta (linear approximation)

The *delta* of an option measures the sensitivity of the option price to changes in the underlying asset price: $\delta = \frac{\partial V}{\partial S}$.

Using a first-order Taylor expansion around S_t:

$$V(S_{t+1}) \approx V(S_t) + \delta(S_{t+1} - S_t).$$

Defining the (log) return on the underlying as $R_{t+1} \approx \frac{S_{t+1} - S_t}{S_t}$, the dollar change in the option value is approximated by:

$$\Delta V_{PF,t+1} \approx \delta S_t R_{t+1}.$$

An option portfolio behaves approximately like a stock portfolio holding δ shares of the underlying.

7.3 Delta-normal VaR

Assuming conditional normality of returns, $R_{t+1} \sim N(0, \sigma_{t+1}^2)$, the variance of the option portfolio change is:

$$\text{Var}(\Delta V_{PF,t+1}) \approx \delta^2 S_t^2 \sigma_{t+1}^2.$$

The dollar VaR at confidence level p is:

$$\text{VaR}_{t+1}^p = -|\delta| S_t \sigma_{t+1} \Phi^{-1}(p).$$

For a portfolio with m options, $\Delta V \approx (m\delta) S_t R_{t+1}$, and similarly for the VaR, the δ factor is multiplied by m in the formula above.

This linear approximation works well only for small price changes and is unreliable for options near-the-money or during periods of high volatility.

7.4 Option gamma (quadratic approximation)

To capture curvature in option prices, we include the second derivative: $\gamma = \frac{\partial^2 V}{\partial S^2}$.

A second-order Taylor expansion gives:

$$V(S_{t+1}) \approx V(S_t) + \delta(S_{t+1} - S_t) + \frac{1}{2}\gamma(S_{t+1} - S_t)^2.$$

In terms of returns:

$$\Delta V_{PF,t+1} \approx \delta S_t R_{t+1} + \frac{1}{2}\gamma S_t^2 R_{t+1}^2.$$

Gamma is largest when the option is at-the-money; delta-only risk models are particularly misleading for ATM options.

7.5 Delta-Gamma VaR

Because ΔV_{PF} is now nonlinear in returns, no closed-form VaR exists. Instead we:

1. Simulate M scenarios of K-day returns $\{\hat{R}_{K,h}\}_{h=1}^M$ using HS or Monte Carlo.
2. Compute hypothetical portfolio changes, $\Delta\hat{V}_{PF,h} = \delta S_t \hat{R}_{K,h} + \frac{1}{2}\gamma S_t^2 \hat{R}_{K,h}^2$.
3. Estimate VaR as the empirical percentile:

$$\text{VaR}_{t+1:t+K}^p = -\text{Percentile}\left(\{\Delta\hat{V}_{PF,h}\}_{h=1}^M, 100p\right).$$

Limitation: Delta–gamma models still rely on local Taylor approximations and ignore higher-order effects.

7.6 Full valuation method

Full valuation avoids Taylor approximations entirely. The algorithm is:

1. Simulate future returns $\{\hat{R}_{K,h}\}$ using HS or Monte Carlo.
2. Construct future prices: $\hat{S}_{K,h} = S_t e^{\hat{R}_{K,h}}$.
3. Adjust time to maturity for calendar time decay: $\tilde{T}_{\text{new}} = \tilde{T} - \tau$, where τ is the risk horizon in calendar days.

4. Reprice the option using the pricing model (*e.g.* BS):

$$\Delta \hat{V}_{PF,h} = m\big[c(\hat{S}_{K,h}, r_f, X, \tilde{T}_{\text{new}}, \sigma) - c_{\text{mkt}}\big].$$

5. Compute VaR as:

$$\text{VaR}^p_{t+1:t+K} = -\text{Percentile}\left(\{\Delta \hat{V}_{PF,h}\}_{h=1}^{M}, 100p\right).$$

Full valuation is preferred because it has no linear or quadratic approximation, captures all Greeks, and is accurate for large moves and long horizons.

7.6.1 Full valuation with multiple risk factors

When multiple risk factors matter (*e.g.* S_t, σ_t, r_t), each must be jointly simulated:

$$\Delta \hat{V}_{PF,h} = c(\hat{S}_h, \hat{r}_{f,h}, X, \tilde{T} - \tau, \hat{\sigma}_h) - c_{\text{mkt}}.$$

The joint dependence (correlation) among risk factors must be preserved, either through Historical Simulation or multivariate Monte Carlo.

Overall, full valuation is computationally intensive but conceptually simple and is the benchmark method for option risk management in the course.

8 Backtesting & Stress Testing

8.1 VaR exceedances and the hit sequence

Suppose we compute a one–day ahead VaR forecast VaR^p_{t+1} at confidence level p. Let $R_{PF,t+1}$ denote the realized portfolio return.

Define the *hit sequence*:

$$I_{t+1} = \begin{cases} 1, & \text{if } R_{PF,t+1} < -\text{VaR}^p_{t+1} \quad \text{(VaR violation)} \\ 0, & \text{otherwise} \end{cases}$$

Interpretation:

- $I_{t+1} = 1$ means the realized loss exceeded the predicted VaR.
- A VaR model at level p promises that violations occur only $p \times 100\%$ of the time.

If the VaR model is correctly specified, violations should occur with probability p, and be independent over time.

Formally, the null hypothesis would be H_0: $I_{t+1} \sim^{\text{i.i.d.}} \text{Bernoulli}(p)$.

If violations cluster or occur too often (or too rarely), the VaR model is misspecified.

8.2 Unconditional coverage test (Kupiec Test)

This test checks whether the *frequency* of violations matches the nominal level p.

Let T_1 = the number of violations; T_0 = = the number of non-violations; $T = T_0 + T_1$; $\hat{\pi} = T_1/T$ = the observed violation frequency.

Likelihood under the null ($\pi = p$) is, $L(p) = (1-p)^{T_0} p^{T_1}$.

Likelihood under the alternative ($\pi = \hat{\pi}$) is, $L(\hat{\pi}) = (1-\hat{\pi})^{T_0} \hat{\pi}^{T_1}$.

Likelihood ratio statistic:

$$LR_{uc} = -2 \ln \left(\frac{L(p)}{L(\hat{\pi})} \right) = -2 \ln \left[\frac{(1-p)^{T_0} p^{T_1}}{(1-\hat{\pi})^{T_0} \hat{\pi}^{T_1}} \right]$$

Under H_0, $LR_{uc} \sim \chi^2(1)$, so the decision rule is: if LR_{uc} exceeds the critical value (*e.g.* 3.84 at 5%), reject the VaR model.

8.3 Importance of stress testing

VaR and ES are *probabilistic* risk measures:

- VaR answers: "What loss will I exceed only p% of the time?"
- ES answers: "What is the expected loss given a VaR breach?"

This is while "standard" stress testing often does not assign a probability to the scenario.

Historical data, however, is typically short, may not contain extreme market crashes, and may underestimate future tail risk.

Stress testing is important because standard risk models are typically estimated using relatively short historical samples, and those samples may not contain the kinds of extreme events that could plausibly occur in the future (*e.g.* crashes).

As a result, a VaR model that looks "fine" in normal times can still be fragile when market conditions change sharply.

Stress testing addresses this by artificially generating extreme scenarios for the main risk factors driving portfolio returns and then evaluating the model's output under those stressed inputs. In other words, we "stress the model" by exposing it to data different from the data used to specify/estimate it.

Therefore, stress testing is mainly a tool for *robustness*: it helps the risk manager understand *how bad things could get* if markets move in extreme ways, even if the model's estimated probability of such moves is small or uncertain.

Key limitations of standard stress testing:

- Stress scenarios usually have *no probability attached*
- Therefore, results cannot be directly compared to VaR or ES
- This makes portfolio rebalancing decisions ambiguous

8.4 Coherent Stress Testing

This is a model-based approach. To make stress testing compatible with VaR/ES, we must assign probabilities. Let $f(\cdot)$ = the distribution implied by the risk model, and $f_{\text{stress}}(\cdot)$ = the distribution representing a stress scenario.

Assign probability α to the stress scenario. Define the combined distribution:

$$f_{\text{comb}}(x) = \begin{cases} f(x), & \text{with probability } 1-\alpha \\ f_{\text{stress}}(x), & \text{with probability } \alpha \end{cases}$$

The simulation procedure is as follows:

1. Draw $U \sim \text{Uniform}(0, 1)$
2. If $U < \alpha$, draw from f_{stress}
3. Otherwise, draw from f

Once simulated, we compute VaR or ES from draws of f_{comb}, and the resulting risk measure incorporates extreme scenarios coherently.

As a result, stress scenarios are now embedded within a probabilistic framework.

Advantages are that VaR/ES increase smoothly as α increases; the model can be backtested using f_{comb}; and portfolio decisions become well-defined.

8.4.1 Choosing Stress Scenarios

Stress scenarios should reflect:

1. Changed probabilities: extreme events more likely than history suggests.
2. New shocks: events that have never occurred but could.
3. Parameter instability: volatility, correlation, or tail behavior changing.
4. Structural breaks: breakdown of the model itself.

Historical crises are often used as inspiration. Crises may have short-lived impacts (market corrections), or long-lasting effects (regime shifts).

8.5 Example stress test for a VaR system

One practical stress test is a risk-factor shock scenario:

1. Choose the key risk factors that drive the portfolio's P&L (*e.g.* equity index return, yield curve level/slope, FX).
2. Design an extreme scenario by imposing unusually large shocks to these factors—either by:
   - creating shocks that are more likely than the historical database suggests ("change probabilities"),
   - creating shocks that have never occurred but could ("new shocks"),
   - allowing parameters/correlations to change ("change in parameters"),
   - or reflecting potential structural breaks ("change in model structure").
3. Revalue the portfolio under the stress scenario (full valuation or approximation) and compute the implied loss.
4. Compare stressed losses to the VaR produced by the normal risk model. If stressed losses are unacceptably large relative to limits/capital, the portfolio can be rebalanced.

9 Full valuation method example

We provide an example of using full valuation to calculate the risk of a portfolio long 10 units of a put option with current price p_t, assuming the underlying return is the only risk factor.

The lecture notes define full valuation as a two-step procedure:

1. Simulate scenarios for future hypothetical underlying asset prices.
2. Use an option pricing model to compute the future hypothetical option price for each scenario.

This approach is "precise" because it does not rely on Taylor (delta/gamma) approximations, but it is computationally heavier.

9.1.1 Generate scenarios for tomorrow's underlying price

Let the risk horizon be 1 day (so $K = 1$ trading day). We need a scenario set $\{S_{t+1}^{(h)}\}_{h=1}^{M}$ (for a M day window).

Because the underlying is the only risk factor, we only simulate (or resample) the 1-day return.

A standard way is our 'favourite' Historical Simulation on the underlying return:

- compute historical returns $R_{t-\ell}$ for $\ell = 1, \ldots, M$
- create hypothetical prices with log-returns $S_{t+1}^{(h)} = S_t e^{r^{(h)}}$ for all historical returns $r^{(h)} = \{R_{t-l}\}_{l=1}^{M}$.

(Equivalently, we could simulate from an assumed model for the physical distribution of returns, but HS is the simplest "full valuation via HS/MC" method in the notes.)

9.1.2 Reprice the put option under each scenario

For each scenario price $S_{t+1}^{(h)}$, compute the put price using the pricing model.

Using Black–Scholes, the put price at time $t+1$ is: $p_{t+1}^{(h)} = p_{BS}\left(S_{t+1}^{(h)},\, K,\, \tilde{T},\, r_f,\, \sigma\right)$, where the remaining maturity must be reduced by the passage of time: $\tilde{T} = T - \tau$, and τ is the risk horizon in calendar time (for 1 trading day, typically $\tau \approx 1/252$ in years).

9.1.3 Build the P&L distribution for the portfolio (10 puts)

The portfolio is long 10 puts, so under scenario h the dollar P&L over the 1-day horizon is: $\Delta V_{t+1}^{(h)} = 10\left(p_{t+1}^{(h)} - p_t\right)$.

This is full valuation: compute hypothetical option values for each hypothetical underlying value, then compute value changes.

Collect all scenario P&Ls: $\{\Delta V_{t+1}^{(h)}\}_{h=1}^{M}$.

9.1.4 Compute VaR^p_{t+1} from the scenario P&L distribution

For confidence level p (*e.g.* $p = 1\%$), the 1-day VaR is the negative p-quantile of the P&L distribution:

$$\text{VaR}^p_{t+1} = -\text{Quantile}_p\left(\{\Delta V^{(h)}_{t+1}\}^M_{h=1}\right).$$

10 Historical simulation risk for portfolio

- Stocks: model risk through stock returns (or log-returns).
- Bonds: model risk through yield curve movements represented by a term structure factor model (*e.g.* polynomial term structure).

Then, for each historical scenario, we:

1. shock the risk factors using their historical changes, and
2. reprice each asset (stocks trivially, bonds via discounting cashflows under the shocked curve), producing a scenario P&L distribution.

This is "traditional HS" because the scenarios are directly constructed from historical observations, rather than from a parametric model.

10.1.1 Identify the risk factors

Assume the portfolio value at time t is $V_t = \sum_{i=1}^{n_s} w_i S_{i,t} \; + \; \sum_{j=1}^{n_b} q_j B_{j,t}$, where $S_{i,t}$ are stock prices and $B_{j,t}$ are bond prices.

For each stock i: risk factor: the return $R_{i,t+1}$ (or log-return $\Delta \ln S_{i,t+1}$).

Bond prices depend on the term structure $R(t,T)$. We represent the curve by a polynomial factor model: $R(t,T) = \beta_{1,t} + \beta_{2,t}T + \beta_{3,t}T^2 + \beta_{4,t}T^3$ (or sometimes a 3-factor version).

So the bond risk factors are the curve movements (the time series of β's), *i.e.*: $\Delta\beta_{k,t} = \beta_{k,t} - \beta_{k,t-1}, \quad k = 1,\ldots,4$.

10.1.2 Database required

To implement HS you need a historical time series (daily) of:

1. Stock prices (or returns) for each stock in the portfolio.
2. Yield curve data: observed yields across maturities each day (*e.g.* 1M, 3M, 6M, 1Y, 2Y,...), so that you can fit the polynomial curve and estimate β_t each day via cross-sectional OLS.
3. Bond contract details: coupon rate, payment frequency, maturity, face value.

10.1.3 Historical simulation procedure for 1-day VaR

Let the HS window length be N days (*e.g.* 250). At time t: we construct historical shocks.

For stocks, compute historical stock shocks (log-return form is common): $\Delta \ln S_{i,h} = \ln(S_{i,h}/S_{i,h-1}), \quad h = t-N+1,\ldots,t$.

For yield curve, estimate β_h each day from yields, then compute factor shocks: $\Delta\beta_{k,h} = \beta_{k,h} - \beta_{k,h-1}$.

This creates a joint scenario vector each day:

$$(\Delta \ln S_{1,h}, \ldots, \Delta \ln S_{n_s,h}, \Delta\beta_{1,h}, \ldots, \Delta\beta_{4,h}) .$$

Traditional HS uses these historical vectors directly as scenarios, preserving the historical dependence structure (including correlations) between stocks and rates.

Next, apply shocks to today's state to get tomorrow scenarios For each scenario h:

Stock scenario prices: $S^{(h)}_{i,t+1} = S_{i,t} \exp(\Delta \ln S_{i,h})$.

Yield curve scenario: $\beta^{(h)}_{k,t+1} = \beta_{k,t} + \Delta\beta_{k,h}$.

So the shocked curve is: $R^{(h)}(t+1, T) = \beta^{(h)}_{1,t+1} + \beta^{(h)}_{2,t+1}T + \beta^{(h)}_{3,t+1}T^2 + \beta^{(h)}_{4,t+1}T^3$.

Then, full revaluation of assets under each scenario:

- Stocks: $P^{(h)}_{stock,i} = S^{(h)}_{i,t+1}$.
- Bonds: price by discounted cash flows using the shocked term structure: $B^{(h)}_{j,t+1} = \sum_u CF_{j,u} \exp\{-T_{j,u}\, R^{(h)}(t+1, T_{j,u})\}$.

Scenario portfolio P&L distribution: compute the portfolio value in each scenario: $V^{(h)}_{t+1} = \sum_i w_i S^{(h)}_{i,t+1} + \sum_j q_j B^{(h)}_{j,t+1}$.

Then scenario P&L: $\Delta V^{(h)} = V^{(h)}_{t+1} - V_t$.

Compute risk measure (VaR / ES): for a p-level VaR (*e.g.* $p = 1\%$):
$\text{VaR}_p = -\text{Quantile}_p \left(\{\Delta V^{(h)}\}_{h=1}^{N}\right)$.

11 Shocks

A shock is a historical one-period change in a risk factor; in Historical Simulation, we apply historical shocks to today's factor values, reprice the portolio under each shocked scenario, and use the resulting P&L distribution to compute VaR/ES.

A shock is a change in a risk factor over one day (or whatever horizon you're simulating).

11.1.1 Stock example

Today's stock price is S_t. Pick a historical day h and use its log-return shock $\Delta \ln S_h = \ln(S_h/S_{h-1})$. Then, the scenario price is: $S^{(h)}_{t+1} = S_t e^{\Delta \ln S_h}$.

The interpretation of this is - "what if tomorrow the stock moves like it did on day h in the past?".

11.1.2 Bond (yield curve) example

Suppose your curve is described by factors β_t. A historical factor shock is $\Delta\beta_h = \beta_h - \beta_{h-1}$, and the scenario factor is thus $\beta^{(h)}_{t+1} = \beta_t + \Delta\beta_h$.

For the stock case, repricing the stock is trivial as it is just $P^{(h)}_{\text{stock}} = S^{(h)}_{t+1}$, but for bonds, if the curve changes, the discount rates change, so the bond price changes.

If a bond pays cashflows CF_k at maturities T_k, then under scenario h,

$$B^{(h)}_{t+1} = \sum_k CF_k e^{-T_k R^{(h)}(t+1,T_k)}.$$

11.1.3 Producing a scenario P&L distribution

For each scenario h, you now have a scenario portfolio value $V^{(h)}_{t+1}$, then $\Delta V^{(h)} = V^{(h)}_{t+1} - V_t$. Doing this for many historical scenarios gives a whole set of P&Ls: $\{\Delta V^{(h)}\}_{h=1}^{N}$, which is the scenario P&L distribution used to compute VaR/ES.