Financial Engineering Notes

Alexander Rom

University of Southern California

Abstract

These notes cover the main topics of a financial engineering course, from interest rates, bonds, and futures through discrete and continuous-time valuation models, implied volatility, and interest rate derivatives to credit risk and applications such as value-at-risk and economic capital, asset liability management, commodity derivatives, and real options. Each topic is set out with its key definitions, formulas, and figures, together with short explanations of the ideas behind them.

Contents

Chapter 1

Interest Rates, Bonds, and Futures

Financial Derivative: Contracts value derived from the performance of another asset oreconomic variable. Derivative's value arises from **uncertainty**.

⋆ **General Derivative Valuation.**

$$\boxed{V = \sum_x PV \cdot CF(X) \cdot \Pr(X)}$$

- V: value of the derivative security.
- PV: present value factor (risk-free discounting for derivative valuation).
- X: **random outcome** of the underlying variable the payoff is written on.
- $CF(X)$: **cash flow** of the derivative given realization $X = x$.
- $\Pr(X)$: **pricing probability** of outcome X.

Derivative prices are computed as discounted expected cash flows, but the expectation is taken under a pricing distribution rather than the real-world one. These probabilities are not directly observable and are instead inferred from market prices to rule out arbitrage.

⋆ **Real World. Subjective distribution** of future economic outcomes.

- Stylized view of the future; no uniquely correct probability model.
- **Requires** specifying **asset growth rates** and **discount rates**, including **risk premia**.
- Used for: risk management, Value-at-Risk and economic capital, scenario analysis and stress testing.

The real-world measure reflects beliefs about how the economy might evolve and is shaped by judgment and assumptions. It is essential for understanding risk and potential losses, but it is not directly suitable for pricing traded derivatives.

⋆ **Pricing World (Risk-Neutral / Market-Consistent). Implied distribution** of future economic outcomes inferred from market prices.

- Constructed to match today's traded option prices.
- Built under a **no-arbitrage** framework linking spot and derivative markets.
- Asset growth and discounting occur at the **risk-free rate**.

The pricing world is engineered to eliminate arbitrage rather than describe reality. Its probabilities are tools for valuation, ensuring consistency with observed market prices rather than forecasting actual outcomes.

We use the pricing world because derivatives must be priced to prevent arbitrage, not to reflect someone's forecast of the future. Risk-neutral probabilities let us value payoffs using risk-free discounting while still *matching observed market prices, so the result is tradable and internally consistent across securities.*

1.1 Interest Rate Fundamentals

⋆ **SOFR (Secured Overnight Financing Rate).** SOFR is the modern **risk-free reference rate** used for derivative referencing and discounting, built from actual **secured overnight repo transactions** rather than survey-based quotes; it is published by the **Federal Reserve Bank of New York**.
It is transaction-based and secured, so it better reflects real funding conditions and is harder to manipulate, which makes it suitable as a pricing anchor.

LIBOR (London Interbank Offered Rate). LIBOR was the legacy benchmark intended to represent unsecured interbank borrowing costs, based largely on **bank submissions** (panel quotes) and administered by ICE Benchmark Administration, and it was phased out for many uses.
Because it was not reliably observable from liquid trades and became vulnerable to fixing issues, markets transitioned away from it toward transaction-based benchmarks like SOFR.

⋆ **OIS (Overnight Index Swap Rate).** An OIS rate is the fixed rate in an overnight indexed swap, **market-implied** from traded OIS contracts and designed to reflect compounded overnight funding linked to policy rates (e.g., Fed Funds); it is commonly used as the **risk-free discount rate** in collateralized derivative pricing.
OIS is used because it aligns with no-arbitrage funding logic: discounting should match the overnight collateral/funding rate implied by the trading and margining mechanics.

⋆ **U.S. Treasury Rates.** U.S. Treasury yields are default-free government bond rates issued by the U.S. Department of the Treasury and determined in auction markets; they are often used as broad benchmarks, but they are **not** the standard discounting curve **for derivatives**.
Treasuries reflect government borrowing and bond-market supply/demand, while derivative discounting is tied to funding and collateral conditions, so they can differ in practice.

U.S. Treasury (UST) Term Structure. The UST term structure is the schedule of **Treasury yields** across different maturities (e.g., 1M, 3M, 2Y, 10Y, 30Y), describing how the **risk-free interest rate** varies with time to maturity.
It summarizes how markets price time and interest rate risk for government borrowing, forming the baseline yield curve used for macro analysis, bond valuation, and as a reference for constructing other curves.

While the UST term structure is often treated as "risk-free," it reflects Treasury market supply and demand rather than funding or collateral conditions, which is why derivative pricing typically relies on OIS curves instead.

Repurchase Agreement (Repo) Market. The repo market is the market for short-term **collateralized borrowing**, where one party sells securities (typically U.S. Treasuries) for cash and simultaneously agrees to repurchase them at a later date at a higher price.
Economically, a repo functions as a secured loan: the price difference reflects the interest paid on cash, and the use of high-quality collateral makes the market central to funding, liquidity, and the construction of risk-free reference rates such as SOFR.

Treasury Repo Market Segments (3 main types) **Treasury repo transactions** are any repos collateralized by U.S. Treasuries (the overall secured-funding universe). **FICC GCF Repo trades** are standardized, anonymous dealer-to-dealer repos executed on the GCF platform and centrally cleared by FICC (clean benchmark dealer funding). **Cleared bilateral Treasury repo activity** are directly negotiated bilateral repos that are then cleared at FICC (broader, relationship-driven funding with reduced counterparty risk).

Using both GCF and cleared bilateral data gives a rate that is both standardized (benchmark quality) and representative (wide coverage), making SOFR robust and hard to game.

Daily SOFR Rate r_b. The daily SOFR rate r_b is the **volume-weighted median** interest rate of all eligible overnight U.S. Treasury–collateralized repo transactions executed on business day b, aggregated across tri-party repo, FICC GCF repo, and cleared bilateral Treasury repo markets and published by the Federal Reserve Bank of New York.
It is constructed from actual trades, filtered to exclude specials and outliers, and summarized using a median so that the rate reflects the central secured funding cost of the market rather than being distorted by extreme or low-volume transactions.

⋆ **SOFR Calculation.** The annualized SOFR over a period is computed by compounding daily overnight rates and scaling to a calendar-year basis:

$$\boxed{\text{Annualized SOFR } = \left[\prod_{b=1}^{T}\left(1+\frac{r_b\, n_b}{360}\right)-1\right]\frac{360}{d_c}.}$$

T is the number of **business days** in the accrual period,
d_c is the number of **calendar days**,
r_b is the SOFR rate observed on business day b,
n_b is the number of calendar days for which r_b applies (typically $n_b = 1$).
SOFR is built by compounding realized overnight repo funding rates, so it reflects the actual secured cost of borrowing over the period rather than a single quoted yield. The product compounds day-by-day accrual, the -1 turns the growth factor into a realized return, and $360/d_c$ annualizes it under money-market convention, making the measure transaction-faithful and robust for discounting.

SOFR is fully **transaction-based** and incorporates data from **tri-party** **Treasury repo transactions**, **FICC GCF Repo trades**, and **cleared bilateral Treasury repo activity**. *This broad coverage ensures the rate is grounded in deep, observable markets, which is critical for robustness and arbitrage-free pricing.*

Major Risk-Free Reference Rates by Currency. For major currencies, derivative markets use the following overnight reference rates:

USD:	SOFR (overnight **secured** repo rate)
GBP:	SONIA (overnight **unsecured** rate)
JPY:	TONA (overnight **unsecured** rate)
EUR:	€STR / ESTER (overnight **unsecured** rate)
CHF:	SARON (overnight **secured** repo rate)

⋆ **Interest Rate Term Structure.** The **term structure of interest rates** describes interest rates across different maturities observed on a given date.
It summarizes how the price of time varies across horizons and is the foundation for valuing bonds, swaps, and interest-rate derivatives.

Types of Interest Rates. There are three closely related ways to represent the term structure:
1. **Par rates** are the **coupon rates** that set the price of a coupon-paying instrument equal to par (current bond or swap quoted yields).
2. **Zero (spot) rates** are **discount rates** for a single cash flow at a given maturity (zero-coupon yields).
3. **Forward rates** are **implied future spot rates** extracted from today's term structure.

Par rates are what markets quote, while zero and forward rates are analytical objects constructed to value and decompose cash flows consistently.

Market Quoting Convention. Market-traded instruments such as Treasuries and interest rate swaps are typically quoted in terms of **par rates**. Spot (zero) and forward rates are constructed from observed par rates and must be mutually consistent to prevent arbitrage.
Different coupon frequencies (e.g., semi-annual for U.S. Treasuries, annual for SOFR swaps) require careful curve construction so that all representations of the term structure align under no-arbitrage pricing.

⋆ **Risk-Free Discount Curve.** The **discount factor** $P(t,T)$ is the present value at time t of \$1 paid at time T; it is an absolute object and does **not** depend on coupon or compounding conventions.
The discount curve is the most fundamental representation of interest rates, because all cash-flow valuation reduces to multiplying future payoffs by discount factors.

Formally,

$$\boxed{P(t,T) = e^{-R_c(T-t)} = \left(1 + \frac{r_m}{m}\right)^{-m(T-t)},}$$

where R_c is the **continuously compounded discount rate** and r_m is the **m-period compounded rate**. The two are linked by

$$R_c = m \ln\left(1 + \frac{r_m}{m}\right).$$

Different compounding conventions describe the same economic price of time; discount factors provide a convention-free representation that ensures consistency across instruments.

Key properties:

$$\boxed{P(t,t) = 1, \qquad 0 < P(t,T) \le 1 \text{ for } R_c \ge 0.}$$

If interest rates are negative, discount factors can exceed one, reflecting that future cash is worth more than present cash under negative yields.

⋆ **Key Relationships.** The **discount factor** $P(0,T)$ gives the present value of \$1 paid at maturity T. The **spot rate** $R_c(T)$ is the continuously compounded rate that prices a single cash flow at T via

$$P(0,T) = e^{-R_c(T)T}.$$

The **forward rate** $f(0,T)$ is the instantaneous rate for borrowing over an infinitesimal future interval starting at T, extracted from the slope of the discount curve:

$$f(0,T) = -\frac{\partial}{\partial T} \ln P(0,T).$$

Discount factors are the primitive objects; spot rates summarize their level at each maturity, and forward rates summarize their local slope. Knowing $P(0,T)$ uniquely determines both.

Discount factors $P(t, T)$ translate future cash into **present value**: for a fixed maturity, higher **interest rates** imply a lower $P(t, T)$ (faster decay with T). If rates are **negative**, discount factors can exceed one, $P(t, T) > 1$, meaning future cash is worth more than cash today. This curve relationship is the **time value of money**: longer maturities and higher rates make future payments less valuable in present terms.
The monotonic decline of discount factors encodes the time value of money: the steeper the curve, the more expensive it is to receive cash further in the future.

1.2 Forward Rate

⋆ **Forward Rate.** A **forward rate** $F(t, T_1, T_2)$ is the interest rate, agreed upon at time t, that applies to borrowing or lending over the future interval $[T_1, T_2]$.
Forward rates are not forecasts; they are rates implied by today's term structure that ensure no-arbitrage consistency across maturities.

Timing convention.
- t: **valuation time** (usually today; we set $t = 0$).
- T_1: start of the future period.
- T_2: end of the future period, with $T_2 > T_1 > t$.
In practice, we drop t and write forward rates as $F(T_1, T_2)$ because all quantities are observed today.

⋆ **Forward Rate (Continuous Compounding): from Discount Factors and Spot Rates.** Because **discount factors** are prices, no-arbitrage requires consistency across maturities: receiving \$1 at T_2 directly must cost the same as receiving cash at T_1 and rolling it to T_2:

$$P(0, T_2) = P(0, T_1)\, P(T_1, T_2).$$

Under **continuous compounding**, the intertemporal discount factor over $[T_1, T_2]$ is parameterized by the **forward rate**:

$$P(T_1, T_2) = e^{-F(T_1,T_2)(T_2-T_1)}.$$

Substituting and solving yields the forward rate in terms of discount factors:

$$\boxed{F(T_1, T_2) = \frac{\ln P(0, T_1) - \ln P(0, T_2)}{T_2 - T_1}}$$

Spot (zero) rates are a re-expression of discount factors, $P(0, T) = e^{-R(T)T}$, so the same forward rate can be written using spot rates $R_1 = R(T_1)$ and $R_2 = R(T_2)$:

$$\boxed{F(T_1, T_2) = \frac{R_2 T_2 - R_1 T_1}{T_2 - T_1}}$$

Both formulas describe the same object: the annually compounded (continuous) rate implied by today's curve for the future interval $[T_1, T_2]$. Discount factors are the primitive prices; spot rates summarize discounting to each maturity; the forward rate isolates the incremental discounting between maturities and is the rate that equalizes investing to T_2 directly versus rolling from T_1 to T_2.

⋆ **Forward Rate (Discrete / Annual Compounding).** If rates are quoted with **annual compounding**, the same no-arbitrage logic applies, but discounting is expressed differently:

$$\frac{P(0,T_2)}{P(0,T_1)} = (1 + F(T_1,T_2))^{-(T_2-T_1)}.$$

Solving for the forward rate gives

$$\boxed{F(T_1,T_2) = \left(\frac{P(0,T_1)}{P(0,T_2)}\right)^{\frac{1}{T_2-T_1}} - 1}$$

Interpretation: this is the annually compounded rate that, when applied over $T_2 - T_1$ years, reproduces the same intertemporal discounting implied by the discount curve.

⋆ **Forward Rates over Multiple Periods (N-periods)** For a sequence of maturities $0 < T_1 < \cdots < T_N$, discount factors decompose multiplicatively:

$$P(0,T_N) = \prod_{n=1}^{N} P(T_{n-1},T_n).$$

Under continuous compounding,

$$e^{-R_N T_N} = \prod_{n=1}^{N} e^{-F_{n-1,n}(T_n - T_{n-1})}.$$

Taking logs yields

$$\boxed{R_N = \sum_{n=1}^{N} w_n F_{n-1,n}, \qquad w_n = \frac{T_n - T_{n-1}}{T_N}}$$

states that the **spot rate** to maturity T_N is a **time-weighted average** of the forward rates over all sub-intervals leading up to T_N. Here, R_N represents the single rate that discounts a cash flow from T_N back to today, while each $F_{n-1,n}$ is the forward rate applying only to the period $[T_{n-1}, T_n]$. The weights w_n measure the fraction of total time spent in each interval, so longer intervals contribute more to the overall spot rate.

Economically, this relationship reflects how discounting accumulates over time: discount factors multiply across periods, and taking logarithms turns this into an additive decomposition. As a result, the total discounting to maturity can be viewed as an average of the incremental discounting implied by forward rates. This is a no-arbitrage identity, not a statement about expectations or forecasts.

⋆ **Instantaneous Forward Rate.** A standard forward rate $F(T_1,T_2)$ applies over a **finite future interval** $[T_1,T_2]$. Letting this interval shrink to zero defines the **instantaneous forward rate** as a limit:

$$F_{T,T} := \lim_{T_2 \to T_1 \to T} F(T_1,T_2).$$

Using discount factors, this limit yields

$$\boxed{F_{T,T} = f(0,T) = -\frac{\partial}{\partial T} \ln P(0,T) = R(T) + T\,\frac{\partial R(T)}{\partial T}}$$

Interpretation: the instantaneous forward rate is not an average rate over time, but a pointwise rate.

It measures how fast the discount factor is decaying at exactly maturity T, i.e., the marginal time value of money at that point on the curve.

Economically, this rate explains how the price of a zero-coupon bond changes for an infinitesimal increase in maturity. It captures both the level of the spot rate $R(T)$ and the slope of the yield curve at T, making it the local driver of discounting rather than a traded borrowing rate.

Big picture.
Discount factors are prices. Spot rates summarize total discounting from today to a given maturity. Forward rates isolate the incremental discounting between maturities. All three representations describe the same term structure and differ only in how the time value of money is decomposed across time.

⋆ **Discount Function from Forward Rates.** The **discount factor** to maturity T can be written directly in terms of the **instantaneous forward rate** curve:

$$\boxed{P(0,T) = \exp\left(-\int_0^T f(0,t)\,dt\right).}$$

This expression says that discounting is governed by the cumulative effect of instantaneous forward rates over time. Instead of averaging rates over $[0,T]$, we integrate the local marginal cost of time at each maturity.

⋆ **Constant Forward Rate.** If the forward rate is constant across maturities,

$$f(0,t) = r,$$

then

$$\boxed{P(0,T) = \exp\left(-\int_0^T r\,dt\right) = e^{-rT}.}$$

A flat forward curve implies exponential discounting with a constant rate. In this case, spot rates, forward rates, and instantaneous forward rates all coincide.

⋆ **Linear Forward Rate Curve.** If forward rates vary linearly with maturity,

$$f(0,t) = a + bt,$$

then

$$\boxed{P(0,T) = \exp\left(-\int_0^T (a+bt)\,dt\right) = \exp\Big(-(aT + \tfrac{1}{2}bT^2)\Big).}$$

Here, discounting accelerates or decelerates depending on the slope b. A positive slope implies increasing marginal discounting at longer maturities, reflecting an upward-sloping yield curve.

Key interpretation.
Forward rates fully determine discount factors. The discount factor is the exponential of the negative area under the forward-rate curve, so changes in the level or slope of forward rates translate directly into changes in present values. This is why modern term-structure models are built around forward rates rather than spot rates.

$\star$ **Example: Spot $\to$ Discount Factors $\to$ Forward $\to$ Par.** Given a **spot (zero) curve** $\{R_n\}$ (annual compounding, maturity $t_n = n$ years), we can derive the entire term structure objects.

1) Spot rate $\to$ discount factor.

$$\boxed{P(0,t_n) = (1+R_n)^{-t_n}} \qquad (\text{for } t_n = n,\ P(0,n) = (1+R_n)^{-n}).$$

A spot rate is the single annual rate that prices a zero-coupon cash flow at that maturity; the discount factor is the actual price of \$1 at that maturity.

2) Discount factors $\to$ forward rates (between adjacent years). Define the one-period forward rate for $[t_{n-1}, t_n]$ by the ratio of discount factors:

$$\boxed{\mathrm{FWD}(t_{n-1}, t_n) = \frac{P(0,t_{n-1})}{P(0,t_n)} - 1.}$$

This is the implied one-year rate that makes rolling from t_{n-1} to t_n equivalent to investing directly, ensuring no-arbitrage consistency.

3) Discount factors $\to$ par coupon yield. The **par coupon rate** for maturity t_n is the coupon rate that makes a bond price equal face value:

$$\boxed{C_{\mathrm{par}}(t_n) = \frac{1 - P(0,t_n)}{\sum_{i=1}^{n} P(0,t_i)}.}$$

Par rates are how many traded instruments are quoted: they are not primitive, but are implied by discount factors and change whenever the curve shifts.

Key takeaway.

Any one of the four curves (discount factors, spot rates, forward rates, par rates) pins down the others. Discount factors are the pricing primitives; spot/forward/par rates are different ways of quoting the same underlying term structure.

1.3 Bonds

⋆ **Risk-Free Bond Characteristics.** A **risk-free bond** is a fixed-income security with no credit or default risk, so its value depends only on interest rates and time.
Risk-free bonds isolate the time value of money, making them the clean building blocks for discount curves and interest-rate modeling.

The key characteristics are:

- L: **par (face) value**, the amount repaid at maturity.
- C: **coupon payment**, the periodic cash flow paid to the bondholder.
- m: **coupon frequency**, the number of coupon payments per year.
- Fixed / Float: **coupon type**, indicating whether payments are fixed or indexed to a reference rate.
- T: **maturity**, the time in years until principal repayment.
- $B(r,t)$: **market value** of the bond at time t given yield r.

Day-count conventions and accrual bases affect precise cash-flow timing but are often abstracted away when focusing on core valuation logic. Because there is no credit risk, all variation in the bond's price reflects changes in interest rates and remaining time to maturity.

Bond Price	Coupon vs Par Coupon	What we say
$B = L$	$C = C_{\text{par}}$	**At par**
$B > L$	$C > C_{\text{par}}$	**At a premium**
$B < L$	$C < C_{\text{par}}$	**At a discount**

⋆ **Bond Yield (Yield to Maturity).** The **bond yield** y is the single interest rate that equates the present value of the bond's cash flows to its market price, assuming a flat term structure:

$$\boxed{B(y,0) = L\left(\sum_{n=1}^{mT} \frac{C}{m} e^{-yn/m} + e^{-yT}\right),}$$

Yield compresses the entire term structure into one number, making it convenient for quoting but potentially misleading for valuation when rates vary by maturity.

⋆ **Fixed-Rate Bond Pricing.** The price today ($t = 0$) of a **risk-free fixed-rate bond** equals the present value of all coupon payments plus the principal repayment:

$$\boxed{B(\{r_n\},0) = L\left(\sum_{n=1}^{mT} \frac{C}{m} P\left(0, \tfrac{n}{m}\right) + P(0,T)\right) = L\left(\frac{C}{m} A + P(0,T)\right),}$$

where L is face value, C is the annual coupon rate, m is the coupon frequency, and $A = \sum_{n=1}^{mT} P(0, n/m)$ is the **annuity factor**.
This representation makes clear that bond prices are linear combinations of discount factors, which is why discount curves are the fundamental inputs in fixed-income valuation.

⋆ **Par Coupon Yield.** The **par coupon rate** C_{par} is the coupon rate that makes a newly issued bond trade exactly at its **face value** L, meaning the present value of its coupons plus principal equals par:

$$\boxed{L\left(\frac{C_{\text{par}}}{m} \sum_{n=1}^{mT} P\left(0, \tfrac{n}{m}\right) + P(0,T)\right) = L = B,} \qquad C_{\text{par}} = \frac{m\left(1 - P(0,T)\right)}{A}.$$

Conceptually, the par coupon yield answers the question: "What coupon must I set today so that investors are willing to pay exactly par given the current discount curve?" It is not a discount rate itself, but a rate derived from discount factors. Par yields are what we observe in practice when bonds and swaps are quoted and issued, because they provide a clean, standardized way to summarize market conditions at a given maturity. As the discount curve shifts with market rates, the par coupon yield adjusts mechanically to keep the bond priced at par.

⋆ **Spot (Zero) Yield.** The **spot rate** is the yield on a zero-coupon bond with notional L.
Under continuous compounding,

$$L\,P(0,T) = Le^{-y_{\text{spot,cc}}T}, \qquad y_{\text{spot,cc}} = -\frac{\ln P(0,T)}{T}.$$

Under discrete compounding with frequency m,

$$L\,P(0,T) = L\left(1+\frac{y_{\text{spot},m}}{m}\right)^{-mT}, \qquad y_{\text{spot},m} = m\left(P(0,T)^{-1/mT}-1\right).$$

Spot rates directly encode the discount curve and are the cleanest rates for valuation, since each corresponds to exactly one maturity and one cash flow.

⋆ **Floating Rate Bond.** A **floating rate bond** $B^{\text{Flt}}(r,t)$ is a bond whose coupons are tied to the **future realized value** of an underlying interest-rate index (e.g., LIBOR, SOFR). Rather than fixing coupons today, the bond resets its coupon at each payment date using rates observed over the preceding period.
Floating-rate bonds transfer interest-rate risk from the investor to the issuer by continuously updating coupons to prevailing market rates.

Valuation principle (pricing at $t=0$). Under no-arbitrage, valuation assumes that future realized floating coupons equal today's **forward rates**. For a bond with coupon frequency m and maturity T,

$$\boxed{B^{\text{Flt}}(r,0) = L\left(\sum_{n=1}^{mT}\frac{F\left(\frac{t_{n-1}}{m},\frac{t_n}{m}\right)}{m}P\left(0,\tfrac{t_n}{m}\right)+P(0,T)\right)},$$

where $F\left(\frac{t_{n-1}}{m},\frac{t_n}{m}\right)$ is the forward rate (as seen today) applicable to the coupon period $[\frac{t_{n-1}}{m},\frac{t_n}{m}]$.
Each coupon is priced using the forward rate implied by today's curve, reflecting the market-consistent expectation under the pricing measure. **Example: Floating Rate Bond (Step-by-step, Annual Compounding).** Consider a two-period **floating-rate bond** with notional $L=1$. Coupons reset each period to the relevant one-period forward rate. Valued at $t=0$, its cash flows are:

- at $t=1$: coupon $=F_{0,1}$,
- at $t=2$: principal 1 plus coupon $F_{1,2}$, i.e. $1+F_{1,2}$.

Therefore the price is the present value of these cash flows:

$$B_0^{\text{Flt}} = P(0,1)\,F_{0,1}+P(0,2)\,(1+F_{1,2}).$$

Meaning: discount each future floating payment using the discount factors from today.

Rewrite discount factors using forward rates. With annual compounding, no-arbitrage implies discounting over each period is governed by the same forward rates:

$$P(0,1)=\frac{1}{1+F_{0,1}}, \qquad P(0,2)=\frac{1}{(1+F_{0,1})(1+F_{1,2})}.$$

Substitute these into the pricing equation:

$$B_0^{\text{Flt}} = \frac{F_{0,1}}{1+F_{0,1}} + \frac{1+F_{1,2}}{(1+F_{0,1})(1+F_{1,2})}.$$

Meaning: the first term prices the first floating coupon; the second term prices the final principal-plus-coupon using the two-period discount factor.

Put over a common denominator and simplify. Rewrite the first term with the common denominator $(1+F_{0,1})(1+F_{1,2})$:

$$B_0^{\text{Flt}} = \frac{F_{0,1}(1+F_{1,2})}{(1+F_{0,1})(1+F_{1,2})} + \frac{1+F_{1,2}}{(1+F_{0,1})(1+F_{1,2})}.$$

Combine numerators:

$$B_0^{\text{Flt}} = \frac{F_{0,1}(1+F_{1,2}) + (1+F_{1,2})}{(1+F_{0,1})(1+F_{1,2})} = \frac{(1+F_{1,2})(1+F_{0,1})}{(1+F_{0,1})(1+F_{1,2})} = 1.$$

Meaning: because the coupons reset to market rates, the floating coupons and the discounting implied by the same forward rates exactly offset, forcing the bond to trade at par immediately after a reset (ignoring spreads and credit risk).

⋆ **Duration (Modified).** **Duration** measures the **first-order sensitivity** of a bond's price to a small, parallel change in its yield y. Formally,

$$\boxed{D \;=\; -\frac{1}{B}\frac{\partial B}{\partial y}, \qquad \frac{\Delta B}{B} \approx -D\,\Delta y.}$$

Interpretation: duration tells you the approximate percentage price change of a bond for a 1-unit change in yield. A higher duration means the bond is more sensitive to interest rate movements.

Cash-flow interpretation (fixed-rate bond). For a fixed-rate bond with cash flows c_i paid at times t_i, $D \;=\; \sum_{i=1}^{n} t_i \frac{c_i e^{-y t_i}}{B}$.

This shows that duration is a present-value–weighted average time of the bond's cash flows, where later cash flows increase interest-rate sensitivity. The principal repayment is included in the final cash flow.

Effective Duration. When bond prices cannot be differentiated analytically (e.g. embedded options), duration is estimated numerically: $\boxed{D_{\text{eff}} \;=\; -\frac{1}{B}\frac{B_{y+\Delta y} - B_{y-\Delta y}}{2\Delta y}.}$

Effective duration measures price sensitivity using small yield shocks and works even when cash flows depend on interest rates.

⋆ **Convexity.** **Convexity** measures the **second-order sensitivity** of a bond's price to changes in yield:

$$\boxed{C \;=\; \frac{1}{B}\frac{\partial^2 B}{\partial y^2}, \qquad \frac{\partial^2 B}{\partial y^2} = CB.}$$

Convexity captures the curvature of the price–yield relationship and corrects the linear approximation given by duration.

Convexity for a fixed-rate bond. $C \;=\; \frac{1}{B}\sum_{i=1}^{n} c_i t_i^2 e^{-y t_i}$.

Longer-dated cash flows increase convexity more than duration because sensitivity grows with the square of time.

Effective Convexity. Numerical approximation: $\boxed{C_{\text{eff}} \;=\; \frac{1}{B}\frac{B_{y+\Delta y} - 2B + B_{y-\Delta y}}{(\Delta y)^2}.}$

Effective convexity accounts for nonlinearity in bond prices when yields change, especially important for large rate moves.

⋆ **Duration–Convexity Relationship (with interpretation).** Here, B denotes the **bond price** as a function of the yield y, and D is the bond's **modified duration**. The defining relationship

$$\frac{\partial B}{\partial y} = -D\,B$$

states that duration measures the **percentage sensitivity** of the bond price to a small change in yield: a higher D means the price reacts more strongly.

Differentiating both sides with respect to y gives

$$\frac{\partial^2 B}{\partial y^2} = \frac{d(-DB)}{dy}.$$

Applying the product rule, $\frac{d(-DB)}{dy} = (-D)\frac{\partial B}{\partial y} - \frac{dD}{dy}B = (-D)(-DB) - \frac{dD}{dy}B$.
By definition, **convexity** C satisfies $\frac{\partial^2 B}{\partial y^2} = C\,B$,
so dividing by B yields

$$\boxed{C = D^2 - \frac{dD}{dy}.}$$

Interpretation: the term D^2 captures curvature that arises purely from the exponential nature of discounting, while the term $-\frac{dD}{dy}$ accounts for how duration itself changes as yields move. If duration is constant with respect to yield (as for a zero-coupon bond), convexity equals D^2. More generally, convexity reflects both the level of duration and how that sensitivity evolves when interest rates change.

⋆ **Bond Value Sensitivity.** A bond price $B(y,t)$ changes with both the **yield** y and time t. A second-order Taylor expansion gives

$$\Delta B(y,t) \;\approx\; \frac{\partial B}{\partial y}\Delta y + \frac{1}{2}\frac{\partial^2 B}{\partial y^2}(\Delta y)^2 + \frac{\partial B}{\partial t}\Delta t + \cdots$$

For **instantaneous shocks** we ignore all Δt terms, so sensitivity is driven only by Δy.

Using $\frac{\partial B}{\partial y} = -DB$ and $\frac{\partial^2 B}{\partial y^2} = CB$, we obtain the **percentage change** approximation:

$$\boxed{\frac{\Delta B}{B} \;\approx\; -D\,\Delta y + \tfrac{1}{2}C(\Delta y)^2}$$

and the corresponding **dollar change** approximation:

$$\boxed{\Delta B \;\approx\; -BD\,\Delta y + \tfrac{1}{2}BC(\Delta y)^2}$$

DV01 (Dollar Value of 1bp). The first-order dollar sensitivity to a 1 basis point move (1bp = 0.0001) is

$$\text{DV01} = B \cdot D \cdot 0.0001$$

Duration explains the first-order (linear) price move, while convexity explains curvature, meaning price losses from rate increases are smaller than price gains from equal-sized rate decreases.
Fails for Swaps*: Percent-change formulas can fail for contracts with zero value (e.g., swaps) because dividing by B is not meaningful, but dollar sensitivities like DV01 remain well-defined and are the standard risk-management measure.*

Duration and Convexity: Benchmark Bond Cases.

Bond Type	Price P	Duration D	Convexity C
Zero-Coupon Bond	$P = Le^{-rT}$	$D = -\frac{1}{P}\frac{\partial P}{\partial r} = T$	$C = \frac{1}{P}\frac{\partial^2 P}{\partial r^2} = T^2$
Floating-Rate Bond (at reset)	$P = L$	$D = 0$	$C = 0$

Zero-coupon bond *delivers a single cash flow at maturity T, so its entire value is exposed to changes in the discount rate over the full horizon. This makes its duration equal to maturity and its convexity equal to the square of maturity, representing maximal interest-rate sensitivity.*

Floating-rate bond *evaluated immediately before a reset date has future coupons that adjust to current market rates. This adjustment mechanism offsets any change in discount rates, locking the bond's value at par. As a result, the bond exhibits zero duration and zero convexity at the reset point.*

Measure	Value	Interpretation
Duration	Low	Bond price is relatively insensitive to yield changes; interest-rate risk is limited (typical for short-maturity or floating-rate bonds).
Duration	High	Bond price is highly sensitive to yield changes; small rate moves cause large price changes (typical for long-maturity or zero-coupon bonds).
Convexity	Low	Price–yield relationship is close to linear; duration provides a good approximation but offers limited protection for large rate moves.
Convexity	High	Price–yield relationship is strongly curved; bond benefits more when yields fall than it loses when yields rise, improving risk-adjusted performance.
Duration = 2	Example	A 1% parallel increase in yields leads to an approximate 2% decrease in bond price (first-order effect).
Convexity = 2	Example	Indicates moderate curvature; duration slightly underestimates gains when yields fall and overestimates losses when yields rise.

⋆ **Forward Rate Agreement (FRA).** FRA is a contract that locks in a fixed interest rate today for borrowing or lending over a future time interval $[T_1, T_2]$, with settlement based on the realized floating rate over that period. FRA is a **special case of an interest-rate swap**: it represents a single-period fixed-for-floating exchange, whereas a standard swap is a strip of FRAs across many periods.

An FRA isolates interest-rate risk over a single future period by fixing the rate today while settlement occurs in the future.

Economic structure.
An FRA exchanges:

- a predetermined **fixed rate** R_K,
- against the **floating market rate** realized over $[T_1, T_2]$.

Only the **net interest difference** is exchanged; the notional principal is never exchanged.

Conceptually, an FRA is equivalent to a one-period interest-rate swap that settles at the start of the accrual period.

Notation.

- L: notional principal.
- R_K: fixed rate agreed for the period $[T_1, T_2]$.
- $F(T_1, T_2)$: forward rate for $[T_1, T_2]$ implied by today's curve.
- R_2: continuously compounded spot (zero) rate for maturity T_2.

FRA Valuation (Receive Fixed, Pay Float). Under no-arbitrage, valuation assumes the **forward rate is realized**. The value at time $t = 0$ of receiving fixed and paying floating is

$$\boxed{V_{\text{FRA}} = L\Big(R_K - F(T_1, T_2)\Big)(T_2 - T_1)\, e^{-R_2 T_2}.}$$

The FRA's value depends only on the difference between the contracted fixed rate and the market-implied forward rate, discounted to today.

At inception. When the contract is initiated,

$$R_K = F(T_1, T_2) \quad \Rightarrow \quad V_{\text{FRA}} = 0.$$

This mirrors swaps and forwards in general: contracts are struck so they have zero value at inception.

1.4 Boot-Strapping the Yield Curve

⋆ **Bootstrapping the Yield Curve.** **Bootstrapping** is an **iterative no-arbitrage procedure** used to extract **discount factors** $P(0,T)$ from **market-quoted bond and swap prices**.
The key idea is to solve for discount factors sequentially, starting from the shortest maturities where cash flows are simplest, and then using those results to value longer-maturity instruments.

Market Sources for the Curve. Discount factors are extracted from liquid, standardized instruments:

- **SOFR swaps**: annual pay, maturities extending up to 50Y.
- **U.S. Treasury bills and bonds**: semi-annual pay, standard on-the-run maturities (1M, 3M, 6M, 1Y, 2Y, 3Y, 5Y, 7Y, 10Y, 20Y, 30Y).

Quoted market prices and par rates do not directly give discount factors; bootstrapping converts these observable quotes into the pricing primitives $P(0,T)$.

⋆ **General Bootstrapping Framework.**

- Write each instrument's **pricing equation** as a linear combination of discount factors.
- Start with the **shortest maturities** (typically zero-coupon instruments).
- Solve recursively for unknown $P(0,T)$ as longer maturities are added.

At each step, all earlier discount factors are treated as known, ensuring a unique solution when enough instruments are available.

Example: Bootstrapping from Market Prices. Consider the following instruments (principal = 100, annual coupons):

Tenor	Principal	Coupon	Frequency	Price
0.25	100	0	–	98.5
0.5	100	0	–	97.5
1	100	0	–	95.5
2	100	4.30%	Annual	100
3	100	4.10%	Annual	100

Step 1: Zero-Coupon Bonds.
The first three instruments are zero-coupon bonds, which directly identify short-maturity discount factors.
$98.5 = 100\,P(0,0.25) \;\Rightarrow\; P(0,0.25) = 0.985$
$97.5 = 100\,P(0,0.5) \;\Rightarrow\; P(0,0.5) = 0.975$
$95.5 = 100\,P(0,1) \;\Rightarrow\; P(0,1) = 0.955$

Step 2: Coupon Bonds (2-Year).
$4.30\,P(0,1) + 104.30\,P(0,2) = 100 \quad\Rightarrow\quad P(0,2) = 0.919$
Earlier discount factors are already known, so the only unknown is $P(0,2)$.
Step 3: Coupon Bonds (3-Year).
$4.10\,[P(0,1) + P(0,2)] + 104.10\,P(0,3) = 100 \quad\Rightarrow\quad P(0,3) = 0.887$
Bootstrapping continues maturity by maturity until the full curve is recovered.

Tenor	$P(0,T)$	$R(T)$	FWD
0.25	0.985	6.05%	6.05%
0.5	0.975	5.06%	4.08%
1	0.955	4.60%	4.15%
2	0.919	4.20%	3.80%
3	0.887	4.00%	3.61%

⋆ **Bootstrapped Term Structure.** From the discount factors, we compute **spot rates** and **forward rates**:

$$\boxed{R(T) = -\frac{\ln P(0,T)}{T},} \qquad \boxed{F(T_i, T_{i+1}) = \frac{\ln P(0,T_i) - \ln P(0,T_{i+1})}{T_{i+1} - T_i}.}$$

Discount factors determine the entire term structure; spot and forward rates are simply different representations of the same information.

⋆ **Matrix Representation of Bootstrapping.** Let $\mathbf{P}$ denote the vector of **discount factors** at all relevant cash-flow dates and let $\mathbf{B}$ denote the vector of **observed market prices** of traded instruments (bonds, bills, swaps). Each instrument's pricing equation is linear in discount factors, so the system can be written as

$$\boxed{\mathbf{B} = \mathbf{C}\,\mathbf{P}} \qquad \begin{pmatrix} B_1 \\ B_2 \\ \vdots \\ B_M \end{pmatrix} = \begin{pmatrix} c_{1,1} & c_{1,2} & \cdots & c_{1,N} \\ c_{2,1} & c_{2,2} & \cdots & c_{2,N} \\ \vdots & \vdots & \ddots & \vdots \\ c_{M,1} & c_{M,2} & \cdots & c_{M,N} \end{pmatrix} \begin{pmatrix} P(0,T_1) \\ P(0,T_2) \\ \vdots \\ P(0,T_N) \end{pmatrix}$$

- B_i is the **market price** of instrument i,
-$P(0,T_j)$ is the **discount factor** to maturity T_j,
- $c_{i,j}$ is the **cash flow coefficient** of instrument i paid at time T_j (coupon or principal).
Each row represents one traded instrument, and each column represents a discount factor at a specific maturity. The matrix entries encode the timing and size of all cash flows.

The objective of bootstrapping is to **solve for the vector P** of discount factors:
$\boxed{\mathbf{P} = \mathbf{C}^{-1}\mathbf{B} \quad \text{(when the system is exactly determined).}}$

Underdetermined Case and Structural Assumptions. In practice, $M < N$, so there are **more discount factors than traded prices**. The system is underdetermined and admits infinitely many solutions.
To obtain a unique curve, additional structure is imposed, typically by parameterizing the ***forward-rate curve*** *(e.g., piecewise constant, piecewise linear, or log-linear discount factors). The chosen structure selects one economically reasonable solution that exactly fits observed market prices while enforcing smoothness and no-arbitrage.*

⋆ **Piecewise Constant Forward Rates.** Assume forward rates are constant over predefined maturity buckets.

$$f(0,t) = F_k \quad \text{for } t \in [T_k, T_{k+1}].$$

This assumption yields a simple, stable curve and is widely used in practice for initial curve construction.

⋆ **Piecewise Linear Continuous Forward Rates.** Assume forward rates vary linearly between knot points.

$$P(0,T) = \exp\left(-\int_0^T f(0,t)\,dt\right).$$

Discounting becomes the exponential of the area under the forward-rate curve; linear segments correspond to trapezoidal areas.

⋆ **Log-Linear Discount Factor Interpolation.** An alternative is to interpolate $\ln P(0,T)$ linearly across maturities.
Log-linear interpolation preserves positivity of discount factors and produces smooth spot and forward curves, making it popular for production curve systems.

1.5 Swaps

⋆ **Swap** Is a bilateral derivative contract in which two parties agree to exchange sequences of **cash flows** at predetermined future dates, where each leg's cash flows are computed according to specified rules based on underlying **reference variables**.
Economically, a swap allows parties to transform the risk profile of existing or anticipated cash flows without exchanging the underlying assets. The contract itself is structured to have zero value at inception under no-arbitrage pricing.The defining feature is that payoffs are contingent on observable market variables rather than ownership of the underlying assets.

⋆ **Interest Rate Swaps vs Currency Swaps.** An **interest rate swap** exchanges payments based on different interest-rate conventions, typically within the **same currency** (most commonly fixed-for-floating). A **currency swap** exchanges payments denominated in **different currencies** and may involve fixed or floating rates on either leg; in many structures, the economic objective is to transform both **interest-rate exposure** and **FX exposure**.
Interest rate swaps re-allocate term-structure risk inside one currency. Currency swaps re-allocate both term-structure risk and exchange-rate risk across currencies, effectively replacing one funding base with another.
Types of Swaps (by Coupon Structure and Currency).

- **Fixed-for-floating coupon** swaps:
  - **Domestic to domestic:** both legs in the same currency, exchanging a fixed coupon stream for a floating-indexed stream.
  - **Domestic to foreign:** fixed in one currency versus floating in another (a hybrid combining rate and currency transformation).
- **Fixed-for-fixed coupon** swaps:
  - **Domestic to foreign:** fixed payments in one currency swapped for fixed payments in another, primarily a currency transformation tool.
- **Floating-for-floating coupon** swaps:
  - **Domestic to domestic:** floating versus floating within one currency (often used to switch floating indices or reset conventions).
  - **Domestic to foreign:** floating in one currency exchanged for floating in another, combining FX and index exposure changes.

The coupon structure determines which risks are exchanged. Fixed legs concentrate exposure in the discount curve; floating legs concentrate exposure in forward rates and reset mechanics; cross-currency legs add FX risk unless contract design offsets it.

⋆ **Vanilla Swap Properties.** A **vanilla swap** refers to the standard, linear form of a swap contract:

- **No optionality:** payoffs are linear in the underlying reference variable, so the contract behaves like a strip of forward agreements.
- **Zero cost to enter:** swaps are typically struck at a **par rate** so the initial value is 0 (ignoring bid-ask, credit, and funding effects).
- **No initial exchange of principal:** the **notional** is used only to scale coupon payments (principal is not exchanged in a plain-vanilla interest rate swap).
- **Exchange of fixed and floating cash flows:** payments occur at scheduled dates, computed from the fixed rate on one leg and an observed floating index on the other.
- **Underlying is rates and/or currency:** the core vanilla forms are interest-rate swaps and currency swaps, differing mainly in whether cash flows are in one currency or multiple.

Vanilla swaps are the workhorse instruments of fixed income: simple rules, transparent cash flows, and pricing that ties directly to discount factors and forward rates.

⋆ **Bilateral Interest Rate Swap.** A **bilateral interest rate swap** is an interest rate swap entered directly between two counterparties, where one party pays a **fixed rate** and receives a **floating rate** (or vice versa), with cash flows exchanged periodically based on a common **notional principal**.
"Bilateral" means there is no intermediary standing between the parties: each counterparty is directly exposed to the other's credit risk, and the swap cash flows are exchanged according to a single contractual agreement.

⋆ **Swap Cash Flow Illustration (Fixed vs Floating).** The following table shows the realized cash flows for a fixed-for-floating SOFR swap.
Notional: \$100,000; Fixed rate paid: 3%; Floating rate received: 3M **SOFR**;

Accrual End Date	SOFR Rate (%)	Floating Received	Fixed Paid	Net Cash Flow
June 8, 2022	2.20	550	750	-200
Sept 8, 2022	2.60	650	750	-100
Dec 8, 2022	2.80	700	750	-50
Mar 8, 2023	3.10	775	750	$+25$
June 8, 2023	3.30	825	750	$+75$
Sept 8, 2023	3.40	850	750	$+100$
Dec 8, 2023	3.60	900	750	$+150$
Mar 8, 2024	3.80	950	750	$+200$

Positive net cash flow means the floating-rate receipts exceed the fixed payment for that period; negative values indicate the fixed payment dominates.

⋆ **Interest Rate Swap Valuation.** The value of an interest rate swap is the **present value** of its future net cash flows, discounted using the appropriate **risk-free discount curve** (SOFR in collateralized markets).
Because a vanilla swap is constructed to have zero value at inception, valuation later in time measures how changes in the term structure have altered the relative value of fixed versus floating payments.

⋆ **Vanilla Interest Rate Swap Contract Structure** Exchanges fixed interest payments for floating interest payments on a common notional principal.Only net interest payments are exchanged; the notional principal is never exchanged and serves only to scale cash flows.

The standard structure is:

- One leg pays a fixed rate c.
- The other leg pays a floating rate indexed to **SOFR**.
- Payments occur at regular intervals (annual in this setup).
- The swap has **zero cost to enter** at initiation.

The fixed rate c is chosen so that the present value of fixed and floating legs are equal at inception.

⋆ **Valuation Assumptions.** Swap valuation relies on the following pricing assumptions:

- **Forward rates** implied by today's curve are realized on the floating leg.
- Both fixed and floating cash flows are discounted using the **SOFR discount curve**.
- No-arbitrage holds across all maturities.

These assumptions define valuation under the pricing measure and ensure consistency with traded instruments used to build the curve.

Approach 1: Portfolio of Forward Rate Agreements (FRAs). Each payment period of the swap can be interpreted as a single-period **FRA** exchanging a fixed rate for a forward rate.
A swap is economically a strip of FRAs, one for each accrual period, all sharing the same notional.

Approach 2: Difference Between Two Bonds. View the swap as the difference between:

- a **fixed-rate bond** (fixed leg) $= B_{\text{fix}}$,
- a **floating-rate bond** (floating leg).$= B_{\text{flt}}$,

The fixed leg behaves like a coupon bond discounted with the swap curve, while the floating leg behaves like a floating-rate note priced off forward rates and discount factors.

⋆ **Swap Valuation Formula (Annual Payments).** Consider a swap with:

- notional L, (amount on which interest payments are calculated)
- maturity T (years),
- fixed rate c,
- forward rates $F_{(i-1,i)}$,
- discount factors $P_i = P(0, i)$.

The swap value to the fixed-rate receiver is

$$\boxed{V_{\text{swap}} = L\left(\sum_{i=1}^{T} c\,P_i + P_T\right) - L\left(\sum_{i=1}^{T} F_{(i-1,i)}\,P_i + P_T\right) = B_{\text{fix}} - B_{\text{flt}}.}$$

The principal terms cancel, reflecting that no notional is exchanged in a vanilla swap.
Simplifying,

$$\boxed{V_{\text{swap}} = L\sum_{i=1}^{T}\left(c - F_{(i-1,i)}\right)P_i. = \sum_{i=1}^{T} \text{PV of FRA}_i}$$

This expression shows that a swap is a portfolio of FRAs: each term $(c - F_{(i-1,i)})P_i$ is the present value of a single-period fixed-for-floating exchange.

This representation makes clear that the swap's value depends on the difference between the fixed rate and the market-implied forward rates, weighted by discount factors.

Swap Value at Inception. At initiation, the fixed rate c is set equal to the **par swap rate**, satisfying

$$\sum_{i=1}^{T} c\,P_i = \sum_{i=1}^{T} F_{(i-1,i)}\,P_i, \quad \Rightarrow \quad \boxed{V_{\text{swap}}(0) = 0.}$$

This mirrors forwards and futures: swaps are struck so neither party pays upfront, and all value comes from subsequent movements in interest rates.

⋆ **Forward-Starting Swap.** A **forward-starting swap** is an interest rate swap in which the exchange of fixed and floating cash flows begins at a **future start date** $T_s > 0$, rather than immediately at inception. The contractual terms of the swap (fixed rate, notional, maturity, payment frequency) are agreed upon today, but no cash flows occur until T_s.

Economically, a forward-starting swap locks in today the interest rate that will apply to a swap beginning in the future, removing uncertainty about future funding or reinvestment conditions.

⋆ **Variable Notional Swap.** A **variable notional swap** is an interest rate swap in which the **notional principal changes over time**, so cash flows in each period are computed using a period-specific notional L_i rather than a constant amount.

The economic structure of the swap is unchanged; only the scaling of each period's cash flows varies to match an evolving exposure profile.

Common Structures.
- **accreting swap** has a notional that increases over time,
- **amortizing swap** has a notional that decreases over time,
- **roller-coaster swap** allows the notional to move both up and down across periods.

These structures are used to align swap cash flows with assets or liabilities whose principal grows, amortizes, or varies non-monotonically over time.

⋆ **Valuation (Annual Payments).** For a variable notional swap with notional L_{i-1} applying over period $[T_{i-1}, T_i]$,

$$V_{\text{swap}} = B_{\text{Fix}} - B_{\text{Flt}} = \left(c\sum_{i=1}^{T} L_{i-1}P_i + L_{T-1}P_T\right) - \left(\sum_{i=1}^{T} L_{i-1}F_{(i-1,i)}P_i + L_{T-1}P_T\right)$$

$$\boxed{V_{\text{swap}} = \sum_{i=1}^{T} L_{i-1}\big(c - F_{(i-1,i)}\big)P_i.}$$

Compared to a standard swap with constant notional, each period's FRA-style payoff is scaled by a time-varying notional. This allows the swap to hedge assets or liabilities whose principal grows, amortizes, or changes shape over time, while preserving the same no-arbitrage valuation logic.

⋆ **Currency Swap Valuation** A **currency swap** is a contract in which two parties exchange cash flows in **different currencies**, typically involving both periodic interest payments and an exchange of principal amounts, valued under no-arbitrage using interest-rate curves and exchange rates.
Economically, a currency swap transforms cash flows from one currency into another while locking in funding and exchange-rate exposure.

Currency swap valuation relies on expressing all cash flows in a **single currency** and discounting them consistently. Two equivalent representations are used:

$$\text{difference between two bonds} \qquad \text{or} \qquad \text{portfolio of forward contracts.}$$

Both approaches enforce no-arbitrage consistency between interest rates and exchange rates.

Bond Decomposition Approach.
V_D = value of the domestic bond in domestic currency,
V_F = value of the foreign bond in foreign currency,
S_0 = spot exchange rate (domestic per foreign).
The currency swap value is

$$\boxed{V_{\text{swap}} = \begin{cases} V_D - V_F S_0, & \text{receive domestic currency,} \\ V_F S_0 - V_D, & \text{pay domestic currency.} \end{cases}}$$

This representation treats the swap as being long one bond and short another, with foreign cash flows converted into domestic currency at the spot rate.

1.6 Futures and Forwards

Forward Contract Representation. Future foreign-currency cash flows are converted using **forward exchange rates** implied by interest-rate parity:

$$\boxed{F_{0,T} = S_0\, e^{(r-r_f)T},}$$

where r is the domestic risk-free rate and r_f is the foreign risk-free rate.

This ensures that currency conversion and discounting are internally consistent, ruling out arbitrage between money markets and FX markets.

$\star$ **Discounting Principle.** Each leg of the swap is discounted using its **own currency's discount curve**, and foreign-leg values are converted into domestic currency using spot or forward FX rates depending on timing.

Interest-rate risk and exchange-rate risk are separated: discounting reflects time value of money in each currency, while FX conversion reflects relative pricing between currencies.

$\star$ **Cost of Carry** is the **net cost** of holding an asset from today until maturity, equal to the financing cost minus any income earned plus any storage costs. It determines how the spot price grows to the forward price under no-arbitrage. The cost of carry is

$$\boxed{c = r - q + u,} \quad = \quad \text{interest cost } (r) \;-\; \text{income } (q) \;+\; \text{storage cost } (u).$$

Asset Type and the Role of Carry

- **Investment Assets.** Assets held purely for financial return (e.g. stocks, bonds, currencies, investment gold). They provide no benefit from physical possession, so arbitrage fully enforces

$$F_{0,T} = S_0 e^{cT}.$$

- **Consumption Assets.** Assets held for use in production or consumption (e.g. oil, copper, wheat). Physical availability has value, so market participants may prefer holding the asset itself rather than replacing it with a forward. This *convenience yield* lowers the effective carry, implying

$$F_{0,T} \le S_0 e^{cT}.$$

  This inequality reflects the extra value of holding the physical asset today, beyond its financial return.

$\star$ **Valuation of Long and Short Forward Positions**

The forward price is the fair delivery price that makes a new contract have zero value today, while the contract value measures how much an existing contract's fixed price differs from this fair price. Valuing the long and short positions amounts to taking this difference and discounting it to today, since the payoff is realized only at maturity.

Long Forward: (**Buy at T**) Measures how much better today's fair forward price is than your locked-in price, translated back to today.

$$\boxed{f_0^L = (F_t - K)e^{-r(T-t)} = S_t - Ke^{-r(T-t)}}$$

Short Forward: (**Sell at T**) Measures how much better your locked-in price is than today's fair forward price, translated back to today.

$$\boxed{f_0^S = (K - F_t)e^{-r(T-t)} = Ke^{-r(T-t)} - S_t}$$

Initial Fair Pricing At **initiation** ($t = 0$), a forward contract has zero value, so the delivery price equals the fair forward price given by cost of carry.

$$K = F_{0,T} \;\Rightarrow\; \text{Value} = 0 \qquad \Rightarrow \quad \boxed{K = F_0 = S_0 e^{rT}}$$

This ensures neither the buyer nor the seller has an advantage when the contract is created.

⋆ **Term Structure: Contango and Backwardation** On the horizontal axis you have the contract maturity date. On the vertical axis you have the futures price for delivery at that maturity.

Backwardation: A market state in which forward or futures **prices decrease** with maturity and lie below the current spot price, so that

$$\boxed{F_{0,T} < S_0 \quad \text{for longer maturities } T.}$$

Economically, the benefit of holding the asset (convenience yield) exceeds the cost of carry.

Contango: A market state in which forward or futures **prices increase** with maturity and lie above the current spot price, so that

$$\boxed{F_{0,T} > S_0 \quad \text{for longer maturities } T.}$$

Here the forward price mainly reflects cost of carry.

⋆ **Forward Price for Dividend-Paying Assets** The no-income forward price is $F_{0,T} = S_0 e^{rT}$, which assumes the asset only pays off at time T. If the asset pays cash flows before T, this income lowers the cost of holding it and therefore **lowers the fair forward price**, as we will see.

Discrete Income Suppose the asset pays discrete income at known dates before T. Let

$$I = \text{present value at time 0 of all income paid over } [0, T].$$

If you buy the asset today, you pay S_0, but you will receive income with present value I before T. So the net amount you are really carrying forward is $S_0 - I$. The fair forward price becomes:

$$\boxed{F_{0,T} = (S_0 - I)\, e^{rT}}$$

Continuous Income Yield (q): The continuously compounded rate at which an asset pays income, expressed as a proportion of its current price, over the life of the contract. Instead of subtracting a lump sum I, the spot price grows at a net rate $r - q$:

$$\boxed{F_{0,T} = S_0 e^{(r-q)T}}$$

The price grows at the interest rate r because of financing, but is also reduced by the income rate q because you earn money while holding the asset. So,

- **Higher $r \Rightarrow$ higher forward price.**
- **Higher $q \Rightarrow$ lower forward price.**

1.7 Futures Hedging Strategies

⋆ **Basis Risk.** Is the risk that the **spot price** and the **futures price** do not move one-for-one over the hedging horizon. The **basis** at time t for a futures contract maturing at T is

$$\boxed{\text{Basis}_t = S_t - F_{t,T}.}$$

If the hedge is closed before futures maturity, or if the hedged asset differs from the futures underlying, the basis at the close-out date is uncertain. Only when the hedge horizon coincides with futures expiration does the basis converge to zero.

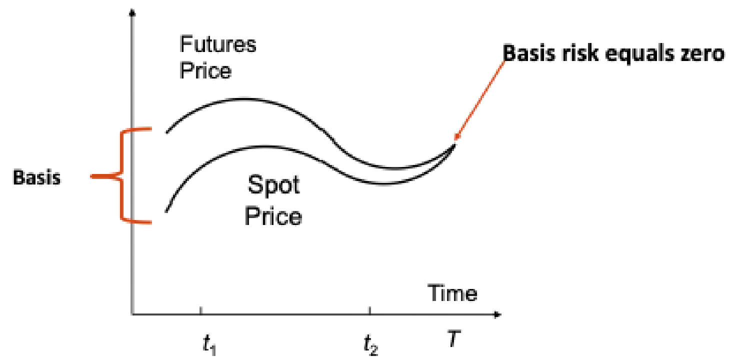

⋆ **Long Hedge: Future Asset Purchase.** A **long hedge** is used when an agent plans to **buy the underlying asset** at a future date and wants to lock in today's price. The hedger goes long a futures contract at time t_1 and purchases the asset at time $t_2 < T$.

- $F_{1,T}$ and $F_{2,T}$ be futures prices at t_1 and t_2,
- S_2 asset Price at t_2 (spot price at purchase)

Then effective cost of the asset is

$$\boxed{\text{Cost} = S_2 - (F_{2,T} - F_{1,T}) = F_{1,T} + (S_2 - F_{2,T}) = F_{1,T} + \text{Basis}_{t_2}.}$$

	Futures contract at t=1	$F_{2,T} - F_{1,T}$
Payoff at $t = 2$:	**Close futures contract**	$F_{2,T} - F_{1,T}$
	Buy asset	$-S_2$

The futures position offsets spot price risk, but the final acquisition cost still depends on the basis at the hedge close-out date. Uncertainty in this basis is the residual risk of the hedge.

⋆ **Short Hedge: Future Asset Sale.** A **short hedge** is used when an agent plans to **sell the underlying asset** at a future date and wants to lock in today's price. The hedger shorts a futures contract at time t_1 and sells the asset at time $t_2 < T$.
Then effective price received is

$$\boxed{\text{Price Received} = S_2 + (F_{1,T} - F_{2,T}) = F_{1,T} + (S_2 - F_{2,T}) = F_{1,T} + \text{Basis}_{t_2}.}$$

	Futures contract at t=1	$F_{1,T} - F_{2,T}$
Payoff at $t = 2$:	**Close futures contract**	$F_{1,T} - F_{2,T}$
	Sell asset	S_2

As with a long hedge, futures lock in the futures price at initiation, but the realized selling price depends on the basis when the hedge is lifted. Basis risk therefore determines the hedge's effectiveness.

⋆ **Cross-Hedging.** Is hedging an exposure when the available futures contract is not written on the exact asset being hedged (**no perfectly matching contract**). Example: needing jet fuel exposure hedged but only crude oil futures are available.

You reduce risk by hedging with a related contract, but because the match is imperfect, some risk remains and the hedge must be scaled to reflect correlation.

⋆ **Optimal Hedge Ratio.** When the hedge asset differs from the exposure, choose the hedge size to minimize the **variance** of the hedged position. The **optimal hedge ratio** is

$$\boxed{h^\star = \rho\,\frac{\sigma_S}{\sigma_F} = \frac{\rho\,\sigma_S\sigma_F}{\sigma_F^2} = \frac{\sigma_{SF}}{\sigma_F^2} = \beta_{S,F}.}$$

$\sigma_S = \mathrm{sd}(\Delta S),\quad \sigma_F = \mathrm{sd}(\Delta F),\quad \rho = \mathrm{Corr}(\Delta S, \Delta F),\quad \sigma_{SF} = \mathrm{Cov}(\Delta S, \Delta F) = \rho\sigma_S\sigma_F$
$\sigma_S^2 = \mathrm{Var}(\Delta S)\quad \sigma_F^2 = \mathrm{Var}(\Delta F)\qquad$ (dollar price changes)

$\beta_{S,F}$ **is the regression slope in** $\Delta S = \alpha + \beta_{S,F}\Delta F + \varepsilon$.

$h^\star$ is the number that scales the futures position so the hedge contract's price moves best offset the exposure's price moves. If correlation is low or futures are much more volatile than the exposure, the optimal hedge is smaller because the hedge would otherwise add noise.

Proof: Define a hedged portfolio Π consisting of a **long position** in one unit of the underlying asset and a **short position** of h units of the futures (or forward) contract. The change in portfolio value over the hedging horizon is

$$\Delta\Pi = \Delta S - h\,\Delta F \qquad \Rightarrow \qquad \mathrm{Var}(\Delta\Pi) = \mathrm{Var}(\Delta S - h\,\Delta F) = \sigma_S^2 - 2h\,\sigma_{SF} + h^2\sigma_F^2,$$

This quadratic form shows how portfolio risk depends on the hedge size: too little hedging leaves exposure risk, while too much hedging amplifies futures volatility.

Variance-minimizing hedge, differentiate with respect to h and set the derivative equal to zero:

$$\frac{\partial}{\partial h}\mathrm{Var}(\Delta\Pi) = -2\sigma_{SF} + 2h\,\sigma_F^2 = h\,\sigma_F^2 - \rho\sigma_S\sigma_F = 0 \qquad \Rightarrow h^\star = \frac{\sigma_{SF}}{\sigma_F^2} = \rho\,\frac{\sigma_S}{\sigma_F}.\square$$

⋆ **Forwards Optimal Hedge** Consider an exposure of Q_A **units of asset** X. The hedge is implemented using forward or futures contracts written on asset Y, where each **future contract represents** Q_F **units of** Y.

Optimal number of hedge contracts is $\boxed{N^\star = \frac{h^\star Q_A}{Q_F}}.$

The hedge ratio $h^\star$ determines how aggressively the exposure should be hedged, while the contract size Q_F scales this into an actual number of contracts traded.

Basis risk considerations. *Basis risk arises if the hedged asset X differs from the futures underlying Y, or if the hedge horizon does not align with contract maturity. When $X = Y$ and maturities align, $h^\star = 1$ and basis risk is eliminated.*

⋆ **Futures Parameter Estimation:** In practice, the optimal hedge ratio is estimated statistically using **historical price data**. Estimation is typically Daily volatility of asset returns (percentage price changes).

$\hat{\sigma}_S = \mathrm{sd}\left(\frac{\Delta S_t}{S_t}\right), \qquad \hat{\sigma}_F = \mathrm{sd}\left(\frac{\Delta F_t}{F_t}\right), \qquad \hat{\rho} = \mathrm{Corr}\left(\frac{\Delta S_t}{S_t}, \frac{\Delta F_t}{F_t}\right)$ (percentage return changes)

$\hat{\sigma}_{\Delta S} = \mathrm{sd}(\Delta S_t), \qquad \hat{\sigma}_{\Delta F} = \mathrm{sd}(\Delta F_t)$ (dollar price changes)

Link (local approximation): $\Delta S_t \approx S_t \left(\frac{\Delta S_t}{S_t}\right), \quad \Delta F_t \approx F_t \left(\frac{\Delta F_t}{F_t}\right) \Rightarrow$

$$\Rightarrow \qquad \hat{\sigma}_{\Delta S} \approx S_t\, \hat{\sigma}_S, \qquad \hat{\sigma}_{\Delta F} \approx F_t\, \hat{\sigma}_F$$

Return volatilities are unit-free and stable for estimation (regressions/hedge ratios). Dollar volatilities scale with the current price level, so even with constant return volatility, dollar risk changes as S_t or F_t move.

Then estimated hedge ratio is $\boxed{\hat{h} = \hat{\rho}\,\frac{\hat{\sigma}_S}{\hat{\sigma}_F} = \hat{\beta}.}$

The optimal hedge ratio coincides with the slope coefficient from a regression of spot returns on futures returns.

$V_A = Q_A P_A =$ value of the exposure being hedged,
$V_F = Q_F P_F =$ value of one futures/forward contract

The corresponding number of hedge contracts is $\boxed{N^* = h^* \frac{V_A}{V_F}}$

Regression-Based Hedge Ratio. The hedge ratio is obtained from the regression

$$\boxed{\frac{\Delta S_t}{S_t} = \alpha + \beta\,\frac{\Delta F_t}{F_t} + \varepsilon_t.}$$

The regression isolates the component of spot price variation explained by futures price movements. The coefficient β directly measures how many futures units are needed to optimally offset spot price risk in a minimum-variance sense.

Why the Optimal Number of Contracts Differs for Forwards and Futures.

- **Hedge ratio** h^* is **dimensionless** and measures the proportional sensitivity of the exposure to the hedge instrument; it does not account for contract size.
- **Forwards hedge quantities.** A forward contract specifies delivery of a fixed number of units, so the hedge is scaled by physical quantities: $N^*_{\text{forwards}} = h^* \frac{Q_A}{Q_F}$.
- **Futures hedge value changes.** Futures are marked to market daily, and hedging occurs through dollar gains and losses, so scaling is done using contract values: $N^*_{\text{futures}} = h^* \frac{V_A}{V_F}$.
- **Economic distinction:** forwards lock in a terminal delivery price for a given quantity, while futures hedge the path of price changes through daily cash flows.

Different contract mechanics imply different scaling rules: quantities for forwards and values for futures, even though the underlying hedge ratio h^* is the same.

⋆ **Rolling the Hedge.** Extending a hedge horizon when no liquid long-dated futures contracts exist by using a sequence of shorter-maturity futures contracts. The hedger closes the expiring contract and enters a new one, repeating the process until the final hedge date.
*Rolling allows long-horizon price locking using short-term contracts, but each rollover introduces **basis risk** and margin cash-flow risk.*

Rolling a 2-Year Hedge with 1-Year Futures. An agent plans to **buy the asset** in 2 years and wants to lock in today's price. Only 1-year futures contracts are available.

- $F_{0,1}$ = 1-year futures price today, $F_{1,2}$ = 1-year futures price in year 1,
- S_0 = spot price today, S_1 = spot price in year 1, S_2 = spot price in year 2,
- $r_{0,1}$, $r_{1,2}$ = 1-year risk-free rates.

Strategy. At $t = 0$, enter a long 1-year futures contract at $F_{0,1}$. At $t = 1$, close the first contract and enter a new long 1-year futures contract at $F_{1,2}$.

Futures payoffs. Year 1 payoff: $S_1 - F_{0,1}$, Year 2 payoff: $S_2 - F_{1,2}$.

Accumulated futures proceeds at year 2.

$$(S_1 - F_{0,1})(1 + r_{1,2}) + (S_2 - F_{1,2}) = S_2 - S_0(1 + r_{0,1})(1 + r_{1,2}).$$

Effective asset purchase price.

$$\boxed{S_0(1 + r_{0,1})(1 + r_{1,2})}$$

With deterministic interest rates, rolling short-dated futures replicates a long-dated futures hedge. The remaining risk comes from interim margin cash flows and basis risk at the roll date, especially at $t = 1$.

⋆ **Tailing the Hedge.** Adjusting the futures hedge to account for interest earned or paid on the futures margin account by hedging the **present value** of the desired exposure. Instead of using $\hat{h}$ futures contracts, adjust as $\boxed{\dfrac{\hat{h}}{(1+r)^T}}$

⋆ **Margin Cash-Flow Risk.** The risk that a futures hedge generates uncertain interim cash inflows or outflows due to daily **mark-to-market** settlement, potentially requiring the hedger to post margin before the underlying exposure is realized.
Even when the hedge is correct in present value, adverse futures price movements can force temporary cash outflows, creating liquidity pressure rather than pricing error.

Daily margin cash flow.

$$\boxed{\text{Margin CF}_t = N\, Q_F\, (F_t - F_{t-1})}$$

Futures gains or losses are realized daily; negative values represent margin calls that must be funded immediately.

Variance of margin cash flows.

$$\boxed{\text{Var(Margin CF)} = (NQ_F)^2\, \text{Var}(\Delta F)}$$

Margin cash-flow risk increases with contract size, number of contracts, and futures price volatility, and is especially relevant for rolling hedges and volatile markets.

Chapter 2

Discrete Valuation Models

2.1 Options Intro

⋆ **Payoff at Maturity.** For maturity T, the option payoffs are given by:

$$\boxed{C(T) = \max\{S(T) - K, 0\} \quad \text{and} \quad P(T) = \max\{K - S(T), 0\}}$$

Equivalently, these can be expressed piecewise as:

$$C(T) = \begin{cases} S(T) - K, & \text{if } S(T) > K, \\ 0, & \text{if } S(T) \leq K, \end{cases} \qquad P(T) = \begin{cases} K - S(T), & \text{if } S(T) < K, \\ 0, & \text{if } S(T) \geq K. \end{cases}$$

At maturity, a call is exercised only if the asset price exceeds the strike price (profitable to buy), and a put is exercised only if the asset price is below the strike price (profitable to sell).

⋆ **Moneyness of an Option.**

- A **CALL** option is **OTM** if $S(t) < K$, **ATM** if $S(t) = K$, and **ITM** if $S(t) > K$.
- A **PUT** option is **OTM** if $S(t) > K$, **ATM** if $S(t) = K$, and **ITM** if $S(t) < K$.

"Moneyness" describes whether exercising the option would currently yield a profit. It determines how valuable or likely the option is to be exercised at maturity.

⋆ **Intrinsic Value.** Value of the option if it could be exercised immediately (time 0):

$$\boxed{C_{\text{int}}(0) = \max\{S_0 - K, 0\}, \qquad P_{\text{int}}(0) = \max\{K - S_0, 0\}}$$

This is the "exercise-now" profit. If exercising right now is not profitable, intrinsic value is 0.

⋆ **Time Value.** Difference between the option value and the intrinsic value:

$$C_0 = C_{\text{int}}(0) + C_{\text{time}}(0), \qquad P_0 = P_{\text{int}}(0) + P_{\text{time}}(0), \Rightarrow$$

$$\Rightarrow \boxed{C_{\text{time}}(0) = C_0 - \max\{S_0 - K, 0\}, \qquad P_{\text{time}}(0) = P_0 - \max\{K - S_0, 0\}}$$

This is what you pay for the possibility that the option becomes more favorable before expiration. It is the extra premium beyond immediate exercise value, driven by uncertainty and time remaining.

Concept	European	American
Call option value	c	C
Put option value	p	P
Exercise (maturity) style	at T	any time $\le T$
Stock price (today)	S_0	
Stock price (maturity)	S_T	
Strike price	K	
Option maturity	T	
Volatility	σ	
Risk-free rate	r	
Dividends (discrete)	D	
Dividend yield (continuous)	q	

Variable	c	p	C	P
S_0	$+$	$-$	$+$	$-$
T	?	?	$+$	$+$
σ	$+$	$+$	$+$	$+$
r	$+$	$-$	$+$	$-$
D	$-$	$+$	$-$	$+$
$C \ge c \ge 0$,		$P \ge p \ge 0$.		
American value $\ge$ European value				

⋆ **Lower Bounds: European Options.** The bounds therefore represent the cheapest price consistent with avoiding arbitrage. If an option were priced below these bounds, an investor could construct a portfolio using the stock, dividends, and a risk-free bond that guarantees a nonnegative payoff at maturity while costing less today, creating a riskless profit.
Dividends lower the stock price by transferring value to shareholders, which reduces call payoffs but increases put payoffs, giving dividends a negative effect on calls and a positive effect on puts.
Discrete Dividends.

$$c \;\ge\; \max\Big(S_0 - D - Ke^{-rT},\, 0\Big), \qquad p \;\ge\; \max\Big(D + Ke^{-rT} - S_0,\, 0\Big).$$

The stock price is reduced by the present value of expected dividends since option holders do not receive dividends. The lower bound reflects the payoff from holding the discounted intrinsic value strategy.

Continuous Dividends.

$$c \;\ge\; \max\Big(S_0e^{-qT} - Ke^{-rT},\, 0\Big), \qquad p \;\ge\; \max\Big(Ke^{-rT} - S_0e^{-qT},\, 0\Big).$$

With a continuous dividend yield, the stock price is effectively discounted at rate q. This captures the continuous erosion of value due to dividend payouts over the option's life.

⋆ **Translation Between Dividend Models.** $\boxed{S_0 - D \;=\; S_0e^{-qT}}$
A discrete dividend model and a continuous dividend yield model are equivalent when the present value of discrete dividends equals the continuously discounted stock price reduction. This allows bounds and pricing formulas to be translated across dividend specifications.

Proof (Call Lower Bound, Discrete Dividends). Compare outcomes of two portfolios.

Portfolio A: one European call c plus cash $D + Ke^{-rT}$.
Portfolio B: one share of stock.

Payoff at maturity T.

$$V_A(T) = \begin{cases} (S_T - K) + De^{rT} + K, & S_T > K, \\ 0 + De^{rT} + K, & S_T \le K, \end{cases} \qquad V_B(T) = S_T + De^{rT}.$$

- If $S_T > K$: $V_A(T) = (S_T - K) + De^{rT} + K = S_T + De^{rT} = V_B(T)$.

- If $S_T \le K$: $V_A(T) = K + De^{rT} \ge S_T + De^{rT} = V_B(T)$.

At T, Portfolio A matches Portfolio B when the call is exercised and dominates it when the call is not exercised. Hence $V_A(T) \ge V_B(T)$ for all S_T.

No-arbitrage implication (today). Since $V_A(T) \ge V_B(T)$ state-by-state, no-arbitrage implies $V_A(0) \ge V_B(0)$:

$$c + \Big(D + Ke^{-rT}\Big) \ \ge\ S_0 \qquad \Rightarrow \qquad c \ \ge\ S_0 - D - Ke^{-rT}.$$

Also $c \ge 0$, so

$$\boxed{c \ \ge\ \max\Big(S_0 - D - Ke^{-rT},\, 0\Big).}$$

If c were below this bound, Portfolio A would cost less than Portfolio B today while never paying less at maturity, creating an arbitrage. $\square$

⋆ **Early Exercise: American Call** (No Dividends) An American call on a non-dividend-paying stock is never exercised early.

$$\boxed{C \ \ge\ c \ \ge\ S_0 - Ke^{-rT} \ \ge\ S_0 - K \ =\ \text{early exercise value}}$$

Early exercise of an American call on a non-dividend-paying stock is never optimal because exercising destroys time value while providing no compensating benefit. By waiting, the holder keeps the option's upside convexity and delays paying the strike K, whose present value is lower than its face value when $r > 0$. Since the stock pays no dividends, there is no cash flow gained by exercising early, so holding the call strictly dominates immediate exercise.

⋆ **Early Exercise: American Put** An American put should be exercised early when the immediate exercise value exceeds the value of waiting.

$$\textbf{Immediate exercise value} = K - S_0.$$

$$\textbf{Continuation value} = P_{\text{continuation}} = \mathbb{E}^{\mathbb{Q}}\Big[e^{-r(T-t)}(K - S_T)^+\Big].$$

For a **deep in-the-money** put, the payoff is almost certain and approximately linear, so the continuation value is well approximated by $P_{\text{continuation}} \ \approx\ Ke^{-r(T-t)} - S_0$.
Thus the early exercise condition becomes

$$\boxed{K - S_0 \ >\ Ke^{-r(T-t)} - S_0 \quad\iff\quad K > Ke^{-r(T-t)}.}$$

Exercising early allows the holder to receive the strike K immediately rather than waiting for its discounted value in the future. When interest rates are positive and the put is sufficiently in the money, waiting only delays cash without meaningful upside, making early exercise optimal.

Put–Call Parity Arbitrage relationship between **European** put and call options with the same strike K and maturity T.

⋆ **Discrete Dividends.** $\boxed{c + Ke^{-rT} \;=\; p + S_0 - D.}$

⋆ **Continuous Dividend Yield.** $\boxed{c + Ke^{-rT} \;=\; p + S_0 e^{-qT}.}$

Conceptually, the put–call parity expresses a fundamental no-arbitrage condition linking call, put, stock, and bond positions. It shows that holding a call option and lending an amount equal to the present value of the strike is equivalent to holding a put and the underlying asset. When dividends are included, their present value (or continuous yield adjustment) must be subtracted since they reduce the asset's effective value to the option holder. This parity ensures that no combination of positions can produce a risk-free profit, preserving market equilibrium.

Proof (Put–Call Parity, Discrete Dividends). Compare two portfolios both long

Portfolio A: European call c plus cash $\big(Ke^{-rT} + D\big)$.
Portfolio B: European put p plus one share of stock S.

Payoff at maturity T. Cash $\big(Ke^{-rT} + D\big)$ grows to $K + De^{rT}$. The stock held from 0 to T delivers value $S_T + De^{rT}$ when accounting for dividend cash flows.

$$V_A(T) = \max(S_T - K, 0) + (K + De^{rT}) = \max(S_T + De^{rT},\, K + De^{rT}),$$
$$V_B(T) = \max(K - S_T, 0) + (S_T + De^{rT}) = \max(S_T + De^{rT},\, K + De^{rT}).$$

Both portfolios produce the same terminal payoff in every state: if $S_T \geq K$, the call is exercised and the put expires; if $S_T < K$, the put is exercised and the call expires. In both cases, the total value equals the larger of $S_T + De^{rT}$ and $K + De^{rT}$.

No-arbitrage implication (today). Since $V_A(T) = V_B(T)$ state-by-state, their time-0 costs must be equal:

$$c + \big(Ke^{-rT} + D\big) = p + S_0 \qquad \Rightarrow \qquad \boxed{c + Ke^{-rT} = p + S_0 - D.} \qquad \square$$

If the prices differed today, one could buy the cheaper portfolio and sell the more expensive one, locking in the same payoff at T for a riskless profit.

Put–Call Parity Extensions. Put–call parity generalizes to other underlyings by adjusting the present value of holding the underlying asset.

⋆ **Stock Index Options.**

$$c + Ke^{-rT} \;=\; p + S_0 e^{-qT}.$$

Stock indices typically pay dividends continuously at yield q. Holding the index is therefore equivalent to holding a non-dividend-paying asset with an effective discounted price S_0e^{-qT}, so parity replaces S_0 with its dividend-adjusted value.

⋆ **Foreign Exchange Options.**

$$c + Ke^{-rT} \;=\; p + X_0 e^{-r_f T}.$$

Holding foreign currency earns the foreign risk-free rate r_f. Thus, the present value of holding the foreign asset is discounted at r_f, while the strike is discounted at the domestic rate r, preserving no-arbitrage parity across currencies.

⋆ **Futures Options.**

$$c + Ke^{-rT} \;=\; p + F_0 e^{-rT}.$$

A futures position requires no initial investment, so its current value is zero and its future price F_0 is already a forward-looking quantity. Both sides of parity are discounted at the same risk-free rate, yielding a simplified expression.

⋆ **Chooser Option.** Option expiring at time T_1 with the right to choose either a European call or a European put, both with strike K and maturity T_2 ($T_2 > T_1$).

Contract Terms. -T_1: maturity of the chooser option.
- At T_1, holder chooses either a call or a put expiring at T_2.
- Both options have strike K.

The chooser option delays the decision of whether upside or downside protection is more valuable until more information about the stock price is revealed.

Payoff at T_1. $\boxed{V_{T_1} = \max\{c, p\},}$
where c and p are the values at T_1 of a call and a put expiring at T_2. *At the decision date, the holder simply selects the more valuable option.*

Put–Call Parity Chooser Option at T_1. $\boxed{p = c + Ke^{-r(T_2-T_1)} - S_1 e^{-q(T_2-T_1)}}$
Put–call parity relates the values of the call and put at the choice date, adjusting the stock price for dividends and discounting the strike over the remaining life.
Payoff via Parity.

$$\max\{c, p\} = \max\Big\{c,\; c + Ke^{-r(T_2-T_1)} - S_1 e^{-q(T_2-T_1)}\Big\} = c + e^{-q(T_2-T_1)} \max\Big\{e^{-(r-q)(T_2-T_1)}K - S_1,\; 0\Big\}.$$

The chooser payoff decomposes into the call value plus an additional option that captures the value of waiting to decide.

2.2 Option Strategies

⋆ **Covered Call / Put Strategies.**

a) Long stock, short call
b) Short stock, long call
c) Long stock, long put
d) Short stock, short put

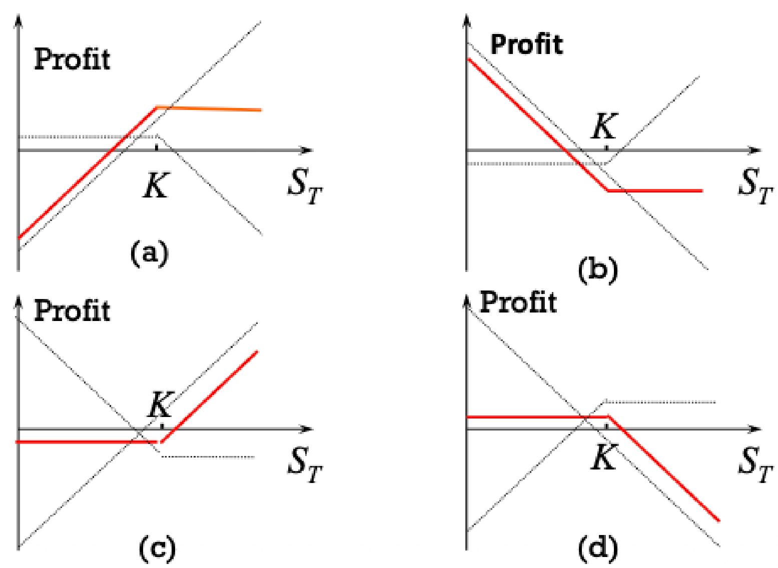

⋆ **Bull Spread.** Profits in an up market and Capped and floored profit/loss

Call Options

- Long call at K_1
- Short call at K_2.

Put Options

- Long put at K_1
- Short put at K_2.

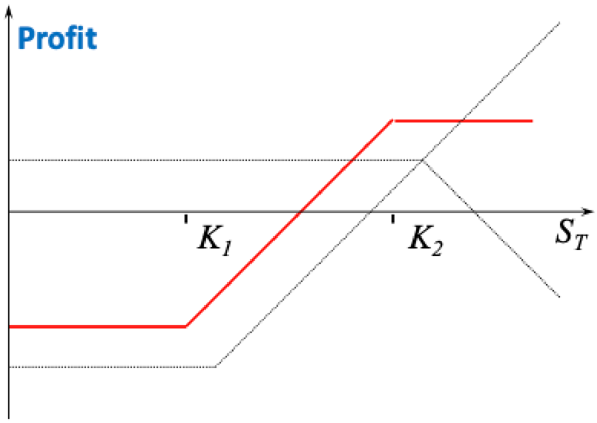

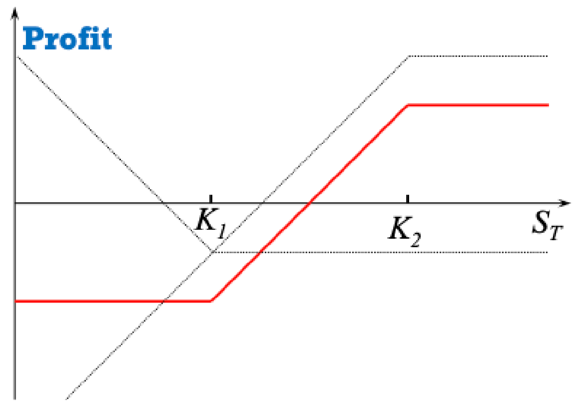

⋆ **Bear Spread.** Profits in a down market with capped upside and limited downside.

Call Options

- Short call at K_1.
- Long call at K_2.

Put Options

- Short put at K_1.
- Long put at K_2.

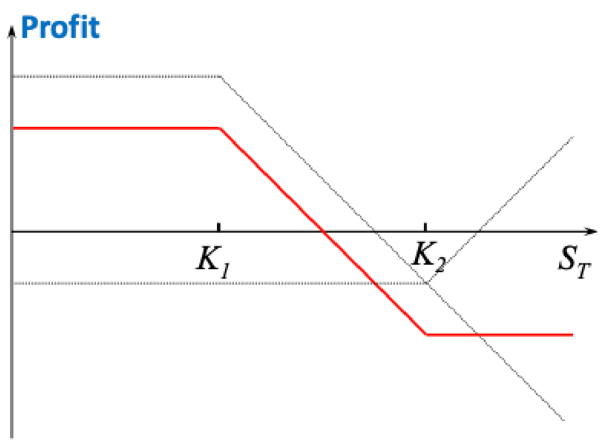

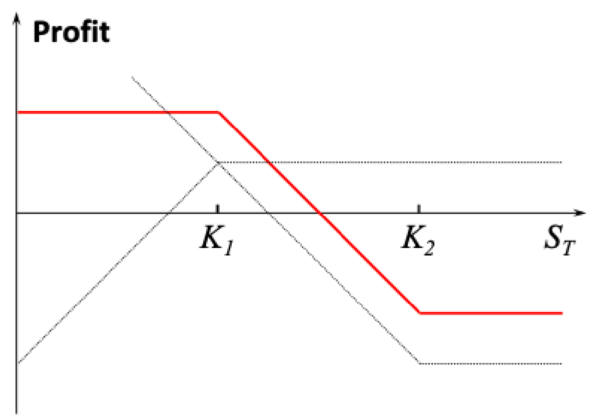

⋆ **Box Spread.** Option portfolio combining a bull spread and a bear spread with the same strikes (K_1, K_2) and maturity T.

Payoff at $T = K_2 - K_1$ (riskless). equals the difference between strike prices in all states:

$\boxed{\text{Value today} = e^{-rT}(K_2 - K_1).}$

A box spread replicates a risk-free bond: it pays a fixed amount $K_2 - K_1$ regardless of S_T. Its price must therefore be the present value of that fixed payoff, otherwise an arbitrage exists.

Call-Based Box 1) long bull call spread at K_1 and K_2 2) long bear call spread at K_1 and K_2

Maturity	Bull Spread	Bear Spread	Total
$S_T \leq K_1$	0	$K_2 - K_1$	$K_2 - K_1$
$K_1 < S_T < K_2$	$S_T - K_1$	$K_2 - S_T$	$K_2 - K_1$
$S_T \geq K_2$	$K_2 - K_1$	0	$K_2 - K_1$

Longing both spreads cancels the stock-direction exposure and leaves a pure fixed payoff. This is why box spreads are used as an arbitrage pricing check.

⋆ **Butterfly Spread**

Call Options

- Long call at K_1.
- Long call at K_3.
- Short 2 calls at $K_2 = 0.5(K_1 + K_3)$.

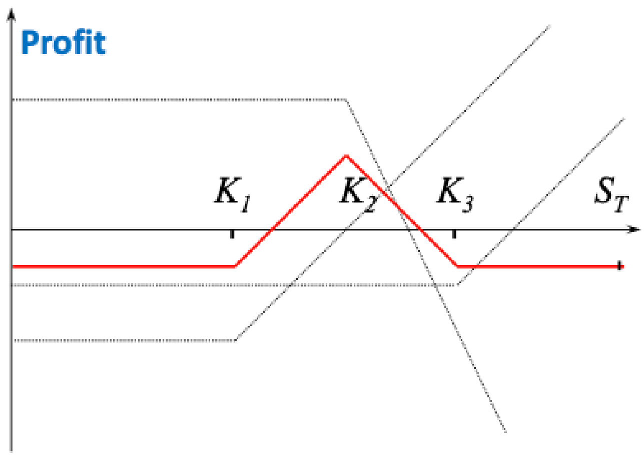

Put Options

- Long put at K_1.
- Long put at K_3.
- Short 2 puts at $K_2 = 0.5(K_1 + K_3)$

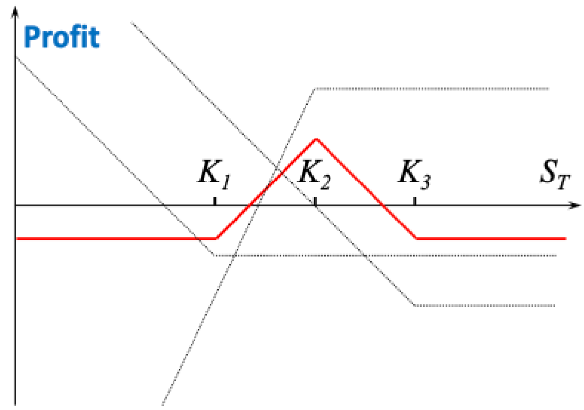

⋆ **Straddle.**

Positions

a) Long call at K

b) Long put at K

A straddle profits from large movements in either direction. Losses are limited to the total premium paid, and the payoff is symmetric around the strike.

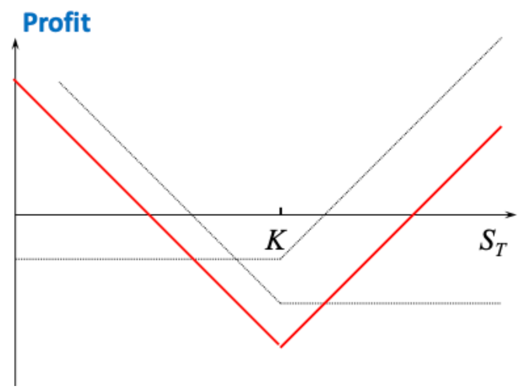

Positions

a) Long call at K_2

b) Long put at K_1, with $K_1 < K_2$

A long strangle profits from large price movements in either direction, like a straddle, but requires a bigger move to break even because the strikes differ. The cost is lower than a straddle, but the payoff region is wider and flatter around the center.

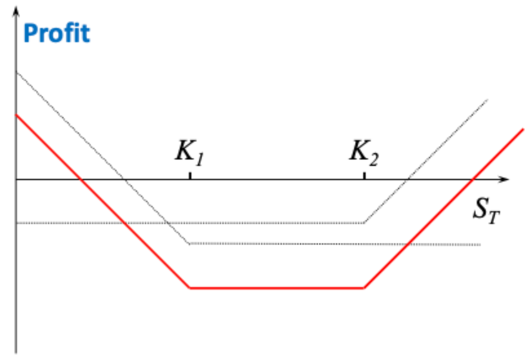

⋆ **Risk Reversal.**

Positions

- Long call at K_2
- Short put at K_1, with $K_1 < K_2$

A risk reversal creates asymmetric exposure: limited downside between K_1 and K_2 with full upside participation beyond K_2. The strategy embeds a directional view while reducing upfront cost by selling downside protection.

Application

- Hedge the risk of buying the underlying outside the interval $[K_1, K_2]$.

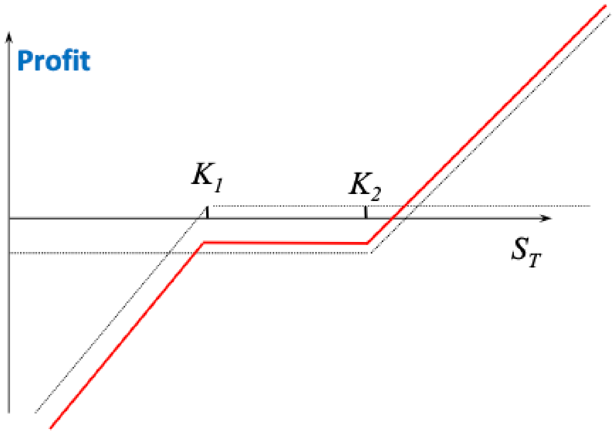

⋆ **Calendar Spread.**

Sell a short-term option and buy a longer-term option with the same strike K. Both options are either calls or puts.

Positions (using calls)

a) Short call at strike K with maturity T_1

b) Long call at strike K with maturity T_2, where $T_2 > T_1$

The long-dated option costs more than the short-dated option because it contains more time value. The strategy is constructed to profit from differences in time decay and volatility sensitivity across maturities.

Profit characteristics

- Profit is calculated at time T_1.
- Most profitable when the stock price remains near the strike K.

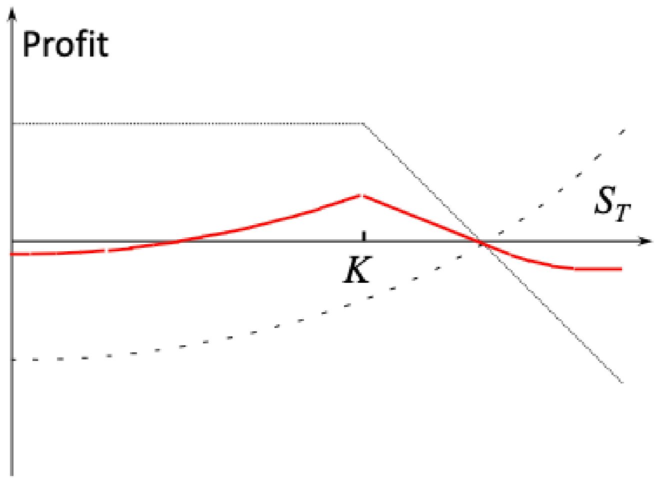

Why use a calendar spread

- **Time decay**: the short option loses value faster than the long option.
- **Volatility exposure**: the long option is more sensitive to volatility than the short option.
- **Defined risk**: maximum loss is known and limited to the net premium paid.

2.3 Binomial Pricing Model

⋆ **Binomial Pricing Model (Discrete-Time Option Pricing).** A **binomial model** describes asset price dynamics over discrete time, where the asset moves up or down each period and derivative prices are determined by **replication** and **no-arbitrage**.

The core idea is that uncertainty is simplified into two possible moves per period, yet prices are still pinned down uniquely by arbitrage arguments rather than investor preferences.

- **Asset price evolution.** At each step, the asset price moves:

$$S_0 \longrightarrow \begin{cases} S_0 u & \text{with probability } p, \\ S_0 d & \text{with probability } 1-p, \end{cases} \qquad u > 1,\ d < 1.$$

- **Volatility representation.** The magnitudes u and d encode asset volatility; higher volatility corresponds to larger separation between u and d.
- **Discrete-time structure.** Time is divided into periods, and pricing is performed node-by-node on the tree.

The binomial tree approximates continuous price paths while keeping the mathematics transparent and constructive.

⋆ **Redundant Security.** An option is a **redundant security** because it can be replicated using the underlying asset and cash.

Since the option adds no new payoff possibilities beyond what trading already provides, its price is fully determined by no-arbitrage.

⋆ **Risk-Neutral Pricing Logic.** Derivative values are obtained by forming a **riskless portfolio** of the underlying asset and cash that replicates the option's payoff.

- **Replication.** The option payoff can be exactly reproduced by trading the underlying asset and a risk-free bond.
- **Riskless portfolio.** At each node, choose holdings so that portfolio payoffs are identical in up and down states.
- **Discounting.** All certain cash flows are discounted at the **risk-free rate**.
- **Risk-neutral measure.** Under pricing, the asset's expected growth rate equals the risk-free rate (net of payouts).

Risk preferences disappear because any payoff that can be replicated must have the same price, regardless of who holds it.

⋆ **One-Period Binomial Tree: Notation.**

- S_0: current asset price
- u: up-move factor
- d: down-move factor
- f: derivative value at time 0
- f_{up}: derivative value in the up state
- f_{dn}: derivative value in the down state

$$S_0 \longrightarrow \begin{cases} S_0 u, \\ S_0 d, \end{cases} \qquad f \longrightarrow \begin{cases} f_{\text{up}}, \\ f_{\text{dn}}. \end{cases}$$

This single-step tree is the building block for multi-period binomial models and makes replication and pricing completely explicit.

⋆ **Valuing a Derivative in the Binomial Model (Stock options General Recipe).** To value any derivative in the binomial model, construct a **risk-free portfolio** that replicates the derivative's payoffs and price it by **no-arbitrage**.

Everything is driven by no-arbitrage: once you force the portfolio payoff to be the same in both states, its value must grow at the risk-free rate.

Long Δ shares,
Short 1 derivative (f).

Step 1: Specify payoffs (up/down states). $S_0 \longrightarrow \begin{cases} S_0 u, \\ S_0 d, \end{cases} \qquad f \longrightarrow \begin{cases} f_{\text{up}}, \\ f_{\text{dn}}. \end{cases}$

Step 2: Form a hedged portfolio. $\boxed{\Pi_0 = \Delta S_0 - f}$

Step 3: Portfolio value at maturity T (up/down). $\Pi_T = \begin{cases} \Pi_{\text{up}} = S_0 u\,\Delta - f_{\text{up}}, \\ \Pi_{\text{dn}} = S_0 d\,\Delta - f_{\text{dn}}. \end{cases}$

Step 4: Eliminate risk (match up = down).

$$\Pi_{\text{up}} = \Pi_{\text{dn}} \quad \Longleftrightarrow \quad \boxed{S_0 u\,\Delta - f_{\text{up}} = S_0 d\,\Delta - f_{\text{dn}}} \quad \Longrightarrow \quad \boxed{\Delta = \frac{f_{\text{up}} - f_{\text{dn}}}{S_0(u-d)}.}$$

Δ is chosen so the hedge cancels state-by-state uncertainty. This is the whole trick.

Step 5: Price the risk-free payoff by discounting.

$$\Pi_0 = e^{-rT}\Pi_T \qquad \Rightarrow \qquad \boxed{\Pi_0 = e^{-rT}\Pi_{\text{up}} = e^{-rT}\Pi_{\text{dn}}.}$$

Step 6: Recover the derivative price.

$$\Pi_0 = \Delta S_0 - f \qquad \Rightarrow \qquad \boxed{f = \Delta S_0 - e^{-rT}\Big(S_0 u\,\Delta - f_{\text{up}}\Big) = \Delta S_0 - e^{-rT}\Big(S_0 d\,\Delta - f_{\text{dn}}\Big).}$$

Once you can compute f_{up}, f_{dn}, you can price any payoff by repeating these steps.

⋆ **Risk-Neutral Valuation Formula.**
Equivalently, Substituting for Δ we obtain discounted **risk-neutral expected payoff**

$$\boxed{f = e^{-rT}\Big[p\,f_{\text{up}} + (1-p)\,f_{\text{dn}}\Big], \qquad p = \frac{e^{rT} - d}{u-d}.}$$

The probability p is chosen so the expected asset growth equals the risk-free rate; it is a pricing device, not a belief.

⋆ **Risk-Neutral Probability.** The **risk-neutral probability** p and $1-p$ are used to weight up and down payoffs in pricing, not to describe real-world beliefs.

$$\boxed{p = \frac{e^{rT} - d}{u-d}, \qquad 1 - p = \frac{u - e^{rT}}{u-d}.}$$

The **option value** is the present value of the expected payoff under the **risk-neutral world**:

$$\boxed{f = e^{-rT}\Big[p\,f_{\text{up}} + (1-p)\,f_{\text{dn}}\Big].}$$

This representation avoids explicit construction of Δ; pricing can be done directly from terminal payoffs without building the hedge or drawing the tree.

⋆ **Risk-Neutral Valuation (RN World).** Under the **risk-neutral measure**, the expected stock price at time T satisfies

$$\boxed{\mathbb{E}^{RN}[S_T] = S_0 e^{rT}}$$

In the RN world, all assets grow on average at the risk-free rate; risk premia are not priced through expected returns.

Risk-Neutral Probabilities and Stock Dynamics. With up and down movements,

$$S_T \longrightarrow \begin{cases} S_0 u & \text{with probability } p, \\ S_0 d & \text{with probability } 1-p, \end{cases} \qquad p = \frac{e^{rT} - d}{u - d}$$ These probabilities are chosen so that

$$\boxed{\mathbb{E}^{RN}[S_T] = p\, S_0 u + (1-p)\, S_0 d = S_0 e^{rT}.}$$

This equation is exactly what pins down the value of p; it is derived from no-arbitrage, not beliefs about the real world.

⋆ **Risk-Neutral Valuation of a Stock.** The stock value equals the discounted expected payoff under the **risk-neutral measure**:

$$V_{\text{STK}} = e^{-rT}\Big[p\, S_0 u + (1-p)\, S_0 d\Big].$$

Substitute the risk-neutral probability $\quad p \quad$ and $\quad 1-p$ to obtain

$$V_{\text{STK}} = e^{-rT}\left[\frac{e^{rT} - d}{u - d} u S_0 + \frac{u - e^{rT}}{u - d} d S_0\right] = S_0 e^{-rT}\left[\frac{ue^{rT} - ud}{u - d} + \frac{ud - de^{rT}}{u - d}\right]. \Rightarrow$$

$$\Rightarrow V_{\text{STK}} = S_0 e^{-rT}\left[\frac{e^{rT}(u - d)}{u - d}\right] = S_0 e^{-rT} e^{rT} = S_0.$$

The RN framework prices the stock consistently with its observed market value. The same probability p that prices derivatives also ensures the stock's expected growth equals the risk-free rate, confirming no-arbitrage consistency. Risk-neutral valuation does not predict prices; it enforces consistency between current prices and future payoffs.

Economic Implications.

- **Expected returns are irrelevant** for option pricing; current prices are taken as correct.
- **Risk is encoded in volatility**, through u and d, not through expected growth.
- **Real-world probabilities** do not enter derivative valuation.

The same risk-neutral probability p derived from replication ensures both correct stock pricing and correct derivative pricing.

⋆ **Interpretation of Δ Hedge Ratio = Option Equity Sensitivity** In the binomial model, Δ is the ratio of the change in the derivative value to the change in the underlying asset price between the **up** and **down** nodes.

$$\boxed{\Delta = \frac{f_{\text{up}} - f_{\text{dn}}}{S(u - d)} \approx \frac{\Delta V}{\Delta S}}$$

This is a **finite-difference slope:** it measures how sensitive the derivative value is to movements in the underlying asset over that local interval.

Multi-period note. In a multi-step tree, Δ is computed **at each node** using that node's S and continuation values $(f_{\text{up}}, f_{\text{dn}})$, so Δ generally **changes from node to node**. *This happens because the option's sensitivity depends on both the current underlying level and time-to-maturity (and the payoff's nonlinearity), so the local slope $\Delta V/\Delta S$ is different at different nodes.*

Sensitivity interpretation. Locally, a one-unit increase in the asset price changes the derivative value by approximately Δ units over that interval.

⋆ **American Option Valuation (Binomial Framework).** An **American option** can be exercised at any time up to maturity, so its value at each node must account for the possibility of **early exercise**.

Unlike European options, pricing requires comparing holding the option versus exercising immediately at every node.

Terminal Payoffs (at Maturity). At maturity T, the option value equals its **intrinsic value**:

$$V_T(S_T) = \begin{cases} \max(K - S_T, 0) & \text{(put)}, \\ \max(S_T - K, 0) & \text{(call)}. \end{cases}$$

At the final nodes there is no continuation value, so exercise is automatic if in the money.

Value at Node. At any node with underlying price S, the **continuation value** is

$$V^C(S) = e^{-r\Delta t}\Big[p\, V_{\text{up}} + (1-p)\, V_{\text{dn}}\Big], \qquad p = \frac{e^{r\Delta t} - d}{u - d}.$$

This is the value of keeping the option alive and pricing it as if it were European over the next step.

Early Exercise Value. The **early exercise (intrinsic) value** at the same node is

$$V^{EX}(S) = \begin{cases} K - S & \text{(put)}, \\ S - K & \text{(call)}. \end{cases}$$

This is the payoff obtained by exercising immediately at that node. EE is optimal only when exercising now dominates the value of waiting, i.e., when time value is sufficiently small.

⋆ **Node Value (American Option).** The value of an American option at each node is

$$\boxed{V(S) = \max\Big\{V^{EX}(S),\ V^C(S)\Big\}.}$$

Early exercise is optimal exactly when intrinsic value exceeds continuation value.

⋆ **Vlaution Backward Induction Procedure.**

1) Compute terminal payoffs at maturity. $V_T(S_T) = \begin{cases} \max(K - S_T, 0) & \text{(put)}, \\ \max(S_T - K, 0) & \text{(call)}. \end{cases}$

2) Move backward one step at a time. Let S denote the stock price at a generic interior node with time step Δt.

3) At each node compute:

Continuation value via RN expectation : $V^C(S) = e^{-r\Delta t}\Big[p\, V_{\text{up}} + (1-p)\, V_{\text{dn}}\Big],$

Intrinsic (early exercise) value: $V^{EX}(S) = \begin{cases} K - S & \text{(put)}, \\ S - K & \text{(call)}. \end{cases}$

Then set the node value $\boxed{V(S) = \max\{V^{EX}(S),\, V^C(S)\}.}$

4) Continue until time 0

This backward recursion is the defining feature of American option pricing.

⋆ **BL: Continuous Dividend Yield q.** For a dividend-paying stock, the **risk-neutral growth rate** of the stock price equals the **forward growth rate**:

$$\boxed{\text{Equity growth rate in RN world} = r - q.}$$

Dividends reduce capital gains: part of the return is paid out as cash, so the stock price grows more slowly than the risk-free rate.

Expected stock price satisfies: $\mathbb{E}^{RN}[S_T] = p\, S_0 u + (1-p)\, S_0 d = S_0 e^{(r-q)T}$.

The expected growth rate is reduced by q because dividends are paid out continuously rather than reinvested in the stock price.

RN Probability with Dividends: $\boxed{p = \dfrac{e^{(r-q)T} - d}{u - d},}$

This choice of p ensures no-arbitrage: after accounting for dividends, the expected stock price grows at the forward rate $r - q$.

⋆ **Corporate Securities as Options (Structural / Merton View).** (using the firm's value to value its securities) Model the firm's **total asset value** as V. The firm's securities are **claims** on V, and total firm value satisfies

$$\boxed{V = E + D.}$$

Equity and debt split the firm's asset value according to priority: debt gets paid first, equity gets what is left.

Priority of Claims: Makes equity look like an option: equity holders benefit from upside but can walk away in default

1) **Debt** has **first claim** on the firm's assets (limited upside).
2) **Equity** has the **residual claim** after debt is paid.

Payoffs at Debt Maturity T (No Coupons). Let F be the face value of debt due at T. Then terminal payoffs are

$$\boxed{D_T = \min(F, V_T),} \qquad \boxed{E_T = \max(V_T - F, 0).}$$

Debt is like a risk-free bond minus a put on the firm; equity is like a call on the firm's assets with strike F.

RN Binomial Dynamics for Firm Value. Assume V treated as the underlying traded asset and follows a risk-neutral binomial pricing process over step Δt:

$$V \longrightarrow \begin{cases} uV & \text{with prob. } p, \\ dV & \text{with prob. } 1-p, \end{cases} \qquad p = \frac{e^{r\Delta t} - d}{u - d}.$$

Once V is treated as the underlying traded asset, equity and debt values follow by standard risk-neutral pricing.

⋆ **Valuation by Backward Induction (General Case).** At any node with current firm value V, compute values one step ahead and discount.

Debt Valuation (with Coupon c per step). $\boxed{D^C(V) = e^{-r\Delta t}\Big[p\, D_{\text{up}} + (1-p)\, D_{\text{dn}}\Big] + c.}$

Equity Valuation (with Dividend δ per step). $\boxed{E^C(V) = e^{-r\Delta t}\Big[p\, E_{\text{up}} + (1-p)\, E_{\text{dn}}\Big] + \delta.}$

Coupons and dividends are simply added as cash flows at each step; the continuation value is discounted risk-neutral expectation.

Terminal Conditions at Final Time T.

$$D_T = \min(F, V_T)\ (+c_T \text{ if paid at } T), \qquad E_T = \max(V_T - F, 0)\ (+\delta_T \text{ if paid at } T).$$

Consistency Check. At every node (including time 0), $\boxed{V = E + D.}$

The decomposition is an accounting identity and also a pricing check: if your computed E and D do not sum to V at each node, something is wrong in the recursion or cash-flow timing.

2.4 Futures Options

⋆ **Futures Contract.** Requires no initial investment; the futures contract is zero-cost to enter. $K = F_0 = S_0 e^{rT} = 0 \Rightarrow$ 1) Expected return on the futures price is zero. 2) Futures price growth rate is therefore zero.

Because no cash is exchanged at initiation, the futures price itself must have zero expected growth under the risk-neutral measure.

⋆ **Futures Option** An option written on a futures contract, giving the holder the right (but not the obligation) to enter a long or short futures position at a specified futures price K at maturity.Option payoffs depend only on F_T, not directly on the spot price.

Unlike stock options, the underlying is a futures contract rather than the asset itself, so the option payoff depends on the futures price at maturity, not the spot price.

$F_T =$ futures price at option maturity, $\qquad K =$ option strike (futures price).

Call Option on a Futures Contract. $\boxed{\text{Payoff at } T: \quad \max(F_T - K,\, 0).}$

Exercising a futures call gives a long futures position at price K, which is valuable when the market futures price exceeds K.

Put Option on a Futures Contract. $\boxed{\text{Payoff at } T: \quad \max(K - F_T,\, 0).}$

Exercising a futures put gives a short futures position at price K, which is valuable when the futures price falls below K.

⋆ **Binomial Tree for Futures Price.** Let $F_0 =$ futures price at $t = 0$, $\quad f =$ option price on the futures. At maturity T,

$$F_0 \to \begin{cases} F_0 u & \text{up state,} \\ F_0 d & \text{down state,} \end{cases} \qquad f \to \begin{cases} f_{\text{up}}, \\ f_{\text{dn}}. \end{cases}$$

Replicating (Riskless) Portfolio. A portfolio constructed to eliminate uncertainty in payoffs across states:

1) Long Δ futures contracts 2) Short one futures option we wre Replicating risk free Portfolio

Portfolio Payoff at Maturity $-f \longrightarrow \begin{cases} \Delta(F_0 u - F_0) - f_{\text{up}}, \\ \Delta(F_0 d - F_0) - f_{\text{dn}}. \end{cases}$

The futures position is chosen to exactly hedge the option's exposure to changes in the futures price. By selecting Δ so that payoffs are identical in the up and down states, the combined portfolio delivers the same value in all states and therefore replicates a risk-free payoff.

Riskless Portfolio Condition. Choose Δ so the payoffs are equal in the up and down states:

$$\boxed{\Delta = \frac{f_{\text{up}} - f_{\text{dn}}}{F_0 u - F_0 d}.}$$

With this choice of Δ, the gains from the futures position exactly offset the difference in option payoffs across states, so the portfolio's terminal value is identical whether the futures price moves up or down. Because the payoff no longer depends on the realized state, the portfolio is riskless and must earn the risk-free rate.

Option Value Today (Riskless). Since the constructed portfolio is riskless and a futures contract has zero value at initiation, the value of the option today equals the **discounted value of the portfolio's certain payoff**:

$$\boxed{f = -\Big[\Delta(F_0u - F_0) - f_{\text{up}}\Big]e^{-rT}.}$$

The term in brackets is the common payoff of the hedged portfolio at maturity, which is identical in all states because Δ was chosen to eliminate state dependence. The minus sign appears because the option is held short in the replicating portfolio. By no-arbitrage, a riskless payoff must be discounted at the risk-free rate r, and the futures position contributes no initial cost, leaving the option price as the present value of this guaranteed payoff.

$\star$ **Risk-Neutral Valuation Formula (Futures Option).** Substituting the hedge ratio Δ into the value of the riskless replicating portfolio yields

$$\boxed{f = \Big[pf_{\text{up}} + (1-p)f_{\text{dn}}\Big]e^{-rT},} \qquad \boxed{p = \frac{1-d}{u-d}.}$$

The risk-neutral probability differs from stock options because futures have zero expected drift under the risk-neutral measure (futures price equal to zero). Under this probability, the expected payoff of the option equals the payoff of the riskless replicating portfolio, so the option value is the discounted expected payoff at the risk-free rate. The formula therefore prices the option by replacing uncertainty with a weighted average of future payoffs that enforces no-arbitrage, rather than by forecasting actual market outcomes.

2.5 Analytical Binomial Model

⋆ **Choosing u and d.**: **Binomial Process.** A one-period binomial model is defined by three parameters (p, u, d). The goal is to match the first two moments of the underlying return distribution using a discrete process.

Moment-Matching Conditions. Let the stock price evolve from S_0 to S_0u or S_0d over a time step Δt.

Stock Value; Mean (expected growth):

$$pS_0u + (1-p)S_0d = S_0e^{\mu\Delta t} \quad \Rightarrow \quad p = \frac{e^{\mu\Delta t} - d}{u-d}.$$

Stock Variance:

$$pu^2 + (1-p)d^2 - \Big[pu + (1-p)d\Big]^2 = \sigma^2\Delta t.$$

These two equations impose only two restrictions on 3 unknowns, so an additional normalization is required.

Risk-Neutral World. Replace the real drift μ with the risk-free rate r: $p = \frac{e^{r\Delta t}-d}{u-d}$.

This enforces the no-arbitrage condition that discounted prices are martingales under the risk-neutral measure.

Variance Check. Substituting into the variance condition:

$$e^{\mu\Delta t}(u+d) - ud - e^{2\mu\Delta t} = \sigma^2\Delta t,$$

which is satisfied to first order in Δt.

Thus the CRR parameters match both the mean (under risk-neutral drift) and variance of continuous-time geometric Brownian motion as $\Delta t \to 0$.

⋆ **Cox–Ross–Rubinstein (CRR) Choice.** Impose the normalization

$$ud = 1 \qquad \Rightarrow$$

$$\boxed{u = e^{\sigma\sqrt{\Delta t}}, \qquad d = e^{-\sigma\sqrt{\Delta t}}.}$$

The CRR choice enforces a symmetric binomial move in log-prices, so an up move followed by a down move returns the price to its original level. *By tying u and d directly to volatility σ and the time step Δt, the binomial process matches the variance of a **geometric Brownian motion** to first order. The mean is then matched separately via the risk-neutral probability p, ensuring no-arbitrage pricing.*

⋆ **Binomial Option Pricing Formula (CRR):**

Probability of j Ups and $(n-j)$ Downs. $\boxed{\Pr(j \text{ ups}) = \binom{n}{j} p^j (1-p)^{n-j}.}$

In an n-step binomial tree, each period is an independent up/down move. There are $\binom{n}{j}$ distinct paths with exactly j up moves, each occurring with probability $p^j(1-p)^{n-j}$ under the risk-neutral measure.

Stock Price at Maturity.

$$\boxed{S_T = S_0 u^j d^{n-j}.}$$

Under the CRR model, prices evolve multiplicatively: each up move scales the price by u, each down move by d. Because the tree recombines, the terminal price depends only on the number of up moves, not their order.

European Call Value with n Steps.

$$\boxed{V_n = e^{-rT} \underbrace{\sum_{j=0}^{n} \binom{n}{j} p^j (1-p)^{n-j}}_{\text{Binomial Prob. } (j \text{ u's, } n-j \text{ d's})} \underbrace{\max\Big(S_0 u^j d^{n-j} - K,\, 0\Big)}_{\text{Option Payoff}}.}$$

This is the discounted risk-neutral expected payoff. Each terminal payoff is weighted by its binomial probability and then discounted at the risk-free rate, reflecting no-arbitrage pricing.

Critical In-the-Money Condition (Call). $\boxed{S_0 u^j d^{n-j} \geq K.}$

Only terminal nodes where the stock price exceeds the strike contribute positively to the option value; all others contribute zero.

Use Cox–Ross–Rubinstein (CRR) specifies $u = e^{\sigma\sqrt{\Delta t}}, \qquad d = e^{-\sigma\sqrt{\Delta t}}, \qquad p = \frac{e^{r\Delta t} - d}{u - d}.$
With these choices, the binomial tree matches the variance of geometric Brownian motion and enforces no-arbitrage via the risk-neutral probability. The formula above is simply the CRR tree written in closed form: instead of backward induction, it aggregates all terminal nodes using binomial probabilities. As $n \to \infty$, this expression converges to the Black–Scholes price.

⋆ **Jarrow–Rudd (JR) Binomial Model.** An alternative binomial parametrization where the probability of an up move is fixed exogenously and the drift is absorbed into the up and down factors. Instead of solving for the risk-neutral probability p (as in CRR), the

Jarrow–Rudd model sets $\boxed{p = \frac{1}{2}}$ and chooses u and d so that the discrete-time log-return matches the mean and variance of the continuous-time lognormal process.

Up and Down Factors. $\boxed{u = \exp\Big(\big(r - q - \tfrac{1}{2}\sigma^2\big)\Delta t + \sigma\sqrt{\Delta t}\Big), \qquad d = \exp\Big(\big(r - q - \tfrac{1}{2}\sigma^2\big)\Delta t - \sigma\sqrt{\Delta t}\Big)}$

Interpretation. The binomial log-return over one time step is symmetric around its mean, with equal probability of moving up or down.

Moment Matching. This construction ensures:

$$\textbf{Expected value of ln(S):} \quad \mathbb{E}[\ln S_{t+\Delta t} - \ln S_t] = (r - q - \tfrac{1}{2}\sigma^2)\Delta t,$$

$$\textbf{Volatility of ln(S):} \quad \mathrm{Var}[\ln S_{t+\Delta t} - \ln S_t] = \sigma^2 \Delta t.$$

Why It Works. By matching the first two moments of the continuous-time lognormal process directly, the JR tree converges to the same limit as Black–Scholes as $\Delta t \to 0$, even though probabilities are fixed.

⋆ **Comparison to CRR.**

- CRR: choose u, d symmetrically and solve for p.
- JR: fix $p = \frac{1}{2}$ and encode drift in u, d.

Both models converge to Black–Scholes, but distribute drift and randomness differently at finite step sizes.

CRR vs. Jarrow–Rudd: Interpretation of the Comparison. Both the Cox–Ross–Rubinstein (CRR) and Jarrow–Rudd (JR) binomial trees are constructed to converge to the same continuous-time limit (Black–Scholes) as the number of steps increases. However, for any finite number of steps, they distribute drift and randomness differently.

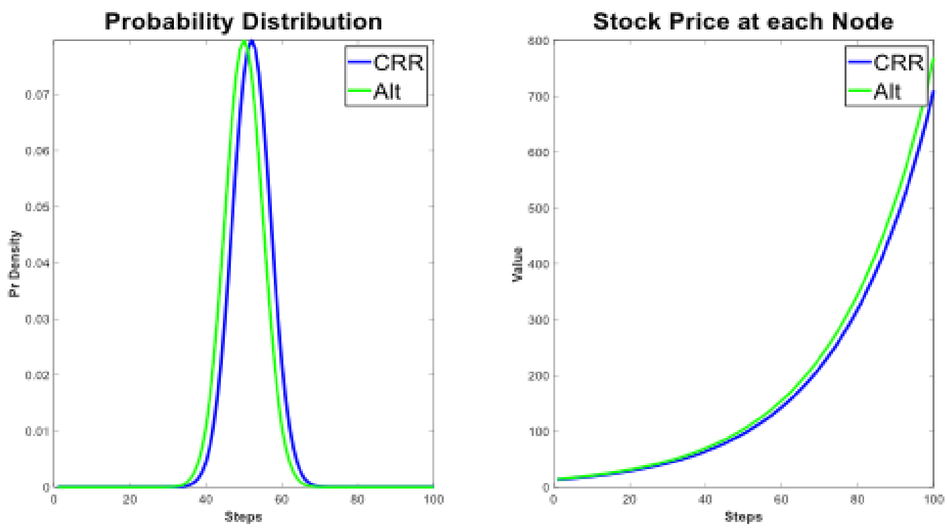

Probability Distribution (Left Panel). In the CRR model, the risk-neutral probability p is endogenously chosen and typically differs from 1/2. This results with smaller step sizes u and d in a slightly higher probability of an up move,

$$p_{\text{CRR}} > \tfrac{1}{2}, \quad \text{while} \quad p_{\text{JR}} = \tfrac{1}{2},$$

JR model adjusts the magnitudes of u and d to encode the drift. As a result, the JR distribution is symmetric in probabilities but asymmetric in step sizes.

Stock Price Paths (Right Panel). Because CRR uses smaller up and down factors with a higher probability of moving up, its stock price path grows more smoothly over time. The JR model uses larger step magnitudes, which leads to slightly more dispersion in node values at finite horizons, even though the expected growth rate is the same.

2.6 Option Tools

⋆ **Implied Dividend Yield (EU)** Dividend yield inferred from quoted European call and put prices with the same strike K and maturity T.

Put–Call Parity. $c + Ke^{-rT} = p + S_0 e^{-qT}$.

Solving for the Implied Dividend Yield. Rearranging parity gives $e^{-qT} = \frac{c-p+Ke^{-rT}}{S_0} \Rightarrow$

$$\Rightarrow \boxed{q = -\frac{1}{T}\ln\left(\frac{c-p+Ke^{-rT}}{S_0}\right).}$$

The implied dividend yield is the value of q that makes observed call and put prices consistent with no-arbitrage. It reflects the market's expectation of dividend payouts embedded in option prices, not necessarily the firm's announced dividend policy.

⋆ **Risk-Neutral Probabilities.** In a one-period binomial model, the **risk-neutral probability** is

$$\boxed{p = \frac{a-d}{u-d}.}$$

Non-dividend-paying stock: $a = e^{r\Delta t}$

Dividend-yield-paying stock: $a = e^{(r-q)\Delta t}$

Currency contract: $a = e^{(r-r_f)\Delta t}$

Futures contract: $a = 1$

The parameter a represents the risk-neutral growth factor of the underlying asset. It differs across assets because no-arbitrage growth depends on carry: dividends reduce expected growth for stocks, foreign interest offsets domestic carry for currencies, and futures have zero expected growth since they require no initial investment.

Chapter 3

Continuous Time Models

3.1 Log Returns

⋆ **Stock Return Decomposition.** A **gross return** over a period $[0, T]$ can be decomposed into a product of sequential sub-period returns ($\frac{S_T}{S_0} = \prod_{t=1}^{T} \frac{S_t}{S_{t-1}}$) →.

$$\frac{S_T}{S_0} = \left(\frac{S_1}{S_0}\right)\left(\frac{S_2}{S_1}\right)\cdots\left(\frac{S_T}{S_{T-1}}\right)$$

This multiplicative property shows that the total growth of a stock is the result of compounded returns over smaller time increments.

⋆ **Log Returns.** The **logarithmic return** is the natural log of the gross return ($\ln\frac{S_T}{S_0} = \sum_{t=1}^{T} \ln\frac{S_t}{S_{t-1}}$).

$$\boxed{\ln\left(\frac{S_T}{S_0}\right) = \sum_{t=1}^{T} \ln\left(\frac{S_t}{S_{t-1}}\right)}$$

Taking the logarithm transforms the multiplicative process of stock prices into an additive process, which is mathematically easier to model using the Central Limit Theorem.

⋆ **Log-Normal Distribution.** If log returns are **normally distributed**, then the **stock price** itself follows a **log-normal distribution**.

- **Bounded at zero**: $S_T \geq 0$, consistent with limited liability.
- **Right-skewed**: Prices have a long right tail (high potential upside).
- **Kurtonic**: Features higher peak and fatter tails than a normal distribution.

Because stock prices cannot drop below zero but have unlimited upside, the log-normal distribution provides a more realistic model than the normal distribution for price levels, capturing the inherent asymmetry of financial markets.

3.2 Probabilty Reviewe

⋆ **Normal Probability Density Function (PDF).** The density $\phi(x|\mu,\sigma)$ describes the likelihood of a continuous random variable X taking a specific value $\boxed{\phi(x|\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}}$

Moments / Statistics:
Mean (μ): Center of the distribution.
Variance (σ^2): Spread of the distribution.
Skew (0): Perfect symmetry.
Kurtosis (3): Measure of "tailedness."
The normal distribution is defined entirely by its mean and variance. In finance, it is the standard model for log returns because of its symmetric properties and mathematical simplicity.

⋆ **Normal Cumulative Distribution Function (CDF).** The function $N(a)$ represents the probability that a normal variable is less than or equal to a.

$$N(a) = \int_{-\infty}^{a} \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} dx$$

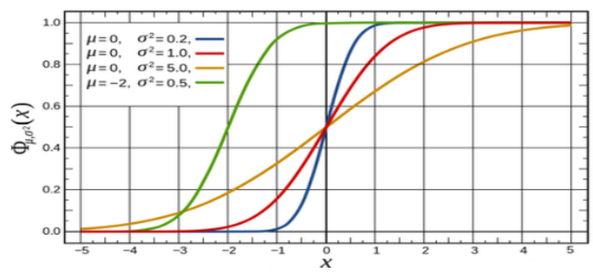

- **Bounded**: Values stay between $[0,1]$.
- **Volatility Impact**: Higher volatility σ leads to a **flatter slope** in the CDF curve.
- **Inverse Function**: $N(N(\alpha))^{-1} = \alpha$ returns the critical value for a given probability.

The CDF is essential for calculating the probability of a stock's log return falling below a certain threshold, such as in Value-at-Risk (VaR) calculations. Visually, as volatility increases, the "S-curve" stretches out, indicating a higher probability of extreme outcomes.

⋆ **Properties of Normal Variables.** If $Y_1 \sim \phi(\mu_1,\sigma_1)$ and $Y_2 \sim \phi(\mu_2,\sigma_2)$:

- **Linear Transformation**: $a + bY_1 \sim \phi(a + b\mu_1, b\sigma_1)$.
- **Summation**: $Y_1 + Y_2 \sim \phi(\mu_1 + \mu_2, \sqrt{\sigma_1^2 + \sigma_2^2 + 2\rho\sigma_1\sigma_2})$.
- **Covariance**: $Cov(Y_1, Y_2) = \sigma_1\sigma_2\rho$ where ρ is **correlation**.

Critical rule in finance is that **variance is additive, but volatility (standard deviation)**is not.

⋆ **Log-Normal Distribution:** Random variable X is **log-normally distributed** if $Y = \ln(X)$ follows a **normal distribution** $Y \sim \mathcal{N}(\mu,\sigma^2)$.

PDF: $\boxed{f(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right)}$

Moments: Unlike the normal distribution, the mean and variance are functions of both μ and σ:

- **Mean (Expectation):** $\mathbb{E}[X] = e^{\mu + \frac{1}{2}\sigma^2}$.
- **Variance:** $\text{Var}(X) = (e^{\sigma^2} - 1)e^{2\mu + \sigma^2}$.
- **Skewness:** $\gamma = (e^{\sigma^2} + 2)\sqrt{e^{\sigma^2} - 1}$.

Asymmetry: The distribution is always **right-skewed**. The degree of skewness increases as σ increases.
Multiplicative Reproducibility: If $X_1, X_2, \ldots, X_n$ are independent log-normal variables, their product $\prod X_i$ is also log-normally distributed.

In financial engineering, these properties imply that while the "average" log return might be μ, the expected price level is boosted by a convexity term ($\frac{1}{2}\sigma^2$). This is a direct consequence of Jensen's Inequality when moving from log-space back to price-space.

⋆ **Jensen's Inequality in Financial Engineering.** Jensen's Inequality describes the relationship between the **average of a function** and the **function of an average**. In finance, this explains the gap between log returns and price levels.

- **Convex Case (Prices)**: Since $f(x) = e^x$ is convex, $\mathbb{E}[e^X] \geq e^{\mathbb{E}[X]}$.
- **Concave Case (Log Returns)**: Since $f(x) = \ln(x)$ is concave, $\mathbb{E}[\ln(S_T)] \leq \ln(\mathbb{E}[S_T])$.

Applying this to stock prices where the log return $X \sim \mathcal{N}(\mu, \sigma^2)$: $\boxed{\mathbb{E}[S_T] = e^{\mu + \frac{1}{2}\sigma^2} > e^{\mu}}$

Implications:

- **Expected Price (LHS)**: $\mathbb{E}[e^X] = e^{\mu + \frac{1}{2}\sigma^2}$. $\rightarrow$
  $\rightarrow$ Implies the actual average future wealth is boosted by volatility.
- **Exponential of Expected Return (RHS)**: $e^{\mathbb{E}[X]} = e^{\mu}$. $\rightarrow$
  $\rightarrow$ Implies the "typical" or median outcome ignores the benefit of positive outliers.
- **The Gap**: $\mu < \mu + \frac{1}{2}\sigma^2$. The difference, $e^{\frac{1}{2}\sigma^2}$, is the **convexity adjustment**. $\rightarrow$
  $\rightarrow$ Implies that to reach a target expected price, the required log return is actually lower than one might intuitively think.

Practical & Engineering Significance:

- **Arithmetic vs. Geometric**: The arithmetic average return (price expectation) is always higher than the geometric average (log return/median outcome).
- **Volatility Drag**: In a symmetric normal distribution of returns, the "typical" or median outcome (μ) is lower than the mean outcome. This is because extreme positive outliers in the right tail of the price distribution pull the average up.
- **Path Underperformance**: Assuming log returns are $\mathcal{N}(\mu, \sigma^2)$ implies that the most likely path (μ) will underperform the expected value of the investment due to the way variance "drags" on compounded growth.

The expected price is higher than the price of the expected return. Volatility creates "upward bias" in the mean price even if the median return stays the same.

⋆ **Brownian Motion.** A continuous-time **random walk** (B_t) comprised of independent increments (dz_t) drawn from a normal distribution.

- **Path Properties**: The process is continuous but **non-differentiable**, meaning paths are jagged at every scale.
- **Independent Increments**: $\boxed{\mathbb{E}[B_t, B_{t-1}] = 0}$ Increments over disjoint time intervals are independent;
- **Normal Increments**: $\boxed{B_t - B_s \sim \mathcal{N}(0, t-s) \text{ for } t > s}$ Follows a normal distribution
- **Zero Drift**: $\boxed{\mathbb{E}[dB_t] = 0}$ The expected change is null
- **Variance**: $\boxed{\mathrm{Var}(B_t) = t}$ Grows linearly with time, such that
- **Initialization**: $\boxed{B_0 = 0}$ The process traditionally begins at the origin,

Brownian motion effectively mimics the constant, random fluctuations of financial markets. Its non-differentiable nature captures the infinite "jaggedness" of price movements where trends can shift at any infinitesimal time step.

⋆ **Stochastic Dynamics.** The evolution of the path is determined by the summation of these random normal shocks.

$$\boxed{B_t = B_{t-1} + dz_t} \quad \text{or} \quad \boxed{dB_t = dz_t}$$

Component Definitions:

- **Increment** (dz_t): Defined mathematically as $\boxed{dz_t = \epsilon_t \sqrt{dt}}$
- **Random Shock** (ϵ_t): A standard normal variable where $\boxed{\epsilon_t \sim \phi(0,1)}$.
- **Step Size** (dt): Represents the infinitesimal time increment.

The inclusion of the $\sqrt{dt}$ term is the most critical feature; it ensures that while the standard deviation (volatility) scales with the square root of time, the variance scales linearly, maintaining statistical consistency across different time horizons.

⋆ **Itô Process.** A stochastic process that incorporates both **deterministic** and **random components**, making it suitable for modeling dynamic variables

$$\boxed{dx = a(x,t)dt + b(x,t)dz}$$

Variables:
- x: The state variable being modeled (e.g., a stock price or interest rate).
- $a(x,t)$: **Drift** component, which is a function of time and the state variable.
- $b(x,t)$: **Volatility** component, also a function of time and the state variable.
- dz: **Random term**, defined as $\epsilon\sqrt{dt}$.

This formula is the "engine" for modeling financial assets. The drift $a(x,t)$ represents the expected trend, while $b(x,t)dz$ adds the unpredictable market noise.

⋆ **Itô's Lemma.** Commonly referred to as the "Calculus of stochastic functions," Itô's Lemma provides the mathematical framework for **finding the derivative** of a function $G(x,t)$ that contains stochastic elements.

Practical Applications:

Equities: If a stock price follows a stochastic process, an equity option is a function of that price. Itô's Lemma allows us to calculate the derivative of the option price relative to the stock.

Fixed Income: If an interest rate follows a stochastic process, a bond price is a function of that rate. The Lemma provides the derivative of the bond price with respect to the interest rate.

Think of Itô's Lemma as the stochastic version of the "Chain Rule." Because market paths are non-differentiable (jagged), standard calculus fails; Itô's Lemma adjusts for the extra variance caused by

Brownian motion.

⋆ **The "Itô Term":** Unlike standard calculus, we must include the **second-order partial derivative** because dx^2 does not vanish in stochastic settings.

$$\boxed{dG(x,t) = \frac{\partial G}{\partial x}dx + \frac{\partial G}{\partial t}dt + \underbrace{\frac{1}{2}\frac{\partial^2 G}{\partial x^2}dx^2}_{\text{Itô Term}}}$$

Why this matters: Standard calculus assumes changes are smooth ($dx^2 \approx 0$). In finance, the "jagged" nature of Brownian motion means small shocks squared (dx^2) actually accumulate into a meaningful "volatility drag" or "convexity adjustment" over time.

⋆ **Intuition for Stochastic Derivatives.** To find the final form of the derivative, we expand the function $G(x,t)$ via a **Taylor series** and evaluate the terms as $dt \to 0$:

- **First-order terms**: dt and dx are kept.
- **Higher-order time**: Any term with dt^b where $b > 1$ (like dt^2 or $dxdt$) tends to 0 and is ignored.
- **The Stochastic Exception**: Because $dz^2 \to dt$, the term dx^2 simplifies to $b(x,t)^2 dt$.

The Final Derivative Form: By substituting dx and dx^2 into the expansion, we get the complete formula for the change in our function G:

$$\boxed{dG = \left(\frac{\partial G}{\partial t} + a(x,t)\frac{\partial G}{\partial x} + \frac{1}{2}b(x,t)^2\frac{\partial^2 G}{\partial x^2}\right)dt + b(x,t)\frac{\partial G}{\partial x}dz}$$

In practice: This allows us to link the movement of an option (G) to the movement of its underlying stock (x). The terms in the parentheses represent the "deterministic" evolution of the option price over time.

⋆ **The Geometry of Itô's Lemma.** The "Itô Term" is fundamentally a **convexity adjustment**. It exists because in a stochastic world, volatility interacting with a curved payoff creates a directional bias that standard calculus misses.

Why "Convexity" $\frac{1}{2}b^2\frac{\partial^2 G}{\partial x^2}$? (Itô / Convexity Term) Convexity refers to the **curvature** of the function $G(x,t)$.

- **Non-Linearity**: If G is convex (curved upward like a Call option), a move "up" in the underlying x gains more value than an equal move "down" loses.
- **Volatility Interaction**: Because Brownian motion is "jagged," the variable x is constantly realized at points away from the mean.
- **The Result**: The Itô term $\frac{1}{2}b^2 G_{xx}$ captures this "bonus" value. It represents the **drift boost** caused by volatility (b^2) hitting that curve (G_{xx}).

This means that even if the underlying asset has zero drift ($a = 0$), a convex derivative (G) will naturally "drift" upward over time simply because of the variance of the underlying.

3.3 GBM

⋆ **Itô's Lemma for GBM.** Geometric Brownian Motion (GBM) is a **log-normal stock price process**. It is defined by the following Itô process:

$$\boxed{dS = \mu S dt + \sigma S dZ}$$

Drift: $a(S,t) = \mu S$. and **Volatility**: $b(S,t) = \sigma S$.

Stochastic process for a function $G(S,t)$: Applying Itô's Lemma to a function of the stock price:

$$\boxed{dG(S,t) = \left(\frac{\partial G}{\partial t} + \mu S \frac{\partial G}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 G}{\partial S^2}\right) dt + \frac{\partial G}{\partial S}\sigma S dZ}$$

Logic of Itô's Lemma for GBM. This formula allows us to calculate the **total change** (dG) in the value of a derivative (like an option) when the underlying stock price S follows Geometric Brownian Motion ($dS = \mu S dt + \sigma S dZ$).

Breaking Down the Terms: The formula is split into **Deterministic Part** (proportional to dt) and a **Stochastic Part** (proportional to dZ):

- $\frac{\partial G}{\partial t}$: **Time Decay (Theta)**. How the value of G changes just because time passes, even if the stock price stays still.

- $\mu S \frac{\partial G}{\partial S}$: **Drift Exposure**. The stock's expected growth (μS) multiplied by the sensitivity of the derivative to the stock price ($\frac{\partial G}{\partial S}$, or Delta).

- $\frac{1}{2}\sigma^2 S^2 \frac{\partial^2 G}{\partial S^2}$: **Convexity Adjustment (Ito Term)**. This accounts for the extra value gained from the stock's volatility ($\sigma^2 S^2$) interacting with the curvature of the payoff ($\frac{\partial^2 G}{\partial S^2}$, or Gamma).

- $\frac{\partial G}{\partial S}\sigma S dZ$: **Random Noise**. The unpredictable part of the change, driven by market shocks (dZ) scaled by the stock's volatility and the derivative's Delta.

Logic: In the "smooth" world of standard calculus, we would ignore the S^2 term. In the stochastic world, because the stock price is "jagged," those small squared moves (dS^2) add up to a significant constant drift (dt), which is why the third term inside the parentheses is required.

⋆ **GBM Properties.** Geometric Brownian Motion (GBM) has specific statistical properties regarding its expected price and returns. Because the distribution is **Log-Normal** (and thus skewed), the mean and median results differ.

1. Price Statistics:

Mean Price: $E[S_T] = S_0 e^{\mu T}$. (get growth rate)

Median Price: For $\alpha = 0.5 \to N(0.5)^{-1} = 0$, the median is $S_T^0 = S_0 e^{(\mu - \frac{\sigma^2}{2})T}$.

Percentile Prices:

Upper Percentile ($+\alpha$): $S_T^{\alpha} = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)T + \sigma N(\alpha)^{-1}\sqrt{T}}$

Lower Percentile ($-\alpha$): $S_T^{-\alpha} = S_0 e^{\left(\mu - \frac{\sigma^2}{2}\right)T + \sigma N(-\alpha)^{-1}\sqrt{T}}$

Represent the confidence intervals for the stock price at time T. They define the range within which the price is expected to fall given a specific probability level α.

2. Return Statistics (Expected Returns): The order of operations (Expectation vs. Logarithm) significantly changes the result:

- **Expected Accumulated Return = Mean**: $\boxed{\dfrac{1}{T}\ln[E(S_T/S_0)] = \mu}$.

  Represents the average value return (Mean).

- **Expected Return (Avg. of Paths) = Median**: $\boxed{\dfrac{1}{T}E[\ln(S_T/S_0)] = \mu - \dfrac{\sigma^2}{2}}$.

  This represents the average path return (Median).

⋆ **Tail Expectation (GBM).** The objective is to find the probability that the stock price exceeds a certain strike K at time T, denoted as $P(S_T > K)$.

1. Start with the log-process: $\ln(S_T) = \ln(S_0) + (\mu - \frac{\sigma^2}{2})T + \sigma z_T$.

2. Set the inequality: $\ln(S_0) + (\mu - \frac{\sigma^2}{2})T + \sigma z_T > \ln(K)$.

3. Isolate the standard normal variable ϵ (where $z_T = \epsilon\sqrt{T}$):

$$\epsilon > \frac{\ln(\frac{K}{S}) - (\mu - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}} = \alpha$$

4. Solve using the CDF $N(x)$: $P(S_T > K) = P(\epsilon > \alpha) = 1 - P(\epsilon < \alpha) = N(-\alpha)$

Tail Expectation: $\boxed{P(S_T > K) = N\left(\dfrac{\ln(\frac{S}{K}) + (\mu - \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}\right)}$ = **Probability of Finishing In-the-Money for Call**

Represents the likelihood that a cash-or-nothing call option ends up "in the money" at expiration. Core component of the Black-Scholes formula.

3.4 Black-Scholes Framework

⋆ **Black-Scholes Assumptions** The Black-Scholes model relies on specific framework and model assumptions to ensure the derivation of a unique option price.

Framework Assumptions: - **Efficient Markets**: Markets are arbitrage-free and liquid. - **Frictionless Trading**: No transaction costs, no taxes, and short selling is unrestricted.
- **Continuous Trading**: Assets can be traded at any time.
- **Unlimited Borrowing**: Risk-free borrowing or lending is unlimited.

Model Assumptions:
-**Interest Rate**: The risk-free rate is constant. - **Volatility**: Volatility is constant.
- **Dividends**: No dividends (initially, though derivation includes continuous yield q). - **Returns**: Stock price returns follow a Log-Normal distribution.

⋆ **Black-Scholes Derivation** The goal is to construct a **riskless portfolio** to eliminate stochastic uncertainty.

1. Definitions and Processes: Let S be the stock price and $f(S,t)$ be an option price dependent on S and time t.

- **Stock Process**: $dS = (\mu - q)Sdt + \sigma Sdz$
- **Option Process (by Itô's Lemma)**:

$$df = \left(\frac{\partial f}{\partial t} + \frac{\partial f}{\partial S}(\mu - q)S + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2}\right) dt + \frac{\partial f}{\partial S}\sigma S dz$$

2. Portfolio Construction (Π): Construct a portfolio consisting of one short derivative and Δ units of the underlying asset.

$$\Pi = -f + \Delta S$$

The value of Δ represents the **sensitivity to the underlying asset** ($\frac{\partial f}{\partial S}$) and the number of shares.

3. Portfolio Change ($d\Pi$): The change in portfolio value over time dt, including dividends received on the Δ shares ($\Delta qSdt$):

$$d\Pi = -df + \Delta dS + \Delta qSdt$$

Substituting df and dS into the equation:

$$d\Pi = -\left(\frac{\partial f}{\partial t} + (\mu - q)S\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2}\right) dt - \sigma S \frac{\partial f}{\partial S} dz + \Delta(\mu - q)Sdt + \Delta\sigma Sdz + \Delta qSdt =$$

$$= -\left(\frac{\partial f}{\partial t} - (\mu - q)S\frac{\partial f}{\partial S} + \Delta(\mu - q)S + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} - \Delta qS\right) dt + \left(\Delta - \frac{\partial f}{\partial S}\right)\sigma S\, dz$$

4. Eliminating Risk (Delta Hedging): To make the portfolio riskless, we set the **stochastic loadings** (terms with dz) to zero:

$$\left(\Delta - \frac{\partial f}{\partial S}\right)\sigma Sdz = 0 \implies \boxed{\Delta = \frac{\partial f}{\partial S}}$$

Substituting Δ back into $d\Pi$ eliminates the dz term, leaving **non stochastic term**: (0 risk protfolio):

$$d\Pi = -\left(\frac{\partial f}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} - qS\frac{\partial f}{\partial S}\right) dt$$

5. Final Partial Differential Equation (PDE): Because the portfolio is riskless, it must earn the risk-free rate ($d\Pi = r\Pi dt$).

$$\left(\frac{\partial f}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2} - qS\frac{\partial f}{\partial S}\right) dt = r(f - \Delta S)dt$$

Substitute $\Delta = \frac{\partial f}{\partial S}$ and rearrange to find the **Black-Scholes PDE**:

$$\boxed{\underbrace{\frac{\partial f}{\partial t}}_{\text{Theta}} + \underbrace{(r-q)S\frac{\partial f}{\partial S}}_{\text{Adjusted Drift}} + \underbrace{\frac{1}{2}\sigma^2 S^2 \frac{\partial^2 f}{\partial S^2}}_{\text{Gamma}} = \underbrace{rf}_{\text{Risk-free Return}}}$$

Conceptual Note: The Black-Scholes PDE The Black-Scholes PDE is the fundamental identity that every derivative price must satisfy in an arbitrage-free market. It represents a "break-even" condition between the time-decay of the option and its risk-adjusted exposure to the stock.

- **Theta** ($\frac{\partial f}{\partial t}$): Represents **Time Decay**. Since options have a fixed expiration, they lose value as time passes, all else being equal.

- **Adjusted Drift** ($(r-q)S\Delta$): This is the expected return of the stock in a **Risk-Neutral World**. The real-world drift μ is replaced by the risk-free rate r minus the dividend yield q.

- **Gamma** ($\frac{1}{2}\sigma^2 S^2 \Gamma$): This is the **Convexity Adjustment**. It captures the value gained from the curvature of the option's payoff as the stock price fluctuates.

- **Risk-free Return** (rf): This represents the opportunity cost of capital. It is the return you would receive if the capital tied up in the option was invested in a risk-free bond instead.
The Fundamental Logic: The PDE dictates that the total change in the option value (Theta + Drift + Gamma) must exactly equal the risk-free interest on the option's price. If the left side was greater than rf, you would buy the option; if it was less, you would sell it. In an

$\star$ **Boundary Conditions and Black-Scholes Formula** While the Black-Scholes PDE is generic to all derivatives, specific contract values are determined by their unique **Boundary Conditions** at maturity T. Solving the PDE with these boundaries yields the analytical pricing formulas for European options.

1. Terminal Boundary Conditions (T) These equations define the payoff of the instrument at expiration:

- **Call Option**: $c(S_T, 0) = \max\{S_T - K, 0\}$
- **Put Option**: $p(S_T, 0) = \max\{K - S_T, 0\}$
- **Forward Contract**: $f(S_T, 0) = S_T - K$

2. Black-Scholes Pricing Equations The solution to the PDE for European calls (c) and puts (p) is:

$$\boxed{c = S_0 e^{-qT} N(d_1) - K e^{-rT} N(d_2)}$$

$$\boxed{p = K e^{-rT} N(-d_2) - S_0 e^{-qT} N(-d_1)}$$

3. Standardized Distances (d_1 and d_2) These terms represent the risk-adjusted distance the stock price is from the strike in log-space:

$$d_1 = \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} \qquad d_2 = \frac{\ln(S_0/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}$$

- **The Probability of Exercise**: Note the term $N(d_2)$. This represents the **risk-neutral probability** that the call option will finish in-the-money ($S_T > K$).
- **The Drift Adjustment**: The only difference between d_1 and d_2 is the sign of the $\sigma^2/2$ term. This accounts for the **volatility drag** or the difference between the geometric mean and arithmetic mean return of the stock path.
- **Replication Logic**: The call formula essentially says: "To replicate this option, buy $e^{-qT}N(d_1)$ shares of stock and finance it by borrowing $Ke^{-rT}N(d_2)$ at the risk-free rate."

⋆ **Properties of Black-Scholes Formulas** The Black-Scholes formulas for calls (c) and puts (p) exhibit specific asymptotic behaviors as the underlying stock price (S_0) and time to maturity (T) reach their limits.

1. Limits as Stock Price (S_0) Changes:

- **As S_0 becomes very large**: The call price c tends to its lower bound $S_0e^{-qT} - Ke^{-rT}$, while the put price p tends to zero.
  When S_0 is extremely high, a call option behaves like a forward contract because exercise is virtually certain.
- **As S_0 becomes very small**: The call price c tends to zero, while the put price p tends to $Ke^{-rT} - S_0e^{-qT}$.

2. Limits as Time (T) Changes:

- **As T becomes large ($r > 0$)**: If the dividend yield $q = 0$, the call price c tends to the current stock price S_0; otherwise, it tends to zero. In both cases, the put price p tends to zero.
- **As T approaches 0**: Both c and p approach their respective **intrinsic values** as the time-value component vanishes.

⋆ **Put-Call Parity Revisited** Put-Call Parity is a **distribution-free** relationship that must hold between the prices of European call and put options with the same strike (K) and maturity (T). It ensures no-arbitrage between the options and the underlying asset.

Identity: Starting with the Black-Scholes analytical formulas for both c and p:

$$c - p = \left[S_0e^{-qT}N(d_1) - Ke^{-rT}N(d_2)\right] - \left[Ke^{-rT}N(-d_2) - S_0e^{-qT}N(-d_1)\right]$$

Simplification: Using the symmetry property of the normal distribution where $N(-x) = 1 - N(x)$:

1. **Group S_0 terms**: $S_0e^{-qT}(N(d_1) + N(-d_1)) = S_0e^{-qT}(N(d_1) + 1 - N(d_1)) = S_0e^{-qT}$
2. **Group K terms**: $-Ke^{-rT}(N(d_2) + N(-d_2)) = -Ke^{-rT}(N(d_2) + 1 - N(d_2)) = -Ke^{-rT}$

Put-Call Parity Formula: $\boxed{c - p = S_0e^{-qT} - Ke^{-rT}}$

This derivation proves that even though the Black-Scholes model assumes log-normal returns, the relationship between calls and puts is fundamentally dictated by the **Law of One Price**.

Note: $(S_0e^{-qT} - Ke^{-rT})$ represents the value of a **forward contract** $\rightarrow$ long call combined with a short put at the same strike is economically identical to being long a forward contract.

⋆ **Black's model** is a variation of the Black-Scholes formula used specifically for pricing options on **forward or futures contracts** rather than spot assets. It treats the forward price (F) as the underlying variable.

Stochastic Process for Forward Price (F) Under the risk-neutral measure, the forward price follows a Geometric Brownian Motion with zero drift, as the current forward price is a martingale:

$$dF = \sigma F dz$$

⋆ **Black's Pricing Formulas** The formulas for European call (c) and put (p) options on a forward contract are:

$$\boxed{c = e^{-rT}[F_0 N(d_1) - K N(d_2)]}$$

$$\boxed{p = e^{-rT}[K N(-d_2) - F_0 N(-d_1)]}$$

Definitions of d_1 and d_2 These terms are simplified versions of the standard Black-Scholes variables, as the cost-of-carry is already embedded in the forward price:

$$\boxed{d_1 = \frac{\ln(F_0/K) + (\sigma^2/2)T}{\sigma\sqrt{T}} \qquad d_2 = \frac{\ln(F_0/K) - (\sigma^2/2)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}}$$

The beauty of Black's Model lies in its **neutrality towards the cost-of-carry** $(r-q)$. By using the forward price $F_0 = S_0 e^{(r-q)T}$, all the complexity of interest rates and dividend yields is "collapsed" into the single observed forward price.

- **Discounting**: Note that the entire formula is discounted by e^{-rT}. This is because no upfront payment is required to enter a forward contract; the only "money" involved is the eventual payoff at T, which must be pulled back to the present.
- **Application**: This is the standard model for pricing **Swaptions**, **Bond Options**, and **Commodity Options**.

⋆ **Currency Options B.S.** Currency options are derivatives where the underlying asset is a foreign currency. These are priced using a variation of the Black-Scholes model where the foreign interest rate acts as a continuous dividend yield.
- S_0: The current spot exchange rate (Domestic price of one unit of foreign currency).
- r_d: The domestic risk-free interest rate.
- r_f: The foreign risk-free interest rate.
- σ: The volatility of the exchange rate.

Pricing Formulas (Garman-Kohlhagen): The European call (c) and put (p) prices are defined by substituting r_f for the dividend yield q in the standard model:

$$\boxed{c = S_0 e^{-r_f T} N(d_1) - K e^{-r_d T} N(d_2)} \qquad \boxed{p = K e^{-r_d T} N(-d_2) - S_0 e^{-r_f T} N(-d_1)}$$

Definitions of d_1 and d_2: $\boxed{d_1 = \frac{\ln(S_0/K) + (r_d - r_f + \sigma^2/2)T}{\sigma\sqrt{T}} \qquad d_2 = d_1 - \sigma\sqrt{T}}$

Logic behind this model is rooted in **Interest Rate Parity**.

Foreign Rate as Dividend: Holding a foreign currency is like holding a stock that pays a dividend. The "dividend" is the interest r_f you earn by investing that currency in its home country.

The Opportunity Cost: The formula balances the benefit of earning the foreign rate ($e^{-r_f T}$) against the cost of borrowing domestic funds to finance the position ($e^{-r_d T}$).

Forward Relationship: If you substitute the forward exchange rate $F_0 = S_0 e^{(r_d - r_f)T}$ into these equations, you will find they simplify directly to Black's Model.

3.5 Risk-Neutral Valuation

⋆ **Feynman-Kac Theorem.** Bridges the gap between physics and finance by proving that the solution to a deterministic **partial differential equation** (like Black-Scholes) is mathematically equivalent to the weighted average of all possible future **pathways** of a random process.

This theorem allows us to solve complex derivative pricing problems by either solving a differential equation or by simulating random paths of the underlying asset and averaging the outcomes.

Risk-Neutral Stochastic Process. The asset price S follows a stochastic process (GBM) where the drift is replaced by the **risk-free rate** r minus the **dividend yield** q (asset grows at rate: r):

$$\boxed{dS = (r - q)Sdt + \sigma S dz}$$

⋆ **Risk-Neutral Valuation Equation.** BS PDE solution given as expectation of the payoff in a **RN ($\mathbb{Q}$)world**. The value of a derivative $V(S_T, T)$ with payoff function $f(S_T)$ is the discounted expected value:

$$V(S_T, T) = e^{-rT}\mathbb{E}^{\mathbb{Q}}[f(S_T)] \approx e^{-rT}\frac{1}{N}\sum_{n=1}^{N}[f(S_T^n)]$$

(Dataset of N scenarios: $\{S_T^1, S_T^2, \ldots, S_T^N\}$)
The idea is to do **N scenarios** and take the average of them. AKA price is simply the average payoff across many simulated paths, discounted back to the present. This forms the basis for Monte Carlo simulation in finance.
Expectation calculation: can be performed via
Analytical (Numerical Integration): Useful when the probability density function of S_T is known
Simulation (Monte Carlo): "Do N scenarios and take average" to approximate the true $\mathbb{E}^{\mathbb{Q}}$.

⋆ **Analytical Black-Scholes Derivation.** To find the price of a **Call Option**, we evaluate the risk-neutral expectation analytically using the integral of the payoff $\max(S_T - K, 0)$ against the log-normal density $f(S)$.

$$C = e^{-rT}\mathbb{E}[\max(S_T - K, 0)] = e^{-rT}\int_K^\infty \underbrace{(S_T - K)}_{\text{payoff at maturity}} f(S)dS = e^{-rT}\left[\int_K^\infty S_T f(S)dS - \int_K^\infty K f(S)dS\right] \rightarrow$$

$$\rightarrow \boxed{C = e^{-rT}[I_1 - I_2]}$$

The integral is split into two parts: I_1 represents the asset-or-nothing component (the expected value of the stock given it finishes in-the-money), and I_2 represents the cash-or-nothing component (the cost of the strike price paid at exercise).

Note: While we use the closed-form formula for speed, the underlying logic is identical to a Monte Carlo simulation or numerical quadrature; we are simply **summing up all discounted risk-neutral outcomes weighted by their standard normal probabilities.**

Log-Normal Dynamics. In the **RN ($\mathbb{Q}$) world**, the log-return follows a normal distribution with a specific drift adjustment:

Risk-Neutral (S): $\ln(S_T) = \left(r - \frac{\sigma^2}{2}\right)dt + \sigma dz$

Mean of ln(S): $m = \ln[E(S_T)] - \frac{\sigma^2 T}{2} = \ln(S_0) + (r - \frac{\sigma^2}{2})T$

Std. Dev of ln(S): $w = \sigma\sqrt{T}$

Standardized variable: $Q = \frac{\ln(S_T) - m}{w}$ (z-type)

Standard Normal Density: $h(Q) = \frac{1}{\sqrt{2\pi}}e^{-\frac{Q^2}{2}}$ **PDF** for a variable $Q \sim \text{Normal}(0, 1)$

$h(Q)$ as the "probability weight" assigned to each standardized scenario of stock movement. It ensures that outcomes near the mean are more influential than extreme outliers when determining the option's fair value.

⋆ **The I_2 Integral (Exercise Probability).** I_2 calculates the expected payment of the strike price K at maturity, weighted by the **probability of exercise**.

$$I_2 = K\int_{\frac{\ln(K)-m}{w}}^\infty h(Q)dQ = KN(d_2) \qquad d_2 = \frac{\ln(S_0/K) + (r - \sigma^2/2)T}{\sigma\sqrt{T}}$$

$N(d_2)$ is the risk-neutral probability that the option will expire in-the-money. Multiplying this by the strike price tells us the expected cost we will incur if we choose to exercise.

⋆ **The I_1 Integral (Expected Asset Value).** I_1 calculates the expected value of the stock, conditional on the stock price being **above the strike K**.

$$I_1 = \int_{\frac{\ln(K)-m}{w}}^\infty (e^{Qw+m})h(Q)dQ = S_0 e^{rT}N(d_1) \qquad d_1 = \frac{\ln(S_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}} = d_2 + \sigma\sqrt{T}$$

$N(d_1)$ is the factor by which the current stock price is multiplied to find the expected value of the asset received at maturity, given it is above K. It accounts for both the probability of exercise and the higher-value outcomes of the stock.

⋆ **Final Black-Scholes Formula.** Combining the discounted components I_1 and I_2 yields the standard pricing equation:

$$\boxed{C = e^{-rT}[I_1 - I_2] = S_0N(d_1) - Ke^{-rT}N(d_2)}$$

The formula essentially says: the value of a call is (Expected value of stock received) minus (Present value of strike price paid), where both are adjusted for the likelihood of the option actually being exercised.

⋆ **Simulation Risk-Neutral Valuation.** A numerical framework that bypasses complex PDEs by modeling asset price paths directly in a **RN world**. The process follows four steps:

1. Assume the expected growth rate of the stock is the risk-free rate r.

2. Generate multiple random price paths to maturity.

3. Calculate the option payoff for each individual trial and discount at the risk-free rate.

4. Average all payoffs and discount at the risk-free rate to find the present value.

This method is a "brute force" application of the Law of Large Numbers; as the number of trials increases, the average payoff converges to the true theoretical value of the option (expectation).

⋆ **Monte-Carlo Simulation** The standard tool for simulating **GBM**. It uses a discrete-time (**time step:** $\Delta t = T/M$, where M is the number of steps in the path) version of the stock price evolution process to project future price paths.

Real-World Process ($\mathbb{P}$-measure): $\boxed{S_t = S_{t-1}e^{\left(\mu - q - \frac{\sigma^2}{2}\right)\Delta t + \epsilon_t \sigma\sqrt{\Delta t}}}$

μ is used for forecasting actual future price levels or calculating Value-at-Risk (VaR), it is never used for derivative pricing because the risk-neutral measure $\mathbb{Q}$ already accounts for risk through the arbitrage-free framework. Here, **μ represents the actual expected return**, which includes a **risk premium to compensate** investors for market danger.

Risk-Neutral Process ($\mathbb{Q}$-measure): $\boxed{S_t = S_{t-1}e^{\left(r - q - \frac{\sigma^2}{2}\right)\Delta t + \epsilon_t \sigma\sqrt{\Delta t}}}$

Used **exclusively for derivative pricing**. By replacing μ with the risk-free rate r, we create a "simplified" world where investors don't demand risk premiums, making arbitrage-free valuation possible.

Transition from $\mathbb{P}$ to $\mathbb{Q}$ (Girsanov's Theorem) *effectively shifts the distribution's mean to the risk-free rate without changing the variance, allowing us to discount all expected payoffs at r.*
Both formulas share the same volatility and randomness, the "drift" (average growth) differs based on whether we are pricing a derivative or forecasting actual market prices.

⋆ **General Framework for Application.** To apply Monte Carlo to any derivative payoff $f(S)$:

$$\boxed{\text{Value} = e^{-rT}\left(\frac{1}{N}\sum_{i=1}^{N}\text{Payoff}_i\right) = e^{-rT}\left(\frac{1}{N}\sum_{i=1}^{N}f(S)_i\right)}$$

Dataset of N simulated payoffs: $\{f(S)_1, f(S)_2, \ldots, f(S)_N\}$
Call Option Payoff: $f(S)_i = \max(\text{Avg}(S_t) - K, 0)$

Put Option Payoff: $f(S)_i = \max(K - \text{Avg}(S_t), 0)$
Vector Representation: Each path i is treated as a **vector** of prices $\mathbf{S}_i = [S_{t_0}, S_{t_1}, \ldots, S_{t_M}]$.

The major strength of Monte Carlo is its flexibility; you can change the payoff function to anything (even complex path-dependent variables) without changing the underlying simulation engine.

⋆ **Parameter Calibration for μ.** The real-world drift μ is typically estimated using the **historical average** of the asset's log returns over a specific look-back period.
Daily Log Returns $u_i = \ln(S_i/S_{i-1})$

Daily Mean: $\bar{u} = \frac{1}{n}\sum u_i$

Annualization: Annual Drift $= \mu - \frac{\sigma^2}{2} = \frac{\bar{u}}{\Delta t}$ (where $\Delta t = 1/252$ for daily data).

Average log return provides an estimate for the $(\mu - \sigma^2/2)$ term directly, which is the actual "compound" growth rate observed in historical data.

⋆ **Multiple Asset Monte-Carlo Simulation.** Extends the univariate GBM process to price derivatives that depend on a **basket** of k underlying assets (e.g., basket options, spread options). Each asset k evolves according to its own risk-neutral drift and volatility while maintaining a specific **correlation structure** with other assets.

$$\boxed{S_t^k = S_{t-1}^k e^{\left(r-q_k-\frac{\sigma_k^2}{2}\right)\Delta t+\epsilon_t^k\sigma_k\sqrt{\Delta t}}}$$

(Random term for asset k: $\epsilon_t^k \sim N(0,1)$)

While each stock follows its own random path, they are not independent; the simulation must capture how these assets tend to move together or apart to accurately price multi-asset risk.

1. Generate Random Terms ϵ_t^k. To maintain the relationship between assets, the random shocks ϵ_t^k must satisfy specific statistical properties:

Zero Mean: $\boxed{\mathbb{E}[\epsilon_t^k] = 0}$ for all assets k and **Correlation Structure:** $\boxed{corr(\epsilon_t^k, \epsilon_t^j) = \rho_{k,j}}$.
Even though the terms are random, they are not independent. The correlation coefficient ρ ensures that if Asset A and Asset B historically move together, their simulated random shocks will reflect that same co-movement in every time step.

Correlated Random Variables. Must transform **uncorrelated standard normal variables u into correlated variables $\boldsymbol{\epsilon}$** using a **Cholesky decomposition** of the correlation matrix ρ.

Bivariate Case ($k = 2$): $\boxed{\epsilon_1 = u_1, \quad \epsilon_2 = \rho u_1 + \sqrt{1-\rho^2}u_2}$

General Case:
$$\underbrace{\begin{bmatrix} \epsilon_{1,1} & \dots & \epsilon_{1,K} \\ \vdots & \ddots & \vdots \\ \epsilon_{N,1} & \dots & \epsilon_{N,K} \end{bmatrix}}_{(N\times K)} = \underbrace{\begin{bmatrix} u_{1,1} & \dots & u_{1,K} \\ \vdots & \ddots & \vdots \\ u_{N,1} & \dots & u_{N,K} \end{bmatrix}}_{(N\times K)} \text{Chol}\left(\underbrace{\begin{bmatrix} \rho_{1,1} & \dots & \rho_{1,K} \\ \vdots & \ddots & \vdots \\ \rho_{K,1} & \dots & \rho_{K,K} \end{bmatrix}}_{(K\times K)}\right)$$

The Cholesky matrix acts as a mathematical "filter" that injects the required level of dependency into independent noise. This ensures that if two stocks are 80% correlated, their simulated paths reflect that same 80% relationship step-by-step.

2. Asset Price Path (Multiple Assets). Simulating the path requires updating the entire **vector of stock prices S_t^k** at each time step t to account for joint movements.

- **Simultaneous Evolution:** All S_t^k calculated using same time step Δt but distinct, correlated ϵ_t^k

- **Markov Property:** The price vector at t is dependent only on the price vector at $t-1$, meaning returns have no memory.

Single Simulation of Price Evolution Matrix S. For a **basket of K stocks over N time steps**, the engine produces an $(N \times K)$ **matrix of prices for every single trial run**, capturing the full evolution of the portfolio.

$$\mathbf{S} = \begin{bmatrix} S_{0,1} & S_{0,2} & \dots & S_{0,K} \\ S_{1,1} & S_{1,2} & \dots & S_{1,K} \\ \vdots & \vdots & \ddots & \vdots \\ S_{N,1} & S_{N,2} & \dots & S_{N,K} \end{bmatrix} \quad \textit{with} \quad S_{t,k} = S_{t-1,k} \cdot \exp\left((r - q_k - \frac{\sigma_k^2}{2})\Delta t + \epsilon_{t,k}\sigma_k\sqrt{\Delta t}\right)$$

The transition from row $t-1$ to row t is calculated by applying the GBM evolution formula to every element in the vector simultaneously:

Deconstructing S_T in the Basket Formula. The notation $\mathbf{S_T}$ represents the state of the entire market at maturity, let $\mathbf{S_T^m}$ represents the specific price of one asset within that market.

Vector $\mathbf{S_T}$: This is the collection of all terminal prices: $\mathbf{S_T} = [S_T^1, S_T^2, \ldots, S_T^M]$.

Individual S_T^m: This is the price of the m^{th} stock in the basket at time T. Each one was evolved using its own volatility (σ_m) and dividend yield (q_m).

Average ($\frac{1}{M}\sum S_T^m$): This operation "collapses" the M different stock prices into a single numerical value representing the basket's performance.

Logic: Think of it like a relay team. $S_T^1, S_T^2, \ldots$ are the individual runners (stocks). The basket payoff doesn't care who ran fastest; it only cares about the average time of the whole team compared to the target K.

3. Multivariate Payoffs $f(\mathbf{S})$. The payoff of a multi-asset derivative is calculated by applying a function to the vector of terminal prices (or averages) of all included stocks.

Basket Call: $$f(\mathbf{S_T}) = \max\left\{\frac{1}{M}\sum_{m=1}^{M} S_T^m - K,\ 0\right\}$$

Basket Put: $$f(\mathbf{S_T}) = \max\left\{K - \frac{1}{M}\sum_{m=1}^{M} S_T^m,\ 0\right\}$$

4. The Simulation Average, Final Value Once all N_{trials} have been completed, we aggregate the individual results to find the expected value of the option today by Averaging and Discounting

$$\boxed{\text{Price} = e^{-rT}\left(\frac{1}{N_{trials}}\sum_{i=1}^{N_{trials}} f(\mathbf{S_T})_i\right)}$$

Each $f(\mathbf{S_T})_i$ represents the outcome of one "parallel universe." By averaging across thousands of these trials, we converge on the risk-neutral price of the basket contract.

⋆ **Simulation Workflow (General Framework).**

1. **Decompose Correlation:** Calculate the Cholesky matrix L such that $LL^T = \rho$.
2. **Generate Noise:** Sample independent standard normals u and transform them into correlated ϵ.
3. **Evolve Paths:** Project all k asset prices simultaneously using their respective S_t^k dynamics.
3. **Apply Payoff:** Calculate $f(\mathbf{S})$ for each simulation trial.
4. **Average and Discount:** Value $= e^{-rT}\mathbb{E}^{\mathbb{Q}}[f(\mathbf{S})]$.

This framework is the "gold standard" for pricing complex structural products because it can scale to any number of assets and any correlation structure without needing a different mathematical model.

⋆ **Time-Varying Parameters & Antithetic Logic for MC** In reality, parameters like r, σ, and q are not constant (but still deterministic). We incorporate "term structures" by allowing these values to change at every discrete time step Δt.

$$S_t = S_{t-1}\exp\left(\left(r_{t-1} - q_{t-1} - \frac{\sigma_{t-1}^2}{2}\right)\Delta t + \epsilon_t \sigma_{t-1}\sqrt{\Delta t}\right)$$

- **Risk-Free Rate (r_t):** Usually follows the **Forward Curve**.
- **Volatility (σ_t):** Follows the **Volatility Term Structure** (e.g., volatility might be expected to rise during earnings season).
- **Dividend Yield (q_t):** Changes based on expected future dividend schedules.

Variance Reduction: Antithetic Variables Monte Carlo is computationally expensive. To speed up "convergence" (getting to the right price faster), we use **Antithetic Variates (Mirroring Random paths)**.

- For every random path generated using shocks $+\epsilon_t$, we *simultaneously* create a "mirror path" using shocks $-\epsilon_t$.
- **The Logic:** If the first path goes "up" due to luck, the mirror path goes "down." When we average them, we cancel out a large portion of the random noise (variance), making the final price much more stable with fewer trials.

By combining time-varying parameters with antithetic variables, the simulation becomes both more realistic (matching the market curves) and more efficient (reaching a stable price with half the random draws).

⋆ **Numerical Integration**

1. The Risk-Neutral Integral Option pricing can be viewed as an integral of the discounted payoff over all possible future states. In the **Risk-Neutral World**, the price C is defined as:

$$\boxed{C = e^{-rT}\int_{-\infty}^{\infty} \max\{S-K,0\}\,dS = e^{-rT}\int_{-\infty}^{\infty} \max\{S(z)-K,0\}\,dz}$$

- z: The standard normal variable driving the uncertainty.
- $S(z)$: The stock price at maturity, which is a function of the normal shock z.

2. Quadrature Approximation When we cannot solve the integral analytically, we use **Numerical Quadrature**. We approximate the continuous integral as a weighted sum of the payoff at specific discrete points:

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} g(z) = \sum_{k=0}^{M} w_k g(F_k)$$

- $g(z)$: The option payoff function (for call: $\max\{S(z)-K,0\}$)

- F_k are the **Points of the Normal Density Function** chosen to evaluate the payoff. Each node F_k is used to calculate a specific terminal stock price: $S(F_k) = S_0 exp((r - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}F_k)$

- w_k w_k represents the **Importance** or relative probability density of each node $w_k = \frac{2^{M-1}M!\sqrt{\pi}}{M^2[H_{M-1}(F_k)]^2}$ Nodes near the center ($F_k \approx 0$) carry the highest weights because those outcomes are most likely.

- M: The number of approximating points used to reconstruct the continuous probability distribution. We **choose M** (higher, more accurate) and look up **given F_k and w_k** for the corresponding M

3. Application: The Call Option For a standard call, we evaluate the stock price at each node F_k using the GBM solution:

$$S(F_k) = S_0 e^{(r-\frac{1}{2}\sigma^2)T+\sigma\sqrt{T}F_k}$$

$$C \approx e^{-rT}\sum_{k=0}^{M} w_k \max\{S(F_k) - K, 0\}$$

⋆ **Generalized Numerical Integration Framework**

The Fundamental Integral In a risk-neutral world, the price of an option V is the discounted expected payoff. For a European option, this is represented as an integral over all possible future asset prices S_T:

$$V = e^{-rT}\int_0^\infty \text{Payoff}(S_T)\cdot\phi(S_T)\,dS_T$$

- **Payoff(S_T):** The contract function (e.g., $\max(S_T - K, 0)$).
- $\phi(S_T)$**:** The probability density function (PDF) of the asset at time T.

Numerical Approximation (Quadrature) Since the PDF often lacks a simple antiderivative, we approximate the integral by a finite weighted sum of the integrand at specific points (*nodes*).

$$\int_a^b f(x)\,dx \approx \sum_{i=1}^{n} w_i f(x_i)$$

- **Nodes (x_i):** The specific stock prices where we evaluate the payoff.
- **Weights (w_i):** The "importance" or probability weight assigned to each node.

Gaussian Quadrature Logic Instead of using equally spaced points (like the Trapezoid rule), Gaussian Quadrature chooses the nodes and weights optimally.

⋆ **Least Squares Monte Carlo (LSMC)** A simulation-based approach (Longstaff and Schwartz, 2001) used to value American options. It solves the optimal early exercise problem by using cross-sectional regression to estimate the expected **continuation value** at each potential exercise date.

Standard Monte Carlo moves forward in time, making it difficult to know the future expected value needed for early exercise decisions. LSMC works backward, using regression across simulated paths to approximate the continuation value, similar to comparing exercise vs. continuation at each node in a binomial tree.

Algorithm Framework:

1. **Simulate Paths:** Generate N forward risk-neutral paths for the underlying asset price S_t.
2. **Terminal Payoff:** At maturity T, determine the standard option payoff for all paths.

$$CF_T = \max(K - S_T, 0) \quad \text{(for a put)}$$

3. **Backward Induction:** For each time step t moving backwards from $T-1$ down to 1:
   - **Identify in-the-money (ITM) paths:** Only consider paths where immediate exercise yields a positive payoff (e.g., $S_t < K$).
   - **Discount future cash flows:** For these ITM paths, take the realized cash flow from the subsequent optimal stopping time (CF_{t+1}) and discount it back to t.
   - **Estimate continuation value:** Regress the discounted future cash flows onto a set of **basis functions** of the current state S_t (e.g., a polynomial basis).

$$V_{cont}(S_t) = \mathbb{E}[e^{-r\Delta t}CF_{t+1}|S_t] \approx \hat{a} + \hat{b}S_t + \hat{c}S_t^2$$

- **Exercise Decision:** Compare the immediate exercise value to the estimated continuation value.

$$\text{Exercise iff: Immediate Payoff} > V_{cont}(S_t)$$

- **Update Cash Flows:** If exercised early, update the path's cash flow at t to the immediate payoff and set all subsequent future cash flows for that path to 0.

4. Final Valuation: At $t = 0$, discount the final optimal cash flows from each path back to present value and average them.

$$\text{American Option Value} = \max\left(\text{Immediate Payoff at } t = 0, \frac{1}{N}\sum_{i=1}^{N} \text{Discounted } CF_i\right)$$

Because the continuation value is an expected value, regression effectively smooths the outcomes across similar in-the-money paths to estimate what the option is worth if kept alive. The framework is highly flexible and can accommodate arbitrary payoff functions and multiple economic factors simply by adding variables to the regression basis.

⋆ **The LSMC Regression.** At each time step t (working backwards), we perform a cross-sectional regression strictly on the **in-the-money (ITM) paths**.

- **Dependent Variable (Y):** The discounted realized cash flow from the *next* optimal stopping point for that path ($e^{-r\Delta t}CF_{t+1}$).
- **Independent Variable (X):** The current state variables of the path, typically expressed through basis functions (e.g., S_t, S_t^2).

The regression mathematically estimates the ***continuation value****, representing the conditional expectation* $\mathbb{E}[e^{-r\Delta t}CF_{t+1}|S_t]$. *We only run the regression on ITM paths because out-of-the-money options will not be exercised anyway; including them would skew the curve and ruin the estimate precisely in the region where the early exercise decision actually matters.*

LSMC Quadratic Regression Model At a specific early exercise date t, we run the following cross-sectional Ordinary Least Squares (OLS) regression strictly across all **in-the-money** paths to estimate the **continuation value**:

$$Y_i = a + bS_{i,t} + cS_{i,t}^2 + \epsilon_i$$

Variables:

- i: Index representing a specific **in-the-money** simulated path at time t.
- Y_i: The **dependent variable**. The realized, discounted future cash flow for path i from the next optimal stopping time ($Y_i = e^{-r\Delta t}CF_{i,t+1}$).
- $S_{i,t}, S_{i,t}^2$: The **independent variables** (basis functions). The underlying asset price (and its square) at time t for path i.
- ϵ_i: The error term representing the noise in the Monte Carlo simulation for that specific path.

Coefficients (a, b, c):

- $\hat{a}, \hat{b}, \hat{c}$: The estimated OLS weights that map the current stock price to the future discounted cash flows.
- **Meaning:** Individually, these coefficients do not represent financial metrics like Delta or Gamma. Together, they define the fitted polynomial curve:

$$\hat{V}_{cont}(S_{i,t}) = \hat{a} + \hat{b}S_{i,t} + \hat{c}S_{i,t}^2$$

  which represents the conditional **expected value** of holding the option instead of exercising it immediately.

⋆ **Fundamental Pricing Equation** A universal pricing framework that links the price of any asset to a chosen numeraire under a specific probability measure.

- v: The asset to be priced (v_0 at time 0 v_T at time T) any asset.
- Q_g: The pricing probability measure associated specifically with the chosen numeraire g.
- g: The **Numeraire asset**, a positive price process that scales all cash flows. Essentially, converting prices to relative (normalized) prices.

For example, when pricing a standard European call option under the risk-neutral measure, the numeraire is typically the risk-free money market account, $g_t = e^{rt}$. or If $g = P(t,T)$ (Zero-Coupon Bond): Then Q_g becomes a different measure, such as the T-forward measure (Q^T).

Fundamental Equation: The equation expresses the core principle of arbitrage-free pricing: the **relative price** of any asset must be a **martingale** under the probability measure uniquely tied to its numeraire.

$$\boxed{\frac{v_0}{g_0} = \mathbb{E}^{Q_g}\left[\frac{v_T}{g_T}\right]}$$

- $\frac{v_0}{g_0}$: The **current relative price**. It represents how many units of the numeraire g it costs to buy exactly one unit of asset v today.
- $\frac{v_T}{g_T}$: The **future relative price**. The payoff or value of asset v at time T, measured strictly in units of g at time T.
- $\mathbb{E}^{Q_g}[\ldots]$: The **expected value** calculated strictly under Q_g, which is the unique, customized probability measure "matched" to the numeraire g.

*Intuitively, think of the numeraire g as a new currency (like converting from Dollars to Gold). If you measure your option's payoff in Gold instead of Dollars, the equation says the option's price today (in Gold) is simply its expected future payoff (in Gold). No separate discount factor is needed because the customized probability measure Q_g mathematically absorbs the discounting and risk adjustments, creating a perfectly balanced "fair game" (martingale) for that specific currency. This equation states that there exists a probability measure Q_g under which the relative price of the asset (v/g) is a **martingale**. By changing the "currency" or reference asset we use to measure prices (the numeraire), we can shift to a probability measure where the asset's current relative price is simply the expected value of its future relative price.*

Example: 1-Period Binomial Model
Stock (S): Today $S_0 = 100$. Tomorrow it's either $S_u = 120$ or $S_d = 90$.
Call Option (v): Strike $K = 100$. Payoffs are $C_u = 20$ and $C_d = 0$.
Numeraire (g): 10% interest rate. Today $g_0 = 1$, tomorrow $g_T = 1.1$.
1. Find the Probability Measure (Q_g): Apply the Fundamental Equation to the stock to find the risk-neutral probability q:

$$\frac{S_0}{g_0} = \mathbb{E}^{Q}\left[\frac{S_T}{g_T}\right] \quad \Rightarrow \quad 100 = \frac{1}{1.1}\Big[q(120) + (1-q)(90)\Big]$$

Solving for q gives $110 = 30q + 90 \Rightarrow q = p = 2/3$.

3. Price the Derivative (v_0): Now apply the exact same equation and measure Q to the Call option to find its price C_0:

$$\frac{C_0}{g_0} = \mathbb{E}^{Q}\left[\frac{C_T}{g_T}\right] \quad \Rightarrow \quad C_0 = \frac{1}{1.1}\left[\frac{2}{3}(20) + \frac{1}{3}(0)\right] = \frac{40}{3.3} \approx 12.12$$

Notice how the Fundamental Equation is used twice: first on the observable underlying asset (the stock) to reverse-engineer the "fake" risk-neutral probabilities ($q = 2/3 = p$), and then on the unpriced derivative (the call option) to calculate its fair present value using those exact same probabilities.

⋆ **Generalized Black-Scholes Framework** The framework revolves around converting absolute prices to relative prices using a numeraire to establish martingales.
Numeraire: A positive price process used to scale all cash flows, effectively converting absolute prices into relative prices. The Black-Scholes model specifically uses the money market account as this numeraire asset.
Martingale: A mathematical property where the expected future value of a function equals its present value: $\mathbb{E}[f(T)] = f(0)$.
Relative Prices: These are martingales when evaluated under the probability measure specific to the chosen numeraire.
Non-Uniqueness: You can choose different numeraires; they will all mathematically reproduce the exact same price, but some are much easier to use depending on the derivative's payoff.

⋆ **Risk-Neutral Numeraire (Black-Scholes)** Applying the fundamental equation specifically to the Black-Scholes money-market setup.
- v: The asset to be priced.
- g: The Risk-Neutral numeraire, which is the money-market account. Today's value is $g_0 = 1$, and its future value assumes continuous compounding: $g_T = e^{\int_0^T r dt} = e^{rT}$.

Applying the Fundamental Pricing Equation: $\boxed{\frac{v_0}{g_0} = \mathbb{E}^{Q_g}\left[\frac{v_T}{g_T}\right] = e^{-rT}\mathbb{E}^{Q_g}[v_T]}$

Since $g_0 = 1$, this simplifies directly to the familiar Feynman-Kac solution: $v_0 = e^{-rT}\mathbb{E}^{Q_g}[v_T]$
This mathematically proves that under the risk-neutral measure, today's option price is simply the expected future payoff discounted at the risk-free rate.

⋆ **Deriving the Risk-Neutral Stock Process** We can use the martingale property of relative prices to find the precise stochastic process for the stock under the risk-neutral measure.

1. Total Return of a Stock (V): Including a continuous dividend yield q, the total return is the capital gains plus dividends: $\boxed{\frac{dV}{V} = \frac{dS}{S} + qdt = (\mu_S + q)dt + \sigma dz}$

2. Define the Relative Price (Z): Let $Z = V/g$ be the relative price of the total stock return with respect to the risk-neutral numeraire g. Because it's a relative price, its drift under the risk-neutral measure (Q_{RN}) must be exactly zero to satisfy the martingale property: $\boxed{\mathbb{E}^{Q_{RN}}\left[\frac{dZ}{Z}\right] = -rdt + (\mu_S + q)dt = 0}$

3. Solve for the Drift (μ_S): By setting the drift to zero, we isolate μ_S: $\boxed{\mu_S = (r - q)}$

4. The Black-Scholes Risk-Neutral Process: Substitute μ_S back into the original stock price dynamic to get the final Black-Scholes risk-neutral stock process: $\boxed{\frac{dS}{S} = (r - q)dt + \sigma dz}$

⋆ **Stock Numeraire** Instead of using the money market account, we can define the stock itself, V, as the numeraire under a new probability measure Q_S.

1. Total Return of the Stock: with a cts. dividend yield q: $\boxed{\frac{dV}{V} = \frac{dS}{S} + qdt = (\mu_S + q)dt + \sigma dz}$

2. Pricing the Money Market Account: Our target asset to price relative to the stock is the money market account R: $\boxed{R = \exp(rT)}$

3. Applying the Martingale Property: Define the relative price $Z = R/V$. For Z to be a **martingale** under the stock numeraire measure Q_S, its expected drift must equal zero. Applying Ito's lemma: $\boxed{\mathbb{E}^{Q_S}\left[\frac{dZ}{Z}\right] = rdt - \mathbb{E}\left[\frac{dV}{V}\right] + \left(\frac{dV}{V}\right)^2 = 0}$

Substitute the terms to solve for zero: $(r - q - \mu_S + \sigma_S^2)dt = 0$

4. The Stock Numeraire Drift (Process): Solving the above equation for the drift μ_S gives:

$$\boxed{\mu_S = (r - q + \sigma_S^2)}$$

Substituting this back into the original stock price dynamic yields the final process under Q_S:

$$\boxed{\frac{dS}{S} = (r - q + \sigma_S^2)dt + \sigma dz}$$

⋆ **Call Price Revisited: Combining Numeraires** The classic Black-Scholes call option formula is fundamentally a combination of the two different probability measures: the **Stock Numeraire** and the **Money Market Numeraire**.

1. The Integral Equation: The call price C written out using both numeraires:

$$C = S_0 \int_K^\infty \left(\frac{S_T}{S_T e^{qT}}\right) f(S^{Q_S}) dS^{Q_S} - \int_K^\infty \frac{K}{e^{rT}} f(S^{Q_{MM}}) dS^{Q_{RN}}$$

- **First Term:** Represents the **Stock Numeraire** evaluation.
- **Second Term:** Represents the **Money Market Numeraire** evaluation.

2. Translating to Probabilities: This complex integral simplifies into the probabilities of the option finishing in-the-money under each respective measure:

$$\boxed{C = S_0 e^{-qT}\text{Prob}^{Q_S}(\ln(S) > \ln(K)) - e^{-rT} K \text{Prob}^{Q_{RN}}(\ln(S) > \ln(K))}$$

3. The Final Black-Scholes Formula: This directly yields the standard formula, where $N(d_1)$ and $N(d_2)$ represent these exact probabilities:

$$C = S_0 e^{-qT} N(d_1) - e^{-rT} K N(d_2)$$

4. Reference Distributions: The underlying distributions for each measure that drive $N(d_1)$ and $N(d_2)$ are:

- **Stock Numeraire (Q_S):** $f(S^{Q_S}) \sim (r - q + \sigma_S^2)dt + \sigma dz$
- **Money Market (Q_{MM}):** $f(S^{Q_{MM}}) \sim (r - q)dt + \sigma dz$

The specific inputs for the cumulative normal distributions are calculated as:

$$d_1 = \frac{\ln(S_0/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} \qquad d_2 = \frac{\ln(S_0/K) + (r - q - \sigma^2/2)T}{\sigma\sqrt{T}}$$

Why Do We Need Them All for the Black-Scholes "Q"? When pricing a complex derivative like a European Call Option, the **payoff has two competing pieces: receiving a risky asset (S_T) and paying safe cash (K)**. Instead of trying to calculate this under one single measure (which creates massive, unworkable integrals), the Numeraire Framework lets us divide and conquer:

- Use the **Stock Measure (Q_S)** to calculate the expected value of the shares we receive. This generates $N(d_1)$.

- Use the **Money Market Measure (Q_{MM})** to calculate the expected cost of the cash we pay. This generates $N(d_2)$.

By combining these two specific probability measures, the Black-Scholes formula cleanly and elegantly solves the pricing problem: $C = S_0 e^{-qT} N(d_1) - e^{-rT} K N(d_2)$

⋆ **Bond Price Numeraire (Q_T)** The Bond Price Numeraire is used to derive the **Black's Formula**, typically used for options on futures, forward contracts, and interest rate products.

1. Bond Price as Numeraire In this measure, we use a zero-coupon bond maturing at time T as our "ruler."

- **Numeraire:** $g = P(t, T)$
- **Current Value:** $P(0, T) = e^{-rT}$
- **Terminal Value:** $P(T, T) = 1$ (A bond is worth exactly its face value at maturity)

2. The Fundamental Martingale Relationship The relative price $Z = V/P(t, T)$ is a martingale under the Q_T measure.

$$Z = \frac{V_0}{P(0,T)} = F_0 = \mathbb{E}^{Q_T}\left[\frac{V_T}{P(T,T)}\right] = \mathbb{E}^{Q_T}[F_T]$$

This implies that the **Forward Price (F)** is a martingale under Q_T:

$$\mathbb{E}^{Q_T}[F_T] = F_0$$

3. Implications for Black's Formula Because the forward price is a martingale, its drift is zero ($\mu_F = 0$):

$$\frac{dF}{F} = \sigma dz$$

- **Significance:** This is the foundation of **Black's Formula**. It allows us to price options on forwards or futures without needing to know the specific path of interest rates, as the interest rate risk is "absorbed" by the bond numeraire.

Logic Check: By switching to the Bond Numeraire, we eliminate the discounting term inside the expectation. Since $P(T,T) = 1$, the option payoff is simply the expectation of the terminal forward price, which is then discounted back to today using the current bond price $P(0,T)$.

⋆ **The Tale of Three Measures: P, Q_{MM}, and Q_S** It is crucial to understand that the **Money Market measure** (Q_{MM}) and the **Risk-Neutral measure** (Q_{RN}) are the exact same thing. We are simply looking at the exact same underlying stock through different "currency lenses."

The Real-World Measure (P)

- **The Baseline:** Reality. Your subjective belief about the stock's trajectory.
- **The Numeraire:** None (or a real-world consumption basket).
- **The Drift:** μ (The actual expected return, including the risk premium).
- **The Process:** $\frac{dS}{S} = \mu dt + \sigma dz$
- **Usage:** Risk management (e.g., Value at Risk) and forecasting.
- **Why it fails for pricing:** Because μ includes a subjective risk premium, different investors will demand different returns, meaning we cannot agree on a single fair price.

The Risk-Neutral / Money Market Measure (Q_{RN} or Q_{MM})

- **The Baseline:** The risk-free bank account (Cash).
- **The Numeraire:** $g_t = e^{rt}$
- **The Drift:** $r - q$
- **The Process:** $\frac{dS}{S} = (r - q)dt + \sigma dz$
- **Usage:** This is our universal pricing tool for the *cash* portion of payoffs (like paying the strike K). It mathematically removes the subjective μ and replaces it with the observable r.

The Stock Measure (Q_S)

- **The Baseline:** The stock itself.
- **The Numeraire:** $V_t = S_t e^{qt}$
- **The Drift:** $r - q + \sigma^2$
- **The Process:** $\frac{dS}{S} = (r - q + \sigma^2)dt + \sigma dz$
- **Usage:** Used when the payoff is directly tied to receiving the stock itself (like the S_T portion of a Call). Measuring in units of volatile stock mathematically shifts the probability distribution to the right by exactly the variance (σ^2).

Other Measures:

The T-Forward Measure (Q_T)

- **Context:** What happens when interest rates are not constant, but stochastic (moving randomly)? If r is volatile, our standard Money Market numeraire ($e^{\int r dt}$) becomes a messy, random variable itself.
- **The Baseline:** A Zero-Coupon Bond maturing at time T.
- **The Numeraire:** $P(t, T)$. The beauty of a zero-coupon bond is that we know exactly what it will be worth at maturity: $P(T, T) = 1$.
- **The Pricing Equation:** $\frac{v_t}{P(t,T)} = \mathbb{E}^{Q_T}\left[\frac{v_T}{P(T,T)}\right] = \mathbb{E}^{Q_T}[v_T]$
- **Why we need it:** By using the zero-coupon bond as our currency, we completely decouple the volatility of the interest rate from the payoff of the derivative. It makes pricing bond options (like caps, floors, and swaptions) mathematically elegant because the discounting factor $P(t, T)$ gets pulled entirely outside the expectation: $v_t = P(t, T)\mathbb{E}^{Q_T}[v_T]$.

The Foreign Measure (Q_f) - For Currency Derivatives

- **Context:** Pricing FX derivatives or "Quanto" derivatives (e.g., an option on a European stock that pays out in US Dollars). You now have two risks: the stock's volatility and the exchange rate's volatility.
- **The Baseline:** The Foreign risk-free bank account, translated into Domestic currency.
- **The Numeraire:** $X_t \cdot e^{r_f t}$, where X_t is the spot exchange rate and r_f is the foreign risk-free

rate.

- **The Drift Adjustment:** Changing from the domestic measure (Q_d) to the foreign measure (Q_f) introduces a "Quanto adjustment." Because the exchange rate and the foreign asset are often correlated, Ito's Lemma forces a covariance penalty into the drift.
- **Math:** The drift of the asset under the domestic measure gets shifted by $-\rho\sigma_S\sigma_X$ (the correlation $\times$ volatility of the stock $\times$ volatility of the FX rate).
- **Why we need it:** If you are a US investor trading a Euro-denominated asset, using the Foreign Measure mathematically absorbs the exchange rate risk. It lets you price the asset as if you were a European investor using local Q_{MM}, and then simply translate the final clean price back to USD using today's exchange rate.

Note: Every time we introduce a new source of randomness (stochastic rates, fluctuating FX), we just define a new Numeraire that perfectly matches and absorbs that specific risk. The math gets cleaner because the Numeraire "eats" the complexity!

3.5.1 For trading Vol

The Log-Measure (Q_{log}): The Volatility Trader's Tool While Q_{MM} and Q_S are great for pricing fixed payoffs, volatility traders care about the **path** and the **variance**. The Log-Measure is the measure associated with the **Log-Numeraire**.

1. The Numeraire: The numeraire is the **Log-Wealth portfolio** (also known as the Kelly Criterion or Growth-Optimal portfolio).

$$g_t = \exp\left(\int_0^t (r + \frac{1}{2}\sigma^2)ds + \int_0^t \sigma dz\right)$$

2. The Connection to Volatility Trading: In volatility trading (like Variance Swaps), we are often interested in the expected log-return of the stock: $\mathbb{E}[\ln(S_T/S_0)]$.

- Under the standard Q_{MM}, the drift of $\ln(S)$ is $(r - q - \frac{1}{2}\sigma^2)$.
- The term $-\frac{1}{2}\sigma^2$ is known as the **volatility drag** or **stratified variance**.

3. Why it's used for Breakout/Vol Strategies:

- **Convexity Adjustment:** Volatility traders profit from the "gamma" of an option. The Log-Measure naturally incorporates the relationship between the stock price path and the variance.
- **The Variance Swap Connection:** To price a Variance Swap, you effectively trade a "log-contract." The Log-Measure allows you to see the fair value of future realized variance directly without needing to hedge every tiny movement in the stock price.
- **Kelly Criterion:** For traders looking for "breakouts" or long-term growth, the Log-Measure represents the measure where the *expected growth rate* of wealth is maximized.

Intuition Check: If you are trading a breakout, you are betting that the realized σ will be higher than the implied σ in the option price. The Log-Measure is the mathematical framework that isolates that σ^2 term from the directional r or μ terms.

⋆ **Volatility Framework: LEAP Straddles & The Log-Measure** Trading LEAPs (Long-term Equity Anticipation Securities) via straddles is a bet on the **Volatility Risk Premium**. The Log-Measure (Q_{log}) provides the mathematical foundation for why this works over long horizons.

1. The Variance "Drag" Equation: In the standard Risk-Neutral world (Q_{MM}), the stock price is a martingale, but its *logarithm* is not. The expected log-return is:

$$\mathbb{E}^{Q_{RN}}[\ln(S_T/S_0)] = (r - q - \frac{1}{2}\sigma^2)T$$

The term $\frac{1}{2}\sigma^2$ is the **variance drag**. As a straddle trader, you are effectively "buying" this drag. You profit if the actual realized σ^2 during the LEAP's life is higher than what was priced into the initial IV.

2. Why use the Log-Measure (Q_{log})? While Q_{MM} is for pricing the entry, Q_{log} is for managing the **Gamma-Theta trade-off** over time.

- **Numeraire:** The Growth-Optimal Portfolio (Kelly Portfolio).
- **Property:** Under Q_{log}, the log-return of the stock is a *true martingale*.
- **Insight:** It isolates the "Vol" component. For a LEAP straddle, your PnL is driven by the **Fundamental Law of Active Management**: the difference between realized variance and implied variance.

3. The "Log-Contract" Replication: A LEAP straddle is a practical approximation of a **Log-Contract**.

- A Log-Contract pays $\ln(S_T/S_0)$.
- Because $\ln(S)$ is concave, you need a portfolio of options (a "strip" of straddles) to replicate it.
- **The Strategy:** By holding the LEAP straddle, you are capturing the *Convexity* of the stock's path. The Log-Measure is the framework that proves that if you delta-hedge perfectly, your PnL is exactly:

$$\text{PnL} = \int_0^T \frac{1}{2}\Gamma S^2(\sigma^2_{realized} - \sigma^2_{implied})dt$$

4. Strategic Application for LEAPs:

- **Time Decay (Theta):** You pay e^{-rT} to "rent" the variance.
- **Gamma Capture:** Because LEAPs have long durations, the Log-Measure helps you identify if the "volatility cluster" (breakout) is sufficient to overcome the high cost of the LEAP premium.
- **Breakout Logic:** In the Log-Measure, a breakout is seen as a "jump" in the quadratic variation of the path.

Intuition Check: Why does this matter for LEAPs? Because over 1-2 years, the "drift" (r) matters less than the accumulated variance (σ^2). The Log-Measure is the only measure that treats the "Variance Drag" as a feature, not a bug, allowing you to price the breakout potential accurately.

$\star$ **The Log-Measure GBM Framework** To value volatility for LEAP Straddles, we simulate the S&P 500 (S_t) under the Q_{log} measure where the Growth-Optimal Portfolio is the numeraire.

1. Implementing the Numeraire: The numeraire L_t is the portfolio that maximizes $E[\ln(L_T)]$. In a Black-Scholes world, this portfolio has a return of:

$$\mu_L = r + \sigma^2$$

By using L_t as the numeraire, the drift of the S&P 500 becomes:

$$\mathbb{E}^{Q_{log}}[dS_t/S_t] = (r + \sigma^2)dt$$

2. The "Physics" of the Path: Under this measure, the log-price $X_t = \ln S_t$ is a **driftless** martingale:

$$dX_t = \sigma dZ_t^{log}$$

The terminal stock price for your simulation is:

$$S_T = S_0 \exp(\sigma\sqrt{T}\epsilon)$$

Note: There is no $(r - \frac{1}{2}\sigma^2)T$ term because the Numeraire has 'absorbed' the drift.

3.6 EXTRA:

3.6.1 Mathematical Derivation of Itô's Lemma.

The objective is to find the differential dG for a function $G(x,t)$, where x is a stochastic state variable (such as a stock price or interest rate) following the Itô process $dx = a(x,t)dt + b(x,t)dz$.

Assumptions and Stochastic Calculus Rules: To solve the expansion, we assume the following rules for infinitesimal terms as $dt \to 0$:

- **Standard rule**: $dt^n = 0$ for all $n > 1$.
- **Stochastic rule**: $dz^2 = dt$.
- **Cross-product rule**: $dz \cdot dt = \epsilon\sqrt{dt} \cdot dt = \epsilon dt^{1.5} = 0$.

Step 1: The Function $G(x,t)$ and Taylor Expansion. We define $G(x,t)$ as a twice-differentiable function of the stochastic variable x and time t. We expand dG using a Taylor series to capture how the function changes over an infinitesimal step:

$$dG(x,t) = \underbrace{\frac{\partial G}{\partial x}dx + \frac{\partial G}{\partial t}dt}_{\text{First-order terms}} + \underbrace{\frac{1}{2}\frac{\partial^2 G}{\partial x^2}dx^2 + \frac{1}{2}\frac{\partial^2 G}{\partial t^2}dt^2 + \frac{\partial^2 G}{\partial x \partial t}dxdt}_{\text{Second-order terms}} + \dots$$

In stochastic calculus, we cannot discard the second-order term $\frac{1}{2}G_{xx}dx^2$ because dx^2 is of order dt.

Step 2: Full Substitution and Expansion of dx^2. To find dx^2, we square the entire Itô process expression $dx = adt + bdz$:

$$dx^2 = (a(x,t)dt + b(x,t)dz)^2$$

Expanding the binomial:

$$dx^2 = [a(x,t)]^2 dt^2 + 2a(x,t)b(x,t)dtdz + [b(x,t)]^2 dz^2$$

Applying our stochastic rules from the assumptions:

1. $[a(x,t)]^2 dt^2 \to 0$ (since $dt^2 = 0$).
2. $2a(x,t)b(x,t)dtdz \to 0$ (since $dt \cdot dz = 0$).
3. $[b(x,t)]^2 dz^2 \to [b(x,t)]^2 dt$ (since $dz^2 = dt$).

Therefore:

$$\boxed{dx^2 = b(x,t)^2 dt}$$

Step 3: Final Substitution into dG. Now we substitute the original dx and our newly derived dx^2 back into the Taylor expansion:

$$dG = \frac{\partial G}{\partial x}(adt + bdz) + \frac{\partial G}{\partial t}dt + \frac{1}{2}\frac{\partial^2 G}{\partial x^2}(b^2 dt)$$

Note: The dt^2 and $dxdt$ terms from the original Taylor expansion are dropped as they are $O(dt^{1.5})$ or higher.

Final Result: The Itô Formula. Factor out dt to group the deterministic components and dz for the random component:

$$\boxed{dG = \left(a\frac{\partial G}{\partial x} + \frac{\partial G}{\partial t} + \frac{1}{2}b^2\frac{\partial^2 G}{\partial x^2}\right)dt + b\frac{\partial G}{\partial x}dz}$$

For us, this means that the change in the function G (like an option price) is driven by three deterministic parts: underlying drift, time decay, and convexity (the Itô term), plus a random shock proportional to the asset's volatility.

3.6.2 GBM with Ito Lemma Application

Example: Forward Price. Consider a forward price contract maturing at time T:

$$G(S,t) = Se^{r(T-t)}$$

Partial Derivatives: $\frac{\partial G}{\partial S} = e^{r(T-t)}$ $\quad \frac{\partial^2 G}{\partial S^2} = 0$ $\quad \frac{\partial G}{\partial t} = -rG$

Resulting Process:

$$\boxed{dG = (\mu - r)Gdt + \sigma G dz}$$

In the risk-neutral world ($\mu = r$), the forward contract has zero drift.

Example: Log-Normal Process. To find the process for log-returns, we set the function

$$G(S,t) = \ln(S)$$

Partial Derivatives: $\frac{\partial G}{\partial S} = \frac{1}{S}$, $\frac{\partial^2 G}{\partial S^2} = -\frac{1}{S^2}$, $\frac{\partial G}{\partial t} = 0$.

The Log Process: $\boxed{d\ln(S) = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dz}$ *Confirms that logged returns are normally dist.*

Exact Transition Process for S. To find the future stock price S_T, we start with the process for log-returns derived via Itô's Lemma and integrate it over the time horizon $[0, T]$.

Step 1: The Differential of $\ln(S)$. we know that the log-process is:

$$d\ln(S) = \left(\mu - \frac{\sigma^2}{2}\right)dt + \sigma dz$$

Step 2: Integration over $[0, T]$. Integrating both sides from $t = 0$ to $t = T$:

$$G(S,t) = \int_0^T dG = \ln S_T - \ln S_0 = \int_0^T d\ln(S) = \int_0^T \left(\mu - \frac{\sigma^2}{2}\right)dt + \int_0^T \sigma dz$$

This yields the change in the log-price: $\boxed{\ln(S_T) - \ln(S_0) = \left(\mu - \frac{\sigma^2}{2}\right)T + \sigma z(T)}$

Step 3: Distribution of the Log-Price. Since $z(T) = \epsilon\sqrt{T}$ where $\epsilon \sim \mathcal{N}(0,1)$, we can write the distribution as:

$$\ln(S_T) = \phi\left[\ln(S_0) + \left(\mu - \frac{\sigma^2}{2}\right)T, \sigma\sqrt{T}\right]$$

Note: ϕ indicates that $\ln(S_T)$ follows a normal distribution with the specified mean and standard deviation.

Step 4: Exponentiating to solve for S_T. Taking the exponential of both sides to isolate S_T:

$$e^{\ln(S_T)} = e^{\ln(S_0)+(\mu-\frac{\sigma^2}{2})T+\sigma\epsilon\sqrt{T}}$$

$$\boxed{S_T = S_0 e^{(\mu-\frac{\sigma^2}{2})T+\sigma\epsilon\sqrt{T}}}$$

Logic: This result proves that while the stock price S_T is log-normally distributed, its log-returns are normally distributed. The $-\frac{\sigma^2}{2}$ term is the critical adjustment that accounts for the difference between the arithmetic mean and the geometric mean in a volatile market.

3.6.3 Multivariate Itô's Lemma: General Form

When dealing with multiple assets, we use a Taylor expansion that includes **cross-partial derivatives** to account for the interaction (correlation) between the variables.

The Vector-Matrix Representation For a function $G(\mathbf{S}, t)$ where $\mathbf{S}$ is a vector of n assets, the change dG can be expressed as:

$$dG = \left(\frac{\partial G}{\partial t} + \sum_i \frac{\partial G}{\partial S_i} a_i + \frac{1}{2} \sum_i \sum_j \rho_{ij} b_i d_j \frac{\partial^2 G}{\partial S_i \partial S_j} \right) dt + \sum_i \frac{\partial G}{\partial S_i} b_i dz_i$$

2-Asset Matrix Logic For two assets S_1 and S_2, the second-order term in the expansion is:

$$\frac{1}{2} \begin{pmatrix} dS_1 & dS_2 \end{pmatrix} \begin{pmatrix} \frac{\partial^2 G}{\partial S_1^2} & \frac{\partial^2 G}{\partial S_1 \partial S_2} \\ \frac{\partial^2 G}{\partial S_2 \partial S_1} & \frac{\partial^2 G}{\partial S_2^2} \end{pmatrix} \begin{pmatrix} dS_1 \\ dS_2 \end{pmatrix}$$

Using the stochastic rule $dz_1 dz_2 = \rho dt$, the matrix expansion simplifies to:

$$\frac{1}{2} \left[b_1^2 \frac{\partial^2 G}{\partial S_1^2} + 2\rho b_1 d_2 \frac{\partial^2 G}{\partial S_1 \partial S_2} + d_2^2 \frac{\partial^2 G}{\partial S_2^2} \right] dt$$

⋆ **Example 1: Product Process ($G = S \cdot X$) Step-by-Step Logic:**

- **Partial Derivatives**: $\frac{\partial G}{\partial S} = X$, $\frac{\partial G}{\partial X} = S$, $\frac{\partial^2 G}{\partial S \partial X} = 1$, $\frac{\partial^2 G}{\partial S^2} = 0$, $\frac{\partial^2 G}{\partial X^2} = 0$.
- **Substitution**: Only the drift and the cross-partial terms remain.
- **Result**: $dG = (SX\mu_X + SX\mu_S + \rho\sigma_S S\sigma_X X)dt + (\sigma_S S dz_S + \sigma_X X dz_X)S$.
- **Final Process**: $\boxed{dG = (\mu_X + \mu_S + \rho\sigma_S\sigma_X)Gdt + \hat{\sigma}Gd\hat{z}}$.

⋆ **Example 2: Ratio Process ($G = S/X$) Step-by-Step Logic:**

- **Partial Derivatives**: $\frac{\partial G}{\partial S} = \frac{1}{X}$, $\frac{\partial G}{\partial X} = -\frac{S}{X^2}$, $\frac{\partial^2 G}{\partial X^2} = \frac{2S}{X^3}$, $\frac{\partial^2 G}{\partial S \partial X} = -\frac{1}{X^2}$.
- **Expansion**:

$$dG = \left(\frac{1}{X}\mu_S S - \frac{S}{X^2}\mu_X X + \frac{1}{2}\sigma_X^2 X^2 \frac{2S}{X^3} - \rho\sigma_S S\sigma_X X \frac{1}{X^2} \right) dt + \dots$$

- **Final Process**: $\boxed{dG = (\mu_S - \mu_X + \sigma_X^2 - \rho\sigma_S\sigma_X)Gdt + \hat{\sigma}Gd\hat{z}}$.

Volatility Calculation ($\hat{\sigma}$): For both cases, the combined volatility is derived from the variance of the log-terms:

$$\hat{\sigma} = \sqrt{\sigma_S^2 + \sigma_X^2 \pm 2\rho\sigma_S\sigma_X}$$

3.6.4 Interest Rate Parity

⋆ **Interest Rate Parity (IRP)** Interest Rate Parity is a no-arbitrage condition which states that the difference in interest rates between two countries must be equal to the difference between the spot and forward exchange rates.

The Fundamental Identity: In an efficient market, the return on a domestic investment must equal the exchange-rate-adjusted return on a foreign investment.

$$(1 + r_d) = \frac{F_0}{S_0}(1 + r_f)$$

Rearranging for the **Forward Rate**:

$$\boxed{F_0 = S_0 \frac{1 + r_d}{1 + r_f}}$$

Continuous Compounding Form: In the context of the Black-Scholes/Garman-Kohlhagen models, we use continuous rates:

$$\boxed{F_0 = S_0 e^{(r_d - r_f)T}}$$

- **The Equilibrium Logic**: If $r_d > r_f$, everyone would want to hold domestic currency. To prevent this "free lunch," the forward rate F_0 must trade at a **discount** ($F_0 < S_0$) to offset the higher interest yield.

- **Covered vs. Uncovered**:

- **Covered IRP**: Uses a forward contract to "lock in" the future rate (Arbitrage-free).
- **Uncovered IRP**: Uses the *expected* future spot rate (Theoretical/Speculative).

- **The Link to Options**: This is why we treat r_f as the "dividend yield" (q) in currency option pricing. Just as a dividend reduces the forward price of a stock, the foreign interest rate reduces the forward price of the currency.

3.6.5 Example: Monte Carlo Simulation Multiple Assets

1. Parameters

- $S_0^1 = 100, \sigma_1 = 0.30$ (High Vol)
- $S_0^2 = 100, \sigma_2 = 0.15$ (Low Vol)
- $r = 0.05, T = 1, \rho_{1,2} = 0.6, K = 100$
- Strategy: **Long Call on** S^1 + **Long Put on** S^2

2. Cholesky Matrix (L) With $\rho = 0.6$, the transform is: $L = \begin{pmatrix} 1 & 0 \\ 0.6 & \sqrt{1-0.6^2} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0.6 & 0.8 \end{pmatrix}$

3. Trial 1 Calculation

3.1 Transform Noise (Cholesky) We start with independent draws $u_1 = 0.50$ and $u_2 = -0.20$.

- $\epsilon_1 = u_1 = \mathbf{0.50}$
- $\epsilon_2 = \rho u_1 + \sqrt{1-\rho^2}u_2 = (0.6)(0.50) + (0.8)(-0.20) = 0.30 - 0.16 = \mathbf{0.14}$

3.2 Evolve Asset 1 (High Vol: $\sigma_1 = 0.30$)

$$S_T^1 = 100 \cdot e^{(0.05 - \frac{0.30^2}{2})(1) + (0.50)(0.30)\sqrt{1}}$$

$$S_T^1 = 100 \cdot e^{(0.05-0.045)+0.15} = 100 \cdot e^{0.155}$$

$$S_T^1 = 100 \cdot 1.16765\cdots \approx \mathbf{116.77}$$

3.3 Evolve Asset 2 (Low Vol: $\sigma_2 = 0.15$)

$$S_T^2 = 100 \cdot e^{(0.05 - \frac{0.15^2}{2})(1) + (0.14)(0.15)\sqrt{1}}$$

$$S_T^2 = 100 \cdot e^{(0.05-0.01125)+0.021} = 100 \cdot e^{0.05975}$$

$$S_T^2 = 100 \cdot 1.06157\cdots \approx \mathbf{106.16}$$

3.4 Calculate Mixed Payoff ($K = 100$)

- **Call on S^1:** $\max(116.77 - 100, 0) = \mathbf{16.77}$
- **Put on S^2:** $\max(100 - 106.16, 0) = \mathbf{0}$
- **Total Payoff:** $16.77 + 0 = \mathbf{16.77}$

4. Trial Generation (5 Trials) We generate 5 pairs of independent noise (u_1, u_2), transform them to (ϵ_1, ϵ_2), and calculate terminal prices S_T and the Payoff P_i.

Trial (i)	u_1	u_2	ϵ_1	ϵ_2	S_T^1	S_T^2	Payoff P_i
1	0.50	-0.20	0.50	0.14	116.77	106.18	$16.77 + 0 = 16.77$
2	-1.10	0.40	-1.10	-0.34	72.25	96.08	$0 + 3.92 = 3.92$
3	0.80	1.20	0.80	1.44	127.76	131.65	$27.76 + 0 = 27.76$
4	-0.30	-0.90	-0.30	-0.90	91.85	88.25	$0 + 11.75 = 11.75$
5	0.10	0.05	0.10	0.10	103.56	105.53	$3.56 + 0 = 3.56$

4. Final Value Calculation The average of our 5 trials is:

$$\bar{P} = \frac{16.77 + 3.92 + 27.76 + 11.75 + 3.56}{5} = \frac{63.76}{5} = 12.752$$

Applying the discount factor e^{-rT}:

$$\text{Final Value (Price)} = 12.752 \cdot e^{-0.05(1)} = 12.752 \cdot 0.9512 \approx \mathbf{12.13}$$

Chapter 4

Implied Volatility and Non-Standard Options

4.1 Implied Volatility and the Volatility Surface

⋆ **Implied Volatility (σ_{Imp})** is the specific **volatility** value that, when plugged into the Black-Scholes formula, produces a theoretical price equal to the observed **market price**.

$$\boxed{Price_{Mkt}^{K,T} = BS(S, K, r, T, q, \sigma_{Imp})}$$

Implied volatility is the **market's "consensus" on future risk.** It is often described as the "wrong number in the wrong formula to get the right price,".

⋆ **Volatility Surface.** A **volatility surface** is the two-dimensional object of **implied volatility** across **strikes (K)** and **maturity (T)**: $\boxed{\sigma_{\text{imp}} = \sigma_{\text{imp}}(K, T)}$ It summarizes how option-implied volatility changes across contracts, not just for one option.

Comprised of two dimensions

- **Term Structure:** relationship between **implied volatility** and **time to maturity** T, describes the time dimension;
- **Skew:** relationship between **implied volatility** and **strike** K, describes the strike dimension

Put and Call Implied Volatility. For a given **strike (K)** and **maturity (T)**, the implied volatility for a call option and a put option must be identical to satisfy **put-call parity** $\boxed{p + S_0 e^{-qT} = c + Ke^{-rT}}$

⋆ **VIX Implied Volatility Index.** The **VIX** is a tradeable index representing a gauge of **near-term** (30-day) volatility derived from traded S&P 500 index option prices ($P(K)$ and $C(K)$).

$$\boxed{VIX = \sqrt{\frac{2e^{rT}}{T}\left(\int_0^F \frac{P(K)}{K^2}dK + \int_F^\infty \frac{C(K)}{K^2}dK\right)}}$$

The integral formulation effectively "sums up" the prices of all options across the strike spectrum. By weighting by $1/K^2$, the calculation isolates the variance of the log-returns, providing a model-independent measure of market sentiment.

1) Construction: It is an **average 30-day volatility** constructed from options with maturities ranging from 23 to 37 days.

2) Weighting: Uses a **weighted average** of **out-of-the-money (OTM) strikes** to capture the full distribution of expected price moves.

4.2 Volatility Term Structure

Volatility Term Structure

⋆ **Volatility Term Structure.** Describes the relationship between **implied volatility** and an option's **time to maturity** (T).
Just as the yield curve shows interest rates for different tenors, the term structure shows how the market prices risk over different horizons. It reflects whether the market expects current volatility levels to persist, mean-revert, or spike in the future.

- **Downward Sloping**: Associated with **stressed economic environments**. Short-term volatility spikes due to immediate uncertainty, while long-term volatility remains lower due to expected **mean reversion**.
- **Upward Sloping**: Generally associated with **low-stress economic environments**. Reflects current calm markets with an expectation that volatility will eventually revert to higher historical averages over time.

Variance is additive over time, so total variance over horizon T is the sum (or integral) of instantaneous variances, and **volatility** is the square root of average variance. This is why term structure is naturally modeled through variance first, then converted back to volatility.

⋆ **Discrete volatility for maturity T.** For T time steps with lengths $\Delta_{t-1,t}$, $\boxed{\sigma(T) = \sqrt{\frac{1}{T}\sum_{t=1}^{T} \sigma_{t-1,t}^2 \, \Delta_{t-1,t}}}$

$\sigma(T)$: **implied term volatility** for maturity T
$\sigma_{t-1,t}$: **instantaneous/step volatility** on interval $(t-1, t)$
The sum builds total variance across intervals. Dividing by T gives average variance per unit time, and the square root gives annualized-style volatility.

⋆ **Continuous volatility over maturity T** $\boxed{\sigma(T) = \sqrt{\frac{1}{T}\int_0^T \sigma_t^2 \, dt}}$

- σ_t: **instantaneous volatility** at time t
- $\int_0^T \sigma_t^2 dt$: **accumulated variance** over $[0, T]$

This is the continuous-time version of the same idea: integrate variance, average it over time, then take the square root.

Parametric representation. A common modeling step is to parameterize instantaneous volatility as a smooth function of time and state variables: $\boxed{\sigma_t = f(t, X)}$

s.t. t: time; X: other drivers / state variables (market factors)

Instead of specifying a separate volatility at every point, we use a smooth function f to capture the shape of the term structure parsimoniously.

Example: Continuous-Time Volatility.
Assume instantaneous volatility is linear in time: $\sigma_t = f(t) = a + bt$
Average variance over $[0, T]$: $\sigma(0,T)^2 = \frac{1}{T}\int_0^T \sigma_t^2 \, dt = \frac{1}{T}\int_0^T (a+bt)^2 \, dt = \frac{(a+bT)^3 - a^3}{3bT}$

Hence $\sigma(0,T) = \sqrt{\frac{(a+bT)^3 - a^3}{3bT}}$

Example: Discrete-Time Volatility. Assumptions (annualized step volatilities). σ_1: year 1; σ_2: years 2 and 3; σ_3: years 4 and 5

With $\Delta_t = 1$, the 5-year accumulated variance is: $\sigma_1^2 + 2\sigma_2^2 + 2\sigma_3^2$

So the 5-year average **volatility** is: $\bar{\sigma} = \sqrt{\frac{1}{5}\left(\sigma_1^2 + 2\sigma_2^2 + 2\sigma_3^2\right)}$

4.2.1 Spot vs Forward Volatility/Variance

⋆ **Spot vs Forward Volatility/Variance .** The key distinction is **time interval**. **Forward** quantities refer to a specific future interval, while **spot (term)** quantities summarize the whole horizon from today to T.

- **Spot (term) vol $\sigma^{\text{spot}}(T)$:** use for pricing/quoting a maturity-T option today **B.S.**
- **Forward vol $\sigma^{\text{fwd}}(T_1, T_2)$:** use for future-period analysis, simulation steps, and decomposing term structure.

Forward variance on interval $(T_{i-1}, T_i]$. The variance rate assigned to a specific future bucket:

$$\boxed{\text{Variance contribution on } (T_{i-1}, T_i] = \sigma^2_{i,\text{fwd}}(T_i - T_{i-1})}$$

This is the building block used period-by-period (useful for simulation and term-structure construction).

Cumulative variance to maturity T. Total variance from time 0 to T:

$$\boxed{V(0,T) = \int_0^T \sigma^2_{i,\text{fwd}}(u)\,du} \text{ (Cts.)} \qquad \boxed{V(0,T_n) = \sum_{i=1}^{n} \sigma^2_{i,\text{fwd}}(T_i - T_{i-1})} \text{ (Disc.)}$$

Variance adds across time intervals, so cumulative variance is just the sum/integral of forward variances.

Spot variance/volatility (term-average variance) to maturity T. Average variance over the full horizon $[0, T]$: $\boxed{\sigma^{(2)\text{spot}}(T) = \frac{V(0,T)}{T}}$ $\qquad \boxed{\sigma^{\text{spot}}(T) = \sqrt{\frac{V(0,T)}{T}}}$

This is the horizon-average volatility from today to T. In option markets, this is often the quoted maturity volatility (for a chosen strike convention, e.g., ATM).

Forward volatility on interval $(T_1, T_2]$. Volatility for a future period only, extracted from variance differences:

$$\boxed{\sigma^{\text{fwd}}(T_1, T_2) = \sqrt{\frac{V(0,T_2) - V(0,T_1)}{T_2 - T_1}}} \quad \text{or with spot vol:} \quad \boxed{\sigma^{\text{fwd}}(T_1, T_2) = \sqrt{\frac{\sigma^{\text{spot}}(T_2)^2 T_2 - \sigma^{\text{spot}}(T_1)^2 T_1}{T_2 - T_1}}}$$

Forward volatility isolates the market-implied volatility for a future slice, not the full path from today.

4.3 Volatility Skew

Volatility Skew

⋆ **Volatility Skew. Volatility skew** is the relationship between **implied volatility** and **strike** for options with the same **maturity** T: $\boxed{K \mapsto \sigma_{\text{imp}}(K, T) \qquad \text{(fixed } T\text{)}}$

Spot-volatility skew quotation. Plot implied volatility against **moneyness** based on spot:

$$\boxed{\text{plot } \sigma_{\text{imp}} \text{ vs } \frac{K}{S_0}}$$

Forward-volatility skew quotation. Plot implied volatility against **forward moneyness**:

$$\boxed{\text{plot } \sigma_{\text{imp}} \text{ vs } \frac{K}{F_0}}$$ Requires accounting for rates and dividends in F_0 $F_0 \approx S_0 e^{(r-q)T}$

Equity skew shape (intuition). For equities (e.g., SPX), lower strikes often have higher implied volatility than ATM or higher strikes: $\boxed{K \downarrow \Rightarrow \sigma_{\text{imp}}(K,T) \uparrow \quad \text{(common equity smirk)}}$

- This is a downside-protection effect: OTM puts become more expensive in volatility terms because crash risk is priced more heavily.

- Short-dated skews are often steeper because near-term tail-risk pricing can be more intense.

⋆ **Implied equity distribution interpretation.** Skew encodes information about the **market-implied distribution** of future equity returns relative to the Black–Scholes lognormal benchmark.

- **Flat volatility (B-S):** implies a lognormal return distribution benchmark
- **Higher left-tail vol:** more implied weight on negative-return outcomes (vs benchmark)
- **Higher right-tail vol:** more implied weight on positive-return tail outcomes (vs benchmark)

Skew is not just a pricing pattern. It is a compact market signal about asymmetry and tail risk in the return distribution.

⋆ **Delta with Skew (IV not constant). Total delta when implied volatility moves with spot.** If **implied volatility** depends on S, the option price changes through both **spot** and **IV**:

$$\boxed{\frac{dC}{dS} = \frac{\partial C}{\partial S} + \frac{\partial C}{\partial \sigma_{\text{IV}}}\frac{\partial \sigma_{\text{IV}}}{\partial S}} \qquad \textit{hence} \qquad \boxed{\Delta_{\text{total}} = \Delta_{\text{BS}} + \text{Vega} \cdot \frac{\partial \sigma_{\text{IV}}}{\partial S}}$$

A spot move can also move the vol surface, so the effective delta is not just Black–Scholes delta.

Equity skew implication. In equities, **spot** (returns) and **implied volatility** are often negatively related: $\boxed{\frac{\partial \sigma_{\text{IV}}}{\partial S} < 0}$ And for long options, Vega > 0, so: $\boxed{\Delta_{\text{total}} < \Delta_{\text{BS}} \quad \text{(typically for equity options)}}$

When spot rises, IV often falls (and vice versa), which adds a vega-driven correction to delta.

⋆ **Delta-Hedged Portfolio P&L (constant σ, $r = 0$).** For a long delta-hedged option, a standard small-step approximation is: $\boxed{P\&L \approx \frac{1}{2}\Gamma(\Delta S)^2 + \Theta\,\Delta t + \text{(higher-order, other terms)}}$ *After delta hedging, the main remaining drivers are **gamma** and **theta**.*

Using Black–Scholes relation (constant σ). For small Δt, with $r = 0$, $\boxed{\Theta \approx -\frac{1}{2}\Gamma S^2 \sigma_{\text{IV}}^2}$

Substitute into P&L: $\boxed{P\&L \approx \frac{1}{2}\Gamma S^2 \left[\left(\frac{\Delta S}{S}\right)^2 - \sigma_{\text{IV}}^2 \Delta t\right]}$

The delta-hedged P&L compares realized squared return over the step to the implied variance paid in the option price.

Key interpretation (IV vs realized vol).

- If realized variance > implied variance, long gamma / delta-hedged option tends to profit.
- If realized variance < implied variance, theta decay dominates and P&L tends to be negative.
- If implied and realized volatility match exactly (ideal continuous hedging), expected P&L is approximately zero.

This is the core logic behind long-vol vs short-vol trading: you are effectively trading realized variance against implied variance.

4.4 Volatility Skew Models

4.4.1 Local Volatility Models

Local Volatility Models

Local volatility models, are s.t. volatility is a deterministic function of **stock price** and **time**: $\boxed{\sigma = \sigma(S,t)}$ which gives flexibility to fit the observed implied-vol surface.

Under the General RN Stochastic Process: $\boxed{\dfrac{dS}{S} = (r-q)\,dt + \sigma(S,t)\,dZ}$

- Can fit a wide range of **skew/smile** patterns
- Deterministic volatility function (not stochastic vol)
- Flexible, but may **overfit** noisy market quotes

⋆ **CEV Model (Constant Elasticity of Variance). Specification.** A simple local-vol model where volatility depends on the stock level via an elasticity parameter β:

$$\boxed{\frac{dS}{S} = (r-q)\,dt + \sigma S^{\beta-1}\,dZ} \quad \text{equivalently} \quad \boxed{dS = (r-q)S\,dt + \sigma S^{\beta}\,dZ}$$

- $\beta = 1$: constant volatility (Black–Scholes case)
- $\beta > 1$: volatility rises as stock price rises; upward-sloping IV vs strike
- $\beta < 1$: volatility falls as stock price rises; downward-sloping IV vs strike (equity-like smirk)

Changing β changes how volatility reacts to the price level, which changes the skew direction.

CEV is a tractable local-vol model that can generate skew through state dependence in volatility.

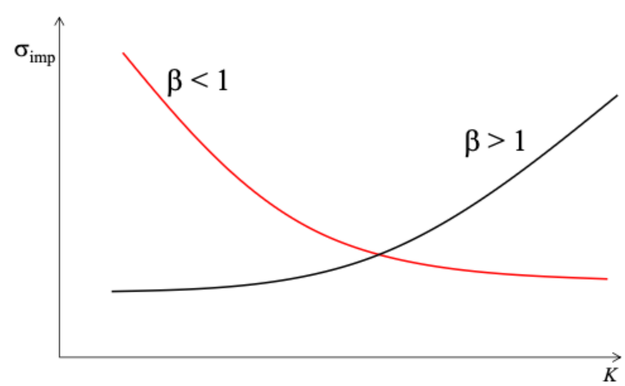

CEV closed-form pricing structure. European call/put prices admit **closed-form** expressions that are Black–Scholes-like in layout, but with cumulative **non-central chi-square** functions replacing the normal CDF.

So CEV remains analytically tractable for European options while allowing state-dependent volatility and skew.

Case $0 < \beta < 1$:

$$\boxed{c = Se^{-qT}\left[1 - \chi^2(a, b+2, c)\right] - Ke^{-rT}\chi^2(c, b, a)} \quad \boxed{p = Ke^{-rT}\left[1 - \chi^2(c, b, a)\right] - Se^{-qT}\chi^2(a, b+2, c)}$$

Case $\beta \geq 1$:

$$\boxed{c = Se^{-qT}\left[1 - \chi^2(c, -b, a)\right] - Ke^{-rT}\chi^2(a, 2-b, c)} \quad \boxed{p = Ke^{-rT}\left[1 - \chi^2(a, 2-b, c)\right] - Se^{-qT}\chi^2(c, -b, a)}$$

Parameter definitions:

$$\boxed{a = \frac{\left[Ke^{-(r-q)T}\right]^{2(1-\beta)}}{(1-\beta)^2\,v}, \qquad b = \frac{1}{1-\beta}, \qquad c = \frac{S^{2(1-\beta)}}{(1-\beta)^2\,v}, \qquad v = \frac{\sigma^2}{2(r-q)(\beta-1)}\left[\beta e^{2(r-q)(\beta-1)T} - 1\right]}$$

The exact special-function form depends on whether $\beta < 1$ or $\beta \geq 1$, but the economic role of β is the same: it controls how volatility scales with the stock level and therefore the skew direction.

⋆ **Implied Volatility Function (IVF) / Local-Vol Parametrization.** Start from the standard constant-vol model: $dS = (r - q)S\,dt + \sigma S\,dZ$ and replace it with a state-time dependent vol: $\boxed{dS = (r - q)S\,dt + \sigma(S,t)S\,dZ}$ The **function $\sigma(S,t)$ is parameterized/calibrated so model prices match observed option prices** across strikes and maturities.

⋆ **Local Volatility Function from Market Option Prices (Dupire-type formula).** A local-vol function can be extracted from the surface of European call prices $c_{\text{mkt}}(K,t)$:

$$\boxed{[\sigma(K,t)]^2 = \frac{2\left(\dfrac{\partial c_{\text{mkt}}}{\partial t} + q(t)c_{\text{mkt}} + K[r(t) - q(t)]\dfrac{\partial c_{\text{mkt}}}{\partial K}\right)}{K^2\dfrac{\partial^2 c_{\text{mkt}}}{\partial K^2}}}$$

This is the key local-vol result: with a smooth arbitrage-free option surface, you can infer a local volatility function consistent with all European option prices.

⋆ **Implied Stock Distribution (Risk-Neutral Density from Option Prices).** Recover the **implied risk-neutral distribution** of terminal stock price S_T from traded European call prices across strikes. Option prices across strikes contain information about how the market prices probability mass at expiration under the risk-neutral measure.

Risk-neutral representation of a call price. Let $g(S_T)$ denote the **risk-neutral density** of S_T. Then: $\boxed{c(K,T) = e^{-rT}\int_K^\infty (S_T - K)\,g(S_T)\,dS_T}$ *A call price is the discounted expected payoff under the risk-neutral density.*

Differentiate with respect to strike K (Breeden–Litzenberger logic).

$$\textit{First derivative: } \frac{\partial c}{\partial K} = -\,e^{-rT}\int_K^\infty g(S_T)\,dS_T \quad \textit{Second derivative:} \frac{\partial^2 c}{\partial K^2} = e^{-rT}g(K)$$

Hence, the **risk-neutral density** is: $\boxed{g(K) = e^{rT}\dfrac{\partial^2 c}{\partial K^2}}$ The curvature of the call-price curve in strike identifies the implied probability density at that strike.

Discrete Approximation from Traded Strikes (Butterfly Approximation). A symmetric call butterfly isolates local **curvature** of call prices versus strike, and that curvature maps directly to **risk-neutral density**. (*Higher butterfly value at a strike means more implied probability mass near that strike*)
With strike spacing δ and call prices (c_1, c_2, c_3) at neighboring strikes $(K - \delta, K, K + \delta)$, approximate the second derivative by a central difference:

$$\boxed{\frac{\partial^2 c}{\partial K^2} \approx \frac{c_1 - 2c_2 + c_3}{\delta^2}} \quad \textit{Hence the implied density estimate is: } \boxed{g(K) \approx e^{rT}\frac{c_1 - 2c_2 + c_3}{\delta^2}}$$

This is why a narrow butterfly spread extracts local risk-neutral probability around strike K.

4.4.2 Jump Diffusion Models (Merton)

Jump Diffusion Models (Merton)

Motivation: Asset returns exhibit more large moves (**jumps**) than a pure log-normal diffusion predicts, which affects option pricing and hedging. A diffusion model alone underestimates tail events; adding jumps helps capture heavy tails and skew/smile features seen in market option prices.

⋆ **Poisson Process (jump counts).** Let $X(t)$ count the number of jump events over t years with intensity λ. Then

$$\boxed{P(X(t) = k) = \frac{(\lambda t)^k e^{-\lambda t}}{k!}} \qquad \boxed{E[X(t)] = \lambda t \quad \text{and} \quad \mathrm{Var}[X(t)] = \lambda t}$$

λ is the expected jump frequency per unit time. The Poisson process controls *how many* jumps occur, not their sizes.

⋆ **Merton Jump-Diffusion Model. Dynamics.** Stock returns follow a standard log-normal diffusion plus a Poisson jump component: $\boxed{\frac{dS}{S} = (r - q - \lambda\kappa)\, dt + \sigma\, dz + (J-1)\, dN}$

- dN: **Poisson jump process** increment
- λ: jump frequency (intensity)
- J: random jump multiplier
- μ_J: expected jump-size parameter
- σ_J: jump-size volatility of $\ln(J)$
- $\kappa = e^{\mu_J + \frac{1}{2}\sigma_J^2} - 1$: expected jump return $E[J-1]$
- $\ln(J) \sim \mathcal{N}(\mu_J, \sigma_J)$: jump-size distribution (log-normal)
- $E[dz, dN] = 0$: independent stochastic shocks

Jump-size specification (slide notation)

$$\boxed{\ln(J) \sim \mathcal{N}(\mu_J, \sigma_J); \qquad \kappa = E[J-1] = e^{\mu_J + \frac{1}{2}\sigma_J^2} - 1; \qquad E[dz, dN] = 0}$$

The term $-\lambda\kappa$ in the drift is the jump compensator, included so the model remains risk-neutral after adding jumps.

⋆ **Option Pricing Merton Jump-Diffusion** European option price can be written as a weighted sum of Black–Scholes prices over possible jump counts:

Poisson-mixture of Black–Scholes prices: $\boxed{V = \sum_{n=0}^{\infty} \frac{e^{-\lambda T}(\lambda T)^n}{n!} BS(S, K, T, r_n, \sigma_n, q)}$

with adjusted parameters $\boxed{r_n = r - q - \lambda\kappa + \frac{n}{T}\left(\mu_J + \frac{1}{2}\sigma_J^2\right)}$ $\boxed{\sigma_n = \sqrt{\sigma^2 + \frac{n\sigma_J^2}{T}}}$

Conditional on n jumps, the option is priced like a Black–Scholes contract with modified drift and volatility; then average across n using Poisson probabilities.

⋆ **Stock-Price Solution (Merton Jump-Diffusion):** $\boxed{S_t = S_0 \exp\left[\left(r - \lambda\kappa - \frac{\sigma^2}{2}\right) t + \sigma\sqrt{t}\,\varepsilon\right] \prod_{i=1}^{n(t)} J_i}$

- $n(t)$: number of jumps up to time t, driven by a Poisson process with intensity λ
- J_i: multiplicative jump sizes

The price is the usual diffusion component multiplied by the product of jump multipliers. For small time steps, it is common to assume at most one jump can occur.

Merton Jump-Diffusion workflow

1. Inputs: $S_0,\ r(T),\ q(T),\ \{(K_i, T_i, C_i^{\text{mkt}})\}_{i=1}^N$ or $\{(K_i, T_i, \sigma_{\text{imp},i}^{\text{mkt}})\}_{i=1}^N$

Use option quotes across strikes/maturities (especially short-dated options, where jumps matter more for smile fit).

2. Calibrated parameters (RN Merton model): $\boxed{\Theta_{\text{MJD}} = (\sigma, \lambda, \mu_J, \sigma_J)}$ (get λ) S.T.

$$\boxed{\Theta^\star_{\text{MJD}} = \arg\min_{\Theta} \sum_i w_i \Big(C_i^{\text{MJD}}(\Theta) - C_i^{\text{mkt}}\Big)^2} \quad \text{(or same in IV space)}$$

- σ: diffusion volatility
- λ: jump intensity (expected jump frequency)
- μ_J, σ_J: jump-size distribution parameters for $\ln(J)$

4. Pricing equation used in calibration (model engine). For each (K_i, T_i), compute Merton model price (Poisson mixture of Black–Scholes prices): $\boxed{C^{\text{MJD}}(K,T;\Theta) = \sum_{n=0}^{\infty} \frac{e^{-\lambda T}(\lambda T)^n}{n!} BS(S, K, T, r_n, \sigma_n, q)}$

with $\boxed{r_n = r - q - \lambda\kappa + \frac{n}{T}\left(\mu_J + \frac{1}{2}\sigma_J^2\right), \qquad \sigma_n = \sqrt{\sigma^2 + \frac{n\sigma_J^2}{T}}}$ $\quad \boxed{\kappa = e^{\mu_J + \frac{1}{2}\sigma_J^2} - 1}$

Equation repeatedly evaluated during calibration. The Poisson weights depend directly on λ, so λ changes the mixture and the smile fit.

5. Model outputs (both are used). $\boxed{C^{\text{MJD}}(K,T;\Theta^\star_{\text{MJD}})} \implies \boxed{\sigma_{\text{imp}}^{\text{model}}(K,T)}$

Merton model naturally outputs a **price**; then convert to Black–Scholes implied volatility for direct comparison to market IV quotes.

λ from historical data (no option calibration).

1. Collect price data and log compute returns: $\boxed{r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)}$

2. Choose a jump-detection rule: Mark a return as a **jump** if it is unusually large relative to local volatility (simple threshold method): $\boxed{|r_t - \bar{r}| > c\,\hat{\sigma}_{\text{loc}}}$ where c is a threshold (e.g., 3 to 5) and $\hat{\sigma}_{\text{loc}}$ is a rolling volatility estimate.

More advanced methods exist (bipower variation, Lee–Mykland, etc.), but a threshold rule is the simplest starting point.

3. Count detected jumps: Let N_{jump} be the number of detected jump days in the sample. And *convert sample length to years (or your time unit)* $\quad T_{\text{sample}} \approx \frac{N}{252}$

5. Estimate jump intensity: $\boxed{\hat{\lambda}_{\text{hist}} = \frac{N_{\text{jump}}}{T_{\text{sample}}}}$ Estimates the average jump frequency (jumps per year) under the historical/physical measure.

6. (Optional) Estimate jump-size parameters. From detected jump returns $\{r_t^{(J)}\}$, estimate:

$$\boxed{\hat{\mu}_J = \text{mean of detected jump log-sizes}, \qquad \hat{\sigma}_J = \text{std. dev. of detected jump log-sizes}}$$

7. Use in Merton model (with caution). You may use $\hat{\lambda}_{\text{hist}}$ (and $\hat{\mu}_J, \hat{\sigma}_J$) as:

- initial guesses for calibration, or
- scenario inputs for risk analysis / simulation.

4.4.3 Stochastic Volatility Models

<u>Stochastic Volatility Models</u>

⋆ **Stochastic Volatility Models (Introduction).** Models, the volatility (or variance) is itself random and evolves over time, rather than being constant as in Black–Scholes. This lets the model generate changing volatility regimes, heavier tails, and realistic skew/smile patterns across strikes.

- **Skew shape** is strongly influenced by the **correlation between stock returns and variance (V).**

- **Fat tails** can arise naturally relative to constant-volatility Black–Scholes.

⋆ **Role of Correlation between V and S:**

Negative Corr V and S: When the stock falls and variance rises, downside strikes become relatively more expensive in implied-vol terms (equity smirk).
Produces a **downward-sloping volatility skew**; common in equity markets

Positive Corr V and S: When price increases are associated with higher variance, higher-strike options can carry higher implied vol.
Produces an **upward-sloping volatility skew**; sometimes observed in commodities

Zero (or near-zero) Corr V and S: With little directional return-vol correlation, the model can still generate tail thickness without a strong left/right skew tilt.
Can produce fatter tails on both sides; often associated with more symmetric smile-like patterns (e.g., some FX markets)

Variance state variable V_t (in stochastic volatility models). V_t is the **instantaneous variance process** of the asset return in the model, so the instantaneous volatility is $\sqrt{V_t}$ $\boxed{dS_t = \mu S_t\,dt + \sqrt{V_t}\,S_t\,dW_t}$
Hence V_t is a model state variable (latent process), not a directly observed market quote.

V_t is usually calibrated (inferred) from option prices / the implied-vol surface by fitting a stochastic-vol model (e.g., Heston), rather than directly measured from one option. $\boxed{\text{Option prices / IV surface} \implies \text{fit model parameters and inferred } V_t}$

⋆ **Three Common Stochastic Variance Dynamics** models mainly differ in how V_t evolves, especially whether it mean-reverts and whether it stays nonnegative.
State variable: Let V_t denote the **variance process** (so volatility is $\sqrt{V_t}$).

- **Heston (CIR-type variance).** $\boxed{dV_t = a(V_L - V_t)\,dt + \xi\sqrt{V_t}\,dZ_V}$
  Mean-reverting variance with state-dependent diffusion $\sqrt{V_t}$, which helps keep variance nonnegative and generates realistic equity-style skew.
- **Stein–Stein.** $\boxed{dV_t = a(V_L - V_t)\,dt + b\,dZ_V}$
  Mean-reverting variance with constant diffusion, but the process can become negative, so it is less natural for variance unless interpreted/modified carefully.
- **Hull–White (vol/variance diffusion form shown on slide).** $\boxed{dV_t = aV_t\,dt + bV_t\,dZ_V}$
  Multiplicative stochastic variance dynamics produce proportional randomness in V_t, giving flexible volatility behavior but without the same explicit mean-reversion form as Heston.

4.4.4 Heston Model (Stochastic Volatility)

⋆ Heston Model (Stochastic Volatility) Option value depends on two correlated sources of uncertainty: the traded asset price **stock** S_t and the non-traded **variance process** v_t.

Risk-neutral Heston dynamics:

$$\boxed{d\ln S_t = \left(r - q - \frac{1}{2}v_t\right)dt + \sqrt{v_t}\,dZ_1} \qquad \boxed{dv_t = \kappa(\theta - v_t)\,dt + \sigma_v\sqrt{v_t}\,dZ_2} \qquad \boxed{\mathbb{E}[dZ_1\,dZ_2] = \rho\,dt}$$

- r: risk-free rate $\quad$ q: dividend yield
- v_t: **instantaneous variance**, so volatility is $\sqrt{v_t}$
- v_0: initial variance
- θ: long-run variance level (mean-reversion target)
- κ: mean-reversion speed
- σ_v: vol of variance (**vol-of-vol**)
- ρ: correlation between stock and variance shocks, is a key driver of skew shape (e.g., negative ρ gives equity-like downside skew)

⋆ **Two-risk-factor hedging idea (derivation logic). Hedged portfolio construction.** Take two options $V(S,v,t)$, $V_1(S,v,t)$, and the stock: $\boxed{\Pi = V - \Delta_1 V_1 - \Delta S}$
Because there are two Brownian shocks (dZ_1, dZ_2), one stock hedge is not enough; a second option is used to hedge volatility risk.

$$d\Pi = \underbrace{\text{Term 1} + \text{Term 2}}_{\text{deterministic } dt \text{ change of } V, V_1} + \underbrace{\text{Equity Uncertainty}}_{\text{stock shock term}} + \underbrace{\text{Volatility Uncertainty}}_{\text{variance shock term}}.$$

$$d\Pi = \left(\frac{\partial V}{\partial t} + (r-q)S\frac{\partial V}{\partial S} + a(v_L - v)\frac{\partial V}{\partial v} + \frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2} + \rho\xi\sqrt{v}\,S\frac{\partial^2 V}{\partial S\partial v} + \frac{1}{2}\xi^2 v\frac{\partial^2 V}{\partial v^2}\right)dt$$

$$- \Delta_1\left(\frac{\partial V_1}{\partial t} + (r-q)S\frac{\partial V_1}{\partial S} + a(v_L - v)\frac{\partial V_1}{\partial v} + \frac{1}{2}vS^2\frac{\partial^2 V_1}{\partial S^2} + \rho\xi\sqrt{v}\,S\frac{\partial^2 V_1}{\partial S\partial v} + \frac{1}{2}\xi^2 v\frac{\partial^2 V_1}{\partial v^2}\right)dt$$

$$\boxed{+\underbrace{\left(\frac{\partial V}{\partial S} - \Delta_1\frac{\partial V_1}{\partial S} - \Delta\right)\sqrt{v}\,S\,dz_1 f}_{\text{Equity Uncertainty}} + \underbrace{\left(\frac{\partial V}{\partial v} - \Delta_1\frac{\partial V_1}{\partial v}\right)\xi\sqrt{v}\,dz_2}_{\text{Volatility Uncertainty}}}$$

Eliminate both sources of uncertainty.(equity uncertainty; volatility uncertainty) Choose hedge ratios so the dZ_1 and dZ_2 terms vanish:

$$\frac{\partial V}{\partial S} - \Delta_1\frac{\partial V_1}{\partial S} - \Delta = 0 \quad \textit{and} \quad \frac{\partial V}{\partial v} - \Delta_1\frac{\partial V_1}{\partial v} = 0$$

$$\textit{Hence solve for} \quad \boxed{\Delta_1 = \frac{\partial V/\partial v}{\partial V_1/\partial v}} \quad \boxed{\Delta = \frac{\partial V}{\partial S} - \frac{\partial V/\partial v}{\partial V_1/\partial v}\frac{\partial V_1}{\partial S}}$$

The second hedge ratio removes variance-shock exposure.

Riskless portfolio condition after solving for Δ and Δ_1. After choosing Δ and Δ_1 to eliminate the **Equity Uncertainty** and **Volatility Uncertainty** terms, the portfolio has only deterministic dt-terms, so it must earn the risk-free rate:

$$d\Pi = \underbrace{\text{Term 1}}_{\text{drift of } V} - \underbrace{\text{Term 2}}_{\text{drift of } (-\Delta_1 V_1)} = \underbrace{r\Pi\,dt}_{\text{risk-free return}} = \underbrace{r\left(V - \Delta S - \Delta_1 V_1\right)dt}_{\text{portfolio growth}}.$$

Portfolio is locally riskless, so it must earn the risk-free rate: $\boxed{d\Pi = r\Pi\,dt}$

Equation must hold for any option V and V_1: both ratios are set equal to the same function $\lambda(S,v,t)$, interpreted as the volatility risk-premium term (independent of contract-specific payoff details).

$$\frac{\underbrace{\left(\frac{\partial V}{\partial t}+\frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2}+\rho v\xi S\frac{\partial^2 V}{\partial S\partial v}+\frac{1}{2}q^2\frac{\partial^2 V}{\partial v^2}+(r-q)S\frac{\partial V}{\partial S}+a(v_L-v)\frac{\partial V}{\partial v}-rV\right)}_{\text{Heston drift terms of } V}}{\underbrace{(\partial V/\partial v)}_{\text{variance sensitivity}}}$$

$$=\frac{\underbrace{\left(\frac{\partial V_1}{\partial t}+\frac{1}{2}vS^2\frac{\partial^2 V_1}{\partial S^2}+\rho v\xi S\frac{\partial^2 V_1}{\partial S\partial v}+\frac{1}{2}q^2\frac{\partial^2 V_1}{\partial v^2}+(r-q)S\frac{\partial V_1}{\partial S}+a(v_L-v)\frac{\partial V_1}{\partial v}-rV_1\right)}_{\text{Heston drift terms of } V_1}}{\underbrace{(\partial V_1/\partial v)}_{\text{variance sensitivity}}}\boxed{=\underbrace{\lambda(S,v,t)}_{\text{vol risk premium}}}$$

The ratio of each derivative's Heston drift terms to its **variance sensitivity must be the same payoff-independent function $\lambda(S,v,t)$.**

This implies a common pricing relation across derivatives and introduces a **volatility risk premium term $\lambda(S,v,t)$,** giving the valuation PDE:

$$\boxed{\underbrace{\frac{\partial V}{\partial t}}_{\text{time decay}}+\underbrace{\frac{1}{2}vS^2\frac{\partial^2 V}{\partial S^2}}_{\text{stock gamma}}+\underbrace{\rho\sigma_v vS\frac{\partial^2 V}{\partial S\partial v}}_{\text{cross term}}+\underbrace{\frac{1}{2}\sigma_v^2 v\frac{\partial^2 V}{\partial v^2}}_{\text{variance convexity}}+\underbrace{(r-q)S\frac{\partial V}{\partial S}}_{\text{stock drift}}+\underbrace{\Big[\kappa(\theta-v)-\lambda(S,v,t)\Big]\frac{\partial V}{\partial v}}_{\text{variance drift}}-\underbrace{rV}_{\Pi\text{return}}=0}$$

Drift of the non-traded variance factor is adjusted by the volatility risk premium λ.

Heston model workflow

1. Inputs: S_0, $r(T)$, $q(T)$, $\{(K_i,T_i,C_i^{\text{mkt}})\}_{i=1}^N$ or $\{(K_i,T_i,\sigma_{\text{imp},i}^{\text{mkt}})\}_{i=1}^N$ *Use market option quotes across many strikes and maturities (the surface slice data you want the model to match). Rates/dividends go into the pricing formula and BS-IV conversion.*

2. Calibrated parameters. $\boxed{\Theta_{\text{Heston}}=(v_0,\kappa,\theta,\sigma_v,\rho)}$ *These are the parameters of the RN Heston dynamics you are fitting, not directly observed market quotes.*

3. Pricing equation you plug into (model engine). For each (K_i,T_i), compute the Heston model price using the Heston pricing formula / characteristic-function implementation (equivalently the Heston PDE solution): $\boxed{C_i^{\text{Heston}}(\Theta)=C^{\text{Heston}}(K_i,T_i;S_0,r,q,\Theta)}$ *This is the equation you repeatedly evaluate during calibration. In practice, you usually use the semi-closed-form Heston price (characteristic functions), not the PDE directly by hand.*

4. Calibration step (which equation is calibrated). Choose Θ_{Heston} so the **model prices** (or model IVs) best match the observed market surface: $\boxed{\Theta^\star_{\text{Heston}}=\arg\min_{\Theta}\sum_i w_i\Big(C_i^{\text{Heston}}(\Theta)-C_i^{\text{mkt}}\Big)^2}$ (or same objective in IV space) *You are calibrating the parameter vector Θ_{Heston}, not solving for a single implied volatility. Use many strikes/maturities so the fitted parameters capture skew/smile and term structure.*

5. Model outputs (both are used). $\boxed{C^{\text{Heston}}(K,T;\Theta^\star_{\text{Heston}})}\implies\boxed{\sigma_{\text{imp}}^{\text{model}}(K,T)}$ *Heston naturally outputs a **price**. You then invert Black–Scholes to report the model-implied volatility and compare directly to market IV quotes.*

4.4.5 Stochastic Volatility and Jump Diffusion Models (Bates Model)

<u>**Stochastic Volatility and Jump Diffusion Models (Bates Model)**</u>

Bates model combines **Heston stochastic volatility** with **Merton jump diffusion** under the risk-neutral measure. AKA It adds jumps to a stochastic-volatility model, so the model can **capture both volatility clustering/skew (from stochastic variance) and sudden large moves (from jumps).**

Bates RN model specification: $$\boxed{\underbrace{\frac{dS}{S}}_{\text{stock return}} = \underbrace{(r-q-\lambda\kappa)\,dt}_{\text{RN drift}} + \underbrace{\sqrt{v}\,dz_s}_{\text{diffusion shock}} + \underbrace{(J-1)\,dN}_{\text{jump term}}}$$

$$\boxed{\underbrace{dv}_{\text{variance change}} = \underbrace{a(v_L - v)\,dt}_{\text{mean reversion}} + \underbrace{\xi\sqrt{v}\,dz_v}_{\text{vol shock}}} \quad \textit{and} \quad \boxed{\underbrace{E[dz_s, dz_v]}_{\text{shock covariance}} = \underbrace{\rho\,dt}_{\text{correlation link}}}$$

- **Heston part**: v is stochastic variance (mean-reverting, correlated with stock shock)
- **Merton part**: dN gives random jump arrivals and J gives jump size
- Jumps are assumed uncorrelated with the other randomness

The drift adjustment $-\lambda\kappa$ is the jump compensator (same logic as Merton), while ρ still controls stock–variance shock correlation (same logic as Heston).

4.5 Non-Standard Options

4.5.1 Compound Option

Compound Option

⋆ **Compound Option.** is an option written on another **option**. The compound option has maturity T_1 and strike K_1; the underlying option has later maturity T_2 and strike K_2, with $T_1 < T_2$.
Interpretation of K_1 vs K_2. K_1 is the price paid at T_1 to acquire the **underlying option**, so it is compared to the **option value** $c(S_{T_1}, K_2, T_2 - T_1)$, not directly to stock price. K_2 is the strike of the underlying call, so it is compared to the **stock price** at T_2.

- **Underlying option**: put/call with (K_2, T_2).
- **Compound option**: put/call on that option with (K_1, T_1).

⋆ **Call on Call (CoC).** A **call on call** gives the right at T_1 to buy a call option expiring at T_2 with strike K_2, by paying K_1.
Payoff at T_1: If $c(S_{T_1}, K_2, T_2 - T_1)$ is the value at T_1 of the underlying call, then

$$\boxed{\mathrm{CoC}_{T_1} = \max(c(S_{T_1}, K_2, T_2 - T_1) - K_1,\ 0)}$$

This looks like a vanilla call payoff, except the "underlying asset" is now the *call value* $c(\cdot)$, not the stock itself.

Underlying call value at T_1: Under Black–Scholes, $\boxed{c(S_{T_1}, K_2, T_2 - T_1) = S_{T_1} N(d_1^*) - K_2 e^{-r(T_2-T_1)} N(d_2^*)}$

with $d_1^* = \frac{\ln(S_{T_1}/K_2) + \left(r + \frac{1}{2}\sigma^2\right)(T_2 - T_1)}{\sigma\sqrt{T_2 - T_1}}$, $\qquad d_2^* = d_1^* - \sigma\sqrt{T_2 - T_1}$.

Critical stock value $S_{T_1}^*$. $S_{T_1}^*$ is the stock price at time T_1 that makes the **underlying option** worth exactly K_1: $\boxed{c(S_{T_1}^*, K_2, T_2 - T_1) = K_1}$

$$\underbrace{c(S_{T_1}^*, T_1)}_{\substack{\text{value of the underlying call}\\ \text{when stock at } T_1 \text{ equals } S_{T_1}^*}} = \underbrace{S_{T_1}^* N(d_1^*)}_{\substack{\text{stock-related part of the}\\ \text{Black–Scholes call value}}} - \underbrace{K_2 e^{-r(T_2-T_1)} N(d_2^*)}_{\substack{\text{present value of paying strike } K_2\\ \text{at } T_2\text{, weighted by exercise probability}}} = \underbrace{K_1}_{\substack{\text{compound-option exercise price}\\ \text{paid at } T_1 \text{ to buy the call}}}$$

Choose the stock level $S_{T_1}^*$ so that the underlying call is worth exactly K_1. That stock level is the exercise cutoff for the compound option.

Hence $\boxed{\text{exercise at } T_1 \iff S_{T_1} \ge S_{T_1}^* \iff c(S_{T_1}, K_2, T_2 - T_1) \ge K_1}$

$S_{T_1}^*$ is the exercise cutoff at T_1: above it, the underlying call is worth paying K_1 for; below it, it is not.

⋆ **Closed-form CoC price.**

$$\boxed{\mathrm{CoC} = S N_2(a_1, b_1, \rho) - K_2 e^{-rT_2} N_2(a_2, b_2, \rho) - K_1 e^{-rT_1} N(a_2)}$$

$$a_1 = \frac{\ln(S/S_{T_1}^*) + \left(r + \frac{1}{2}\sigma^2\right) T_1}{\sigma\sqrt{T_1}}, \qquad a_2 = a_1 - \sigma\sqrt{T_1},$$

$$b_1 = \frac{\ln(S/K_2) + \left(r + \frac{1}{2}\sigma^2\right) T_2}{\sigma\sqrt{T_2}}, \qquad b_2 = b_1 - \sigma\sqrt{T_2}, \qquad \rho = \sqrt{\frac{T_1}{T_2}}.$$

$$\mathrm{CoC} = \underbrace{S\, N_2(a_1, b_1, \rho)}_{\substack{\text{stock-related payoff contribution}\\ \text{from the joint exercise region}}} - \underbrace{K_2 e^{-rT_2} N_2(a_2, b_2, \rho)}_{\substack{\text{present value of paying the underlying}\\ \text{call strike } K_2 \text{ at } T_2}} - \underbrace{K_1 e^{-rT_1} N(a_2)}_{\substack{\text{present value of paying the compound}\\ \text{strike } K_1 \text{ at } T_1}}$$

The formula has three parts: stock contribution, discounted strike K_2 of the underlying call, and discounted first-stage cost K_1. The **bivariate normal** appears because value depends jointly on stock outcomes at T_1 and T_2.

$N_2(a_1, b_1, \rho)$, $N_2(a_2, b_2, \rho)$: joint risk-neutral probabilities involving the T_1 and T_2 exercise events
$N(a_2)$: risk-neutral probability of exercising the compound option at T_1
a_1, a_2: terms tied to the first decision date T_1 and the critical stock level $S^*_{T_1}$
b_1, b_2: terms tied to the underlying call strike K_2 and final maturity T_2
ρ: correlation between $\ln S_{T_1}$ and $\ln S_{T_2}$

Bivariate normal cdf gives the joint probability that two correlated standard normal variables are below given cutoffs: $\boxed{N_2(x, y; \rho) = \Pr(Z_1 \le x,\ Z_2 \le y)}$ where (Z_1, Z_2) are standard normal with correlation ρ. If $\rho = 0$, then $N_2(x, y; 0) = N(x)N(y)$

RN representation CoC: $\boxed{\text{CoC} = e^{-rT_1} \int_{S^*_{T_1}}^{\infty} \left(c(S_{T_1}, K_2, T_2 - T_1) - K_1\right) f(S_{T_1})\, dS_{T_1}}$

Discount the expected payoff at T_1, integrating only over the exercise region $S_{T_1} \ge S^*_{T_1}$.

4.5.2 Asian Options

Asian Options

⋆ **Asian Option:** An option whose payoff depends on an **average price** of the underlying over the life of the contract, rather than only on S_T.
Main payoff types. Let S_{ave} be the contract average price.

- **Average-Price Asian option**: strike is fixed, payoff depends on S_{ave}

$$\boxed{\text{Call} = \max(S_{\text{ave}} - K, 0), \qquad \text{Put} = \max(K - S_{\text{ave}}, 0)}$$

- **Average-Strike Asian option**: strike is the average price

$$\boxed{\text{Call} = \max(S_T - S_{\text{ave}}, 0), \qquad \text{Put} = \max(S_{\text{ave}} - S_T, 0)}$$

Arithmetic average: $\boxed{S^{(A)}_{\text{ave}} = \frac{1}{N}\sum_{i=1}^{N} S_i}$ **Geometric average**: $\boxed{S^{(G)}_{\text{ave}} = \left(\prod_{i=1}^{N} S_i\right)^{1/N}}$

Arithmetic averaging is market-relevant but harder to price exactly. Geometric averaging is easier mathematically because the geometric average of lognormal prices remains lognormal.

⋆ **Arithmetic average option: Turnbull–Wakeman approximation.** Approximate solution by assuming the average is log-normal and than matching first two moments. For a newly issued average-price call/put, $\boxed{c_{TW} = e^{-rT}\Big[F_0 N(d_1) - K N(d_2)\Big], \qquad p_{TW} = e^{-rT}\Big[K N(-d_2) - F_0 N(-d_1)\Big]}$

with $d_1 = \frac{\ln(F_0/K) + \frac{1}{2}\sigma_A^2 T}{\sigma_A \sqrt{T}}, \qquad d_2 = d_1 - \sigma_A \sqrt{T}$

F_0: adjusted **mean** / forward-style level for the average $\boxed{F_0 = S\sum_{i=1}^{T} e^{rt_i}}$ collects the expected contribution of each future sampled stock price into one adjusted level.

σ_A^2: adjusted **variance** of the arithmetic average $\boxed{\sigma_A^2 = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n} S_0^2\, e^{-r(t_i+t_j)} \left(e^{\sigma^2 \min(t_i, t_j)} - 1\right)}$

This formula measures how variable the arithmetic average is by adding all pairwise covariances between sampled prices. The $\min(t_i, t_j)$ *appears because two stock values are correlated only through the part of the path they share up to the earlier time.*

Existing arithmetic-average option. If part of the averaging period has already occurred, with realized average $\bar{S}$ over elapsed time t_1 and remaining averaging time t_2, then

$$\boxed{\max\left(\frac{\bar{S}\,t_1 + S_{\text{ave}}t_2}{t_1+t_2} - K, 0\right) = \frac{t_2}{t_1+t_2}\max(S_{\text{ave}} - K^*, 0)} \quad \text{where} \quad \boxed{K^* = \frac{t_1+t_2}{t_2}K - \frac{t_1}{t_2}\bar{S}}$$

Once part of the average is already fixed, the remaining problem becomes a scaled option on the future average with an adjusted strike K^*.

⋆ **Geometric average option: exact-style Black–Scholes form.** Because the **geometric average** is lognormal, closed-form formulas are available:

$$\boxed{c_{GV} = e^{-(b_A-r)T}SN(d_1) - Ke^{-rT}N(d_2)} \qquad \boxed{p_{GV} = -e^{-(b_A-r)T}SN(-d_1) + Ke^{-rT}N(-d_2)}$$

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

b_A: adjusted drift of the geometric average: $\boxed{b_A = \frac{1}{2}\left(r - \frac{\sigma^2}{6}\right)}$

σ_A: adjusted volatility of the geometric average: $\boxed{\sigma_A = \frac{\sigma}{\sqrt{3}}}$

4.5.3 Multi-Asset Option

Multi-Asset Option

⋆ **Basket/ Multi-asset Option:** Option written on a portfolio of assets. Its payoff depends on a function of several underlying assets, often a weighted sum. Value depends not only on individual volatilities, but also on the **correlation structure across assets**.

⋆ **Moment Matching Style Black-style approximation for basket:**
RN Forward prices F_i, for $i = 1, \ldots, N$

Basket call payoff: for the unweighted basket $\boxed{\text{Payoff} = \max\left(\sum_{i=1}^{N} S_i^T - K,\ 0\right)}$

First moment of the basket forward: $\boxed{M_1 = \sum_{i=1}^{N} F_i}$ mean / forward level of the basket under RN

Second moment of the basket forward: $\boxed{M_2 = \sum_{i=1}^{N}\sum_{j=1}^{N} F_iF_j\, e^{\rho_{ij}\sigma_i\sigma_j T}}$ incorporating pairwise correlation effects. This is where dependence across assets enters. Higher positive correlation increases basket dispersion and typically raises call value.

Moment-matched Forward Volatility: $\boxed{\sigma_{\text{Bsk}}^2 = \frac{1}{T}\ln\left(\frac{M_2}{M_1^2}\right)}$ The basket is approximated as lognormal by matching its first two moments.

$$\boxed{V(F_1,\dots,F_N) = e^{-rT}\Big[M_1 N(d_1) - K N(d_2)\Big]} \quad \text{with } d_1 = \frac{\ln(M_1/K)+\frac{1}{2}\sigma_{\text{Bsk}}^2 T}{\sigma_{\text{Bsk}}\sqrt{T}}, \qquad d_2 = d_1 - \sigma_{\text{Bsk}}\sqrt{T}$$

This has the same form as Black's formula, but the single-asset forward is replaced by the basket forward M_1, and the volatility is replaced by the matched basket volatility σ_{Bsk}.

4.5.4 Foreign Exchange Options

Foreign Exchange Options

⋆ **Option on the exchange rate X.** The underlying is the **exchange rate** X, quoted as domestic currency per unit of foreign currency.
The contract payoff depends only on how the exchange rate moves. It is mathematically the same as a Black–Scholes option on an asset with continuous yield r_f.

- X_0: current exchange rate (domestic / foreign)
- X_T: exchange rate at maturity
- r: domestic risk-free rate
- r_f: foreign risk-free rate
- μ_X: real-world drift of the exchange rate
- σ_X: exchange-rate volatility
- K: strike exchange rate

Forward exchange rate: $\boxed{X_T^F = X_0 e^{(r-r_f)T}}$ *The foreign currency behaves like an asset that earns the foreign interest rate. So under no-arbitrage, the domestic carry is* $r - r_f$.

Real-world dynamics: $dX = \mu_X X\,dt + \sigma_X X\,dZ_X$

Risk-neutral dynamics: $\boxed{dX = (r - r_f)X\,dt + \sigma_X X\,dZ_X}$

For pricing, the real-world drift μ_X *is replaced by the risk-neutral carry* $r - r_f$.

European FX option pricing:

$\boxed{c = X_0 e^{-r_f T} N(d_1) - K e^{-rT} N(d_2)}$ $\boxed{p = K e^{-rT} N(-d_2) - X_0 e^{-r_f T} N(-d_1)}$

with $d_1 = \frac{\ln(X_0/K) + \left(r - r_f + \frac{1}{2}\sigma_X^2\right)T}{\sigma_X\sqrt{T}}$, $\qquad d_2 = d_1 - \sigma_X\sqrt{T}$

This is just Black–Scholes with $S_0 = X_0$ and dividend yield $q = r_f$. The term $X_0 e^{-r_f T}$ is the present value of receiving foreign currency, while $K e^{-rT}$ is the present value of paying domestic currency at exercise.

⋆ **Option on a foreign asset S.** Let S be the **foreign asset price** measured in foreign currency. A domestic investor cares about the domestic value $\boxed{S^* = XS}$ where X converts foreign currency into domestic currency.

- S_0: foreign asset price today
- S_T: foreign asset price at maturity
- q: foreign asset dividend yield
- μ_S: real-world drift of the foreign asset
- σ_S: foreign asset volatility
- K_d: strike in domestic currency
- K_f: strike in foreign currency
- $S^* = XS$: domestic price of the foreign asset
- $\rho_{S,X}$: correlation between asset and exchange-rate shocks

Real-world dynamics: $dS = (\mu_S - q)S\,dt + \sigma_S S\,dZ_S$

Risk-neutral dynamics: $\boxed{dS = (r_f - q)S\,dt + \sigma_S S\,dZ_S}$ FX has procces $dX = (r - r_f)X\,dt + \sigma_X X\,dZ_X$

with $\boxed{\rho_{S,X}\,dt = E[dZ_S\,dZ_X]}$

Under pricing measure, the foreign stock grows at foreign carry $r_f - q$, while the exchange rate grows at domestic-minus-foreign carry $r - r_f$. Correlation matters because the domestic value $S^* = XS$ depends on both shocks simultaneously.

⋆ **Quanto.** A **quanto** contract has payoff determined by a **foreign asset** but paid in a **different currency**, typically at a fixed conversion rule rather than the random future FX rate.

Hence, we get **two risk factors**: foreign asset S, exchange rate X s.t ρ_{SX} between asset and FX shocks. Thus we apply **Quanto adjustment**: correlation changes the domestic pricing drift of the foreign asset.

Quanto-adjusted forward. For a foreign asset payoff paid in domestic currency,

$$\boxed{E_{\text{dom}}[S_T] = E_{\text{fgn}}[S_T]e^{-\rho_{SX}\sigma_S\sigma_X T} = S_0 e^{(r_f - q - \rho_{SX}\sigma_S\sigma_X)T}}$$

$$\boxed{\mu_{\text{quanto}} = r_f - q - \rho_{SX}\sigma_S\sigma_X}$$

Positive correlation lowers the domestic risk-neutral drift of the foreign asset under fixed conversion, while negative correlation raises it. This drift shift is the core quanto adjustment.

⋆ **Foreign equity / exchange option types.** The slides give four standard payoffs:

- c_1**: Foreign equity call struck in foreign currency**: payoff is a standard foreign-stock call with strike in foreign currency, then converted into domestic currency at the random future FX rate.
$\boxed{c_1(S_T, X_T) = X_T \max(S_T - K_f, 0)}$
- c_2**: Foreign equity call struck in domestic currency**: payoff is a call on the **domestic value of the foreign stock** $X_T S_T$, with strike written directly in domestic currency.
$\boxed{c_2(S_T, X_T) = \max(X_T S_T - K_d, 0)}$
- c_3**: Fixed exchange-rate foreign equity call**: payoff is a foreign-stock call with strike in foreign currency, but converted using a **fixed exchange rate**$\bar{X}$.
$\boxed{c_3(S_T, X_T) = \bar{X} \max(S_T - K_f, 0)}$
- c_4**: Equity-linked FX call**: payoff is an FX call multiplied by the foreign stock level S_T, so the stock scales the currency-option payoff.
$\boxed{c_4(S_T, X_T) = S_T \max(X_T - K_X, 0)}$

First identify what is random in the payoff: foreign stock only, product XS, stock with fixed FX conversion, or FX option scaled by stock. Then match the contract to c_1, c_2, c_3, or c_4.

C_1**: foreign equity call struck in foreign currency.** Price the foreign-stock call first in foreign currency, then convert its value into domestic currency using X_0.

Payoff: $\boxed{c_1(S_T, X_T) = X_T \max(S_T - K_f, 0)}$

Value at $t = 0$: $\boxed{c_1(S_0, X_0) = X_0 \left[S_0 e^{-qT} N(d_1^{(1)}) - K_f e^{-r_f T} N(d_2^{(1)})\right]}$

with $\quad d_1^{(1)} = \frac{\ln(S_0/K_f) + \left(r_f - q + \frac{1}{2}\sigma_S^2\right)T}{\sigma_S\sqrt{T}}, \qquad d_2^{(1)} = d_1^{(1)} - \sigma_S\sqrt{T}$

Note: This is just the standard foreign-market stock option value, then converted into domestic currency by multiplying by X_0. There is no quanto adjustment because payoff is translated at the market exchange rate.

C_2**: foreign equity call struck in domestic currency.** Treat XS as the effective domestic underlying and price a call directly on the product $X_T S_T$.

Payoff: $\boxed{c_2(S_T, X_T) = \max(X_T S_T - K_d, 0)}$

Let $S^* = XS$. then $\boxed{\sigma_{SX}^2 = \sigma_S^2 + \sigma_X^2 + 2\rho_{SX}\sigma_S\sigma_X}$

Value at $t = 0$: $\boxed{c_2(S_0, X_0) = S_0 X_0 e^{-qT} N(d_1^{(2)}) - K_d e^{-rT} N(d_2^{(2)})}$

with $\quad d_1^{(2)} = \frac{\ln(S_0 X_0/K_d) + \left(r_d - q + \frac{1}{2}\sigma_{SX}^2\right)T}{\sigma_{SX}\sqrt{T}}, \qquad d_2^{(2)} = d_1^{(2)} - \sigma_{SX}\sqrt{T}$

Note: Here the underlying is the domestic asset value $S^* = XS$. So the problem reduces to Black–Scholes on the product XS, whose volatility includes both volatilities and the correlation term.

$\star$ C_3**: fixed exchange-rate foreign equity call (quanto call).** Price the foreign-stock call under the **quanto-adjusted drift**, then multiply by the fixed exchange rate $\bar{X}$.

Payoff: $\boxed{c_3(S_T, X_T) = \bar{X}\max(S_T - K_f, 0)}$

Value at $t = 0$: $\boxed{c_3(S_0) = \bar{X} e^{-r_d T}\left[S_0 e^{(r_f - q - \rho_{SX}\sigma_S\sigma_X)T} N(d_1^{(3)}) - K_f e^{r_d T} N(d_2^{(3)})\right]}$

with $\quad d_1^{(3)} = \frac{\ln(S_0/K_f) + \left(r_f - q - \rho_{SX}\sigma_S\sigma_X + \frac{1}{2}\sigma_S^2\right)T}{\sigma_S\sqrt{T}}, \qquad d_2^{(3)} = d_1^{(3)} - \sigma_S\sqrt{T}$

Equivalently, the **foreign stock has domestic pricing drift** $\boxed{r_f - q - \rho_{SX}\sigma_S\sigma_X}$

Note: This is the classic quanto result: the FX rate is fixed in the payoff, so FX randomness disappears from terminal payoff, but correlation still shifts the drift used for pricing. The stock volatility stays σ_S; the drift gets the quanto adjustment.

$\star$ C_4**: equity-linked FX call.** Price an FX call whose payoff is scaled by the foreign stock level, so use the FX framework with equity-linked adjustment.

Payoff: $\boxed{c_4(S_T, X_T) = S_T \max(X_T - K_X, 0)}$

Value at $t = 0$: $\boxed{c_4(S_0, X_0) = S_0 e^{-qT}\left[X_0 N(d_1^{(4)}) - K_X e^{-(r - r_f + \rho_{SX}\sigma_S\sigma_X)T} N(d_2^{(4)})\right]}$

with $\quad d_1^{(4)} = \frac{\ln(X_0/K_X) + \left(r - r_f + \rho_{SX}\sigma_S\sigma_X + \frac{1}{2}\sigma_X^2\right)T}{\sigma_X\sqrt{T}}, \qquad d_2^{(4)} = d_1^{(4)} - \sigma_X\sqrt{T}$

Note: This is an FX option scaled by the foreign equity level. Relative to a plain FX option, the drift picks up $+\rho_{SX}\sigma_S\sigma_X$ because the stock multiplies the FX payoff.

4.5.5 Variance and Volatility Swaps

Variance and Volatility Swaps

⋆ **Volatility swap.** Exchanges **realized volatility** against a fixed volatility strike K.
This is the linear contract in volatility itself, but it is not statically replicable as cleanly as a variance swap because the square-root makes the payoff nonlinear in variance.

$$\boxed{V_{\mathrm{vol}}(T) = L_{\mathrm{vol}}(\sigma_R - K)} \quad = \text{Payoff at Maturity}$$

- L_{vol}: volatility notional (often called vega notional)
- K: strike volatility
- σ_R: realized volatility

⋆ **Variance swap.** Exchanges **realized variance** over $[0, T]$ against a fixed variance strike K^2.
It is a pure volatility-trading instrument: payoff rises when realized variance exceeds the preset variance strike. It is easier to replicate than a volatility swap because variance is additive and links directly to option gamma.

$$\boxed{V_{\mathrm{var}}(T) = L_{\mathrm{var}}(\sigma_R^2 - K^2)} \quad = \text{Payoff at Maturity}$$

- L_{var}: variance notional
- K: strike volatility, so K^2 is the strike **variance**
- σ_R^2: realized variance over the contract life

Realized variance. For n trading days, $\boxed{\sigma_R^2 = \frac{252}{n}\sum_{t=1}^{n}\left(\ln\frac{P_t}{P_{t-1}}\right)^2}$ Annualized variance computed from squared log returns. Variance is used because it adds naturally over time and is easier to hedge and replicate.

Variance swap payoff convexity:

$$\sigma_R^2 - K^2 = (\sigma_R - K)^2 + 2K(\sigma_R - K) \approx 2K(\sigma_R - K)$$

near $\sigma_R \approx K$. Hence $\boxed{L_{\mathrm{var}} = \frac{L_{\mathrm{vol}}}{2K}}$ matches a volatility swap locally around the strike.

So a variance swap is the first-order linearization of a volatility swap around strike volatility K. This is why variance swaps are more liquid and often used to proxy volatility exposure.

⋆ **Fair variance strike.** The **fair variance strike** is the strike K^2 that makes the variance swap have zero value at inception: $\boxed{K_{\mathrm{fair}}^2 = E_0[\sigma_R^2]}$

At inception, the fixed strike must equal the risk-neutral expected realized variance, otherwise one side would pay an upfront premium.

Rule-of-thumb fair variance. For short maturity and approximately linear skew,

$$\boxed{K^2 \approx \sigma_{\mathrm{ATMF}}^2\left(1 + 3T \cdot \mathrm{skew}^2\right)} \quad \text{where} \quad \boxed{\mathrm{skew} = \frac{90\mathrm{IV} - 100\mathrm{IV}}{10}}$$

σ_{ATMF} is the at-the-money-forward implied volatility.
Because variance payoff is convex in volatility, downside skew and tail risk push fair variance above squared ATM volatility.

⋆ **Value of a variance swap at time t.** $\boxed{V(t) = e^{-r(T-t)} L_{\text{var}} \left(E_t[\sigma_R^2] - K^2 \right)}$

Expected total variance splits into realized and future parts:

$$\boxed{E_t[\sigma_R^2] = \frac{t}{T}\,\text{Real\_Var}(0,t) + \frac{T-t}{T}\,\text{Fair\_Var}(t,T)}$$

$$\text{so} \quad \boxed{V(0,t,T) = e^{-r(T-t)} L_{\text{var}} \left(\frac{t}{T}\text{Real\_Var}(0,t) + \frac{T-t}{T}\text{Fair\_Var}(t,T) - K^2 \right)}$$

Past realized variance is already locked in; only the remaining future variance is uncertain. So mark-to-market value is determined by realized variance so far plus fair remaining variance.

- $\text{Real\_Var}(0,t)$: realized variance accumulated up to time t
- $\text{Fair\_Var}(t,T)$: fair variance for the remaining life (t,T)

⋆ **Variance swap vega: Sensitivity to Fair Variance**

$$\boxed{\text{Vega} = \frac{\partial V(t)}{\partial \text{Fair\_Vol}} = \frac{\partial V(t)}{\partial \text{Fair\_Var}} \frac{\partial \text{Fair\_Var}}{\partial \text{Fair\_Vol}} = 2\frac{T-t}{T} e^{-r(T-t)} L_{\text{var}} \cdot \text{Fair\_Vol}(t,T)}$$

At $t=0$, $\boxed{\text{Vega} = 2e^{-rT} L_{\text{var}} \cdot \text{Fair\_Vol}(0,T)}$ and if $L_{\text{var}} = \dfrac{L_{\text{vol}}}{2K}$, then at inception $\boxed{\text{Vega} = e^{-rT} L_{\text{vol}}}$

When $t \to T \implies \boxed{\text{Vega} \to 0}$ As maturity approaches, there is less future variance left to move, so sensitivity to implied/fair variance disappears.

⋆ **Realized Variance and Gamma:** $\boxed{RV_T \propto \int_0^T \left(\frac{dS}{S} \right)^2}$

For a delta-hedged option position, $d\Pi \approx \frac{1}{2}\Gamma(dS)^2 = \frac{1}{2}\Gamma S^2 \left(\frac{dS}{S}\right)^2 = \frac{1}{2}\Gamma_\$ \left(\frac{dS}{S}\right)^2$

where $\Gamma_\$ = \Gamma S^2$ is **dollar gamma**.

If you can build a portfolio with approximately constant dollar gamma, the accumulated delta-hedged P&L becomes proportional to realized variance. This is the core replication idea behind variance swaps.

Variance Swap Derivation Replication weights across strikes use a strip of out-of-the-money European puts and calls with weights proportional to $\frac{1}{K^2}$ across strikes. This weighting makes the aggregate option portfolio have approximately constant gamma over a wide range of spot levels, which is what is needed to replicate variance exposure.

$$\frac{dS}{S} = (r-q)dt + \sigma dz, \qquad d\ln S = \left(r - q - \frac{\sigma^2}{2} \right) dt + \sigma dz,$$

subtracting gives $\boxed{\frac{\sigma^2}{2} dt = \frac{dS}{S} - d\ln S}$ averaging over $[0,T]$, $\boxed{\bar{V} = \frac{2}{T}\int_0^T \frac{\sigma^2}{2}\,dt = \frac{2}{T}\int_0^T \frac{dS}{S} - \frac{2}{T}\ln\left(\frac{S_T}{S_0}\right)}$

Taking expectations, $\boxed{E[\bar{V}] = \frac{2}{T}\ln\left(\frac{F_0}{S_0}\right) - \frac{2}{T}E\left[\ln\left(\frac{S_T}{S_0}\right)\right]}$ This converts expected realized variance into a log-payoff expectation. The log contract is then statically replicated using a strip of vanilla options.

Option-integral representation of the log payoff. For any pivot strike S^*,

$$\int_0^{S^*} \frac{1}{K^2} \max(K - S_T, 0)\, dK = \ln\left(\frac{S^*}{S_T}\right) + \frac{S_T}{S^*} - 1$$

$$\int_{S^*}^{\infty} \frac{1}{K^2} \max(S_T - K, 0)\, dK = \ln\left(\frac{S^*}{S_T}\right) + \frac{S_T}{S^*} - 1$$

so $$\boxed{E\left[\ln\left(\frac{S_T}{S^*}\right)\right] = \frac{F_0}{S^*} - 1 - \int_0^{S^*} \frac{e^{rT} P(K)}{K^2}\, dK - \int_{S^*}^{\infty} \frac{e^{rT} C(K)}{K^2}\, dK}$$

where $P(K)$ and $C(K)$ are market European put and call prices with maturity T.
This is the static-replication result: the expected log payoff can be synthesized from a continuum of OTM puts and calls.

⋆ **Fair variance from option prices.** Combining the log-contract identity with the option-integral representation gives the general replication formula:

$$\boxed{E[\bar{V}] = \frac{2}{T} \ln\left(\frac{F_0}{S^*}\right) - \frac{2}{T}\left(\frac{F_0}{S^*} - 1\right) + \frac{2}{T}\left(\int_0^{S^*} \frac{e^{rT} P(K)}{K^2}\, dK + \int_{S^*}^{\infty} \frac{e^{rT} C(K)}{K^2}\, dK\right)}$$

where S^* is typically chosen as the first strike below the forward.

So fair variance is obtained from a strip of OTM options, plus a forward/cash adjustment term. This is the practical pricing formula behind model-free variance swap replication.

Contract value from fair variance: $\boxed{V_0 = L_{\text{var}} e^{-rT}\left(E[\bar{V}] - K^2\right)}$

empty

4.6 EXTRA:

⋆ **Volatility Types (clean map).** *Same symbol σ is used for different objects. The main difference is whether the volatility is observed from market prices, estimated from returns, or inferred as a building block from variance additivity.*

- **Implied Volatility (option-level).** The value of σ that makes a pricing model (usually Black–Scholes) match the observed **option market price** for a specific (K, T); obtained by inverting the pricing formula.

$$\boxed{C_{\text{mkt}} = BS(S_0, K, r, T, q, \sigma_{\text{imp}})} \qquad \Rightarrow \qquad \sigma_{\text{imp}} = \sigma_{\text{imp}}(K, T)$$

- **Volatility Surface.** The full two-dimensional map of **implied volatility** across **strike** and **maturity**; obtained by computing implied vol for many traded options.

$$\boxed{\sigma_{\text{imp}} = \sigma_{\text{imp}}(K, T)}$$

- **Volatility Term Structure (term vol).** The maturity slice of implied volatility as a function of **maturity** only (often at fixed strike or ATM); obtained from option prices across expiries.

$$\boxed{T \mapsto \sigma(T)}$$

This is the curve of "1M vol, 3M vol, 1Y vol, ...".

- **Spot Volatility (instantaneous / current vol).** The volatility at a point in time (or the current instantaneous variance rate in a model); usually model-implied / calibrated, not directly quoted as an option price.

$$\boxed{\sigma_t}$$

In stochastic/local vol models, this is the local-in-time volatility input.

- **Instantaneous Variance.** The square of spot/instantaneous volatility; useful because **variance** (not volatility) adds over time.

$$\boxed{v_t = \sigma_t^2}$$

- **Average (Horizon) Volatility over $[0, T]$.** The variance-averaged volatility over a horizon T; obtained by integrating (or summing) instantaneous/step variances and taking a square root.

$$\boxed{\sigma(0,T) = \sqrt{\frac{1}{T}\int_0^T \sigma_t^2\,dt}}$$

This is the continuous-time term volatility corresponding to maturity T.

- **Discrete Step Volatility / Bucket Volatility.** A volatility assigned to a time interval (bucket), such as year 1 or years 2–3; usually inferred from term vols using variance differences or calibrated in a piecewise-constant model.

$$\boxed{\sigma_1, \sigma_2, \sigma_3, \dots}$$

$$\boxed{\bar{\sigma} = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \sigma_i^2 \Delta_i}}$$

- **Forward Variance.** The market-implied variance expected for a future period (e.g., from T_1 to T_2); obtained from two term vols using variance additivity.

$$\boxed{\sigma^2_{[T_1,T_2]} = \frac{\sigma(0,T_2)^2 T_2 - \sigma(0,T_1)^2 T_1}{T_2 - T_1}}$$

This is one of the most important formulas for extracting future-period variance from quoted term vols.

- **Forward Volatility.** The square root of **forward variance** over a future interval; obtained from the previous formula.

$$\boxed{\sigma_{[T_1,T_2]} = \sqrt{\frac{\sigma(0,T_2)^2 T_2 - \sigma(0,T_1)^2 T_1}{T_2 - T_1}}}$$

- **ATM Implied Volatility.** Implied volatility for an **at-the-money** option at a given maturity; obtained from ATM option quotes and often used to represent the term structure.

$$\boxed{\sigma_{\text{ATM}}(T)}$$

- **Skew (Smile) Volatility.** The variation of implied volatility across strikes for a fixed maturity; obtained from option prices across strikes at the same T.

$$\boxed{K \mapsto \sigma_{\text{imp}}(K,T) \quad (\text{fixed } T)}$$

- **Realized (Historical) Volatility.** Volatility estimated from observed past returns over a sample window; obtained from time series data, not from option prices.

$$\boxed{\hat{\sigma}_{\text{real}} = \sqrt{\frac{A}{n-1}\sum_{i=1}^{n}(r_i - \bar{r})^2}}$$

  A is the annualization factor (e.g., 252 for daily returns).

- **Historical (Rolling) Volatility.** Same idea as realized volatility, but recomputed over a moving window (e.g., 20-day rolling vol); obtained from rolling return samples.

$$\boxed{\hat{\sigma}_{t,\text{roll}} = \text{StdDev}(r_{t-m+1}, \dots, r_t)\sqrt{A}}$$

- **Local Volatility.** A model-implied deterministic volatility surface in state-time space that depends on both current asset level and time; calibrated to match the observed implied vol surface.

$$\boxed{\sigma_{\text{loc}} = \sigma_{\text{loc}}(S, t)}$$

  It is not the same object as $\sigma_{\text{imp}}(K, T)$, even though both are called "volatility surface."

- **Stochastic Volatility (latent vol).** A model where volatility itself evolves randomly over time; parameters are estimated/calibrated from returns and/or option prices.

$$\boxed{dS_t = \mu S_t dt + \sqrt{v_t} S_t \, dW_t, \qquad dv_t = \cdots}$$

  Examples: Heston-type models.

⋆ **What is quoted vs inferred vs estimated?**

- **Quoted / backed out from options:** $\sigma_{\text{imp}}(K, T)$, $\sigma_{\text{ATM}}(T)$, skew, term structure
- **Inferred from term vols (variance math):** forward variance, forward vol, bucket/step vols
- **Estimated from return data:** realized/historical volatility
- **Model objects (calibrated):** spot/instantaneous vol σ_t, local vol $\sigma_{\text{loc}}(S, t)$, stochastic variance v_t

If you remember one thing: ***variance adds over time****, so many "forward" and "bucket" volatilities are extracted by working in variance space first, then taking square roots.*

Chapter 5

Interest Rate Derivatives and Models

5.1 Interest Rate Derivatives

5.1.1 3-Month SOFR Futures

⋆ **3-Month SOFR Futures.** A **SOFR future** is an exchange-traded contract that locks in today the **compounded SOFR rate** that will be realized over a future 3-month accrual window, with daily margin replacing a single forward settlement, which is what makes it unique, namely a forward on a *backward-looking* rate (the rate is only known at the end of the period, not the start, unlike LIBOR which was set up front).
The phrase "compounded SOFR over $[T, T+\delta]$*" means daily overnight SOFR fixings inside the window are chained together via* $(1 + r_1/360)(1 + r_2/360)\cdots$ *to produce one annualized 3-month rate* R*. So the contract pays off on a rate that is built from observed overnights, not quoted in advance, which is the structural shift from the legacy Eurodollar/LIBOR design.*

- **Long SOFR Future:** You profit when rates fall $\rightarrow$ price rises
(Imagine it as mimicking a lender; they benefit by locking in rates before they drop)
- **Short SOFR Future:** You profit when rates rise $\rightarrow$ price falls
(Imagine it as mimicking a borrower who locked in a low rate and now they went up)

Contract specifications.

- **Notional** $L = \$1{,}000{,}000$ Principal on which interest is conceptually accrued, fixed by the exchange. Never changes hands, only sets the dollar scale.
- **Accrual period** $\delta = 0.25$ **years.** Time period during which interest has built up but hasn't been paid yet. Fixed because the underlying is the 3-month rate. The same δ scales rate moves into dollars in caplet and floorlet payoffs later.
- **Quoted price** $Q = 100 - r$, with r **the annualized compounded 3-month SOFR** (in percent). Higher Q, lower rate. A convention that lets the contract trade in price space.
- **DV01** $= L \cdot \delta \cdot 0.0001 = \$25$. Dollar P&L per 1 bp move in the rate. The general rule is **notional** $\times$ **accrual** $\times$ **rate move**, which appears on every rate product.
- **P&L per bp of** $Q = \$25$, since $\Delta Q = -\Delta r$ point-for-point.

Pricing logic and payoff. A futures contract has **zero value at inception**, the quote is a reference level for margin, not a payment.
Before expiry the quote tracks the futures-implied forward rate, $\boxed{Q_t = 100 - F_t^{\text{fut}}}$
which (ignoring convexity) equals the forward rate for the accrual window $[T, T+\delta]$. Entering a long

therefore locks in the current forward r_0.
At expiry the contract settles to $Q_T = 100 - R$, with R the realized compounded 3-month SOFR.

Cumulative P&L on a long held to expiry: $\boxed{\text{P\&L}_{\text{long}} = L\,\delta\,(r_0 - R) \ = \$25 \cdot \Delta Q_{\text{bps}}}$
with $\Delta Q = Q_T - Q_0$ in bps.

Same identity in two units. The long gains when realized rate falls below the locked-in r_0, scaled by $L\,\delta$ to convert percent into dollars. No upfront price because it is a futures, the P&L accumulates as daily variation margin.

Trade lifecycle (procedure).

1. **Enter** long at $Q_0 = 100 - r_0$, locking in forward rate r_0.
2. **Mark to market** daily, \$25 of variation margin per bp move in Q.
3. **At expiry**, realized rate is $r_T = 100 - Q_T$, total cash flow $\$25 \cdot (Q_T - Q_0)$ in bps. Positive if rates fell.

Equivalent to a cash-settled FRA, with the payoff streamed as margin rather than paid in one shot at expiry. The full SOFR-future trade replicates a forward rate agreement settled in cash, with the gain or loss appearing as accumulated margin variation rather than as a single payment at expiry.

5.1.2 Standard Market Model

The market convention for valuing interest-rate options. It is an **interest-rate version of Black's model**, assuming the relevant **forward price is log-normal** and pricing the option from the forward price together with a **forward volatility**. The valuation formula maps directly onto the Black formula for futures options.

General Black Formula
For a European option with strike K, time to maturity T, forward price F, volatility σ, and discount factor $P(0,T)$:

$$c = P(0,T)\Big[F\,N(d_1) - K\,N(d_2)\Big] \quad \textit{and} \quad p = P(0,T)\Big[K\,N(-d_2) - F\,N(-d_1)$$

where

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

Mapping Black's Model Across Products

Product	**Forward** (F)	**Volatility** (σ)	**Discount Factor**
Bond Option	Forward bond price F_B	Bond price vol σ_B	$P(0, T_{\text{opt}})$
Caplet/Floorlet	Forward SOFR rate F_k	Caplet vol σ_t	$L\,\delta\,P(0, t{+}\delta)$
Swaption	Forward swap rate s_F	Swaption vol σ	$L \cdot A$

Market-quoted volatilities may be expressed as *yield volatility* (for bonds), *caplet/cap volatility* (for caps), or *swaption volatility* (for swaptions).

5.1.3 Risk-Free Bond Options

⋆ **Risk-Free Bond Options.** European options on default-free bonds (e.g. Treasuries), priced with Black's model using the **forward bond price** as the underlying. They appear far more often *embedded* inside fixed-income contracts than as standalone derivatives, which is the main reason this pricing matters in practice.

Type	Option holder	Exercised when	Payoff
Callable bond	Issuer	Rates fall, $B > K$	$\max(B - K, 0)$
Puttable bond	Holder	Rates rise, $B < K$	$\max(K - B, 0)$
Prepayment option	Borrower	Rates fall, $B > K$	$\max(B - K, 0)$

All three are bond options in disguise: the bond payoff is capped or floored by an option whose strike is the call/put price. Pricing reduces to pricing the underlying European bond option with Black, then attaching it to the bond cash flows.

Volatility convention. Markets quote bond option volatility as **yield volatility** σ_y, not price volatility. Black's model expects **price volatility** σ_B, so a conversion step is required.

Why yield vol is the market quote: yields are stationary and roughly mean-reverting, so one yield vol number is comparable across maturities. Price vol scales with duration and is maturity-specific, less convenient as a quoting unit.

⋆ **European Bond Option Black Formula.** Treat the forward bond price F_B as the log-normal underlying, then plug into Black.

- F_B: forward bond price (dirty, includes accrued interest)
- $\sigma_B \equiv \sigma_{F_B}$: forward bond price volatility, two interchangeable names for the same quantity
- K: strike on the dirty forward bond price
- T: option maturity, distinct from the bond's own maturity

$$\boxed{c = P(0,T)\Big[F_B\, N(d_1) - K\, N(d_2)\Big]} \qquad \boxed{p = P(0,T)\Big[K\, N(-d_2) - F_B\, N(-d_1)\Big]}$$

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

Same structure as Black for futures options. The product-specific work happens before this formula: **building F_B from the spot bond and converting $\sigma_y \to \sigma_B$.**

Output units. Option prices c, p come out in the **same dollar units** as F_B and K. Plug face-value dollars in, get face-value dollars out (e.g. \$9.49 per \$1,000 face). Plug points in (% of face), get points out.

- σ_B should be **larger** than σ_y for any bond with $D > 1$, since duration amplifies yield moves into price moves. If converted vol comes out smaller than quoted yield vol, the conversion is wrong.
- K must be on the **dirty** forward price (with accrued interest). Bloomberg quotes are usually clean, add accrued interest at T before using as K.
- Pre-expiry coupons must be stripped from B_0 when building F_B, otherwise F_B is inflated and the call price biased upward.

⋆ **Forward Bond Price.** Price agreed today for delivery of the bond at option maturity T:

$$\boxed{F_B = \frac{B_0 \; - \; PV(\text{Coupons})}{P(0,T)}}$$

where B_0 is the current dirty bond price and $PV(\text{Coupons})$ sums the present value of every coupon paid **before** T.

Standard cost-of-carry, current price minus carry income, grossed up to time T. Coupons paid before expiry belong to the current holder so they are stripped out: the forward delivers a bond that has already shed those coupons.

⋆ **Price vs Yield Volatility.**
Converting between vol units uses the duration approximation $\Delta B/B \approx -D\,\Delta y$.

- D: duration of the forward bond at option maturity
- $y_{F,0}$: the yield that discounts the forward bond's remaining cash flows back to F_B
- σ_y: forward yield volatility, quoted as log-yield vol $\sigma_y = \text{vol}(\Delta y/y)$
- σ_{F_B}: forward price volatility (the input Black needs)

Conversion.

$$\frac{\Delta F_B}{F_B} \approx -D\,\Delta y_F = -D\,y_{F,0}\,\frac{\Delta y_F}{y_{F,0}} \qquad \Longrightarrow \qquad \boxed{\sigma_{F_B} = D\,y_{F,0}\,\sigma_y}$$

Each factor has a job: D **amplifies a yield move into a price move,** $y_{F,0}$ **rescales the log-yield vol back into absolute yield vol, and the product is price vol.** Long-duration bonds at high yields are the most price-sensitive to yield shocks, exactly what the formula reflects.

⋆ **Bond Option Pricing Procedure.** Given a market yield vol and bond market data, the standard workflow has five steps, each fixing one input Black needs:

1. **Convert vol units**: $\sigma_{F_B} = D\,y_{F,0}\,\sigma_y$. Without this, Black is fed the wrong volatility.
2. **Discount factor**: $P(0,T) = e^{-rT}$, evaluated at the option's expiry T.
3. **Strip pre-expiry coupons**: compute $PV(\text{Coupons})$ for every coupon paid before T. These do not travel into the forward.
4. **Forward bond price**: $F_B = (B_0 - PV(\text{Coupons}))/P(0,T)$, the underlying for Black.
5. **Apply Black**: compute d_1, d_2, then c and p from the boxed formulas above.

The Black formula at the end is mechanical. The real work is preparation: convert volatility, strip carry, build the forward. Skipping the coupon strip inflates F_B and biases the call price upward; skipping the vol conversion just uses the wrong σ entirely. Always check both before plugging in.

Call vs put symmetry. For a bond trading below par with strike at par, the put is deep in-the-money at the forward and the call is out-of-the-money, so $p \gg c$. This is the typical pattern for a discount bond and a useful check on numerical output.

5.1.4 Caps and Floors

⋆ **Caplet.** European **call option on a future SOFR rate**, used to cap interest paid on a floating rate loan over one accrual window $[t,\, t+\delta]$. **Rate is observed at t, payoff settles at $t+\delta$:**

$$\boxed{\text{Caplet}(t) \;=\; \underbrace{L}_{\text{notional}} \cdot \underbrace{\delta}_{\substack{\text{accrual}\\ \text{fraction}}} \cdot \max(\underbrace{F_{t,t+\delta} - K}_{\substack{\text{rate excess}\\ \text{over strike}}},\; 0)}$$

⋆ **Floorlet.** European **put option on a future SOFR rate**, used to floor income received on a floating-rate investment over one accrual window $[t,\, t+\delta]$. Rate is observed at t, payoff settles at $t+\delta$. **Rate is observed at t, payoff settles at $t+\delta$:**

$$\boxed{\text{Floorlet}(t) \;=\; \underbrace{L}_{\text{notional}} \cdot \underbrace{\delta}_{\substack{\text{accrual}\\ \text{fraction}}} \cdot \max(\underbrace{K - F_{t,t+\delta}}_{\substack{\text{rate shortfall}\\ \text{below strike}}},\; 0)}$$

- L: notional principal
- δ: SOFR tenor (3M standard, $\delta = 0.25$)
- t: option maturity, when rate is observed
- $t+\delta$: payment date, when payoff is settled
- $F_{t,t+\delta}$: SOFR forward rate at time t for accrual $[t,\, t+\delta]$
- K: strike rate

Same call/put structure as equity options, but the underlying is a rate. The factor $L\delta$ converts the rate move (in percent) into dollar interest on the notional. Caplet protects a borrower against rate spikes, floorlet protects a lender against rate drops.

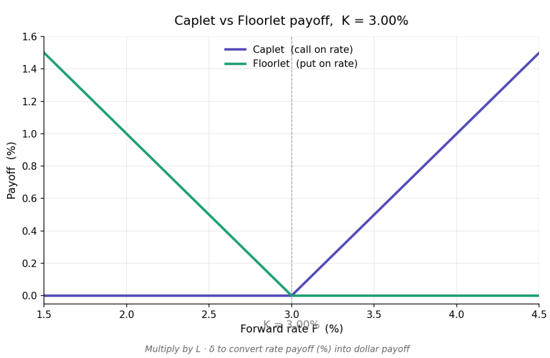

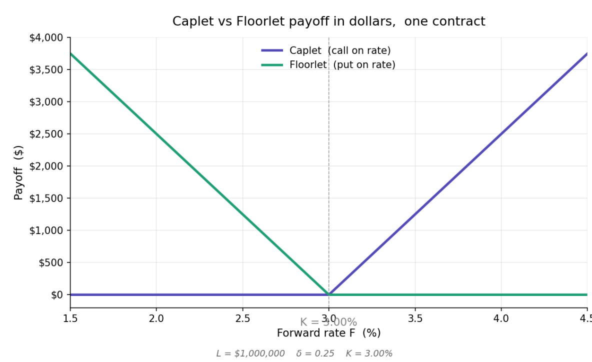

Same payoff, two views. Left axis is the rate move $F - K$ in %, contract-size independent. Right axis is the dollar payoff for one standard contract ($L = \$1\text{M}$, $\delta = 0.25$, $K = 3\%$), so the slope past the strike is $L \cdot \delta = \$2{,}500$ per 1% rate move, exactly the same DV01 scaling as the SOFR future. Multiplying the left chart by $L \cdot \delta$ produces the right chart.

⋆ **Caplet/Floorlet Valuation (Black's Model).** Treat $F_{t,t+\delta}$ as the log-normal underlying. Same Black structure, with discount factor $L\,\delta\,P(0, t+\delta)$.

$$\boxed{\text{Caplet} = L\,\delta\,P(0,t+\delta)\Big[F_{t,t+\delta}\,N(d_1) - K\,N(d_2)\Big]}$$

$$\boxed{\text{Floorlet} = L\,\delta\,P(0,t+\delta)\Big[K\,N(-d_2) - F_{t,t+\delta}\,N(-d_1)\Big]}$$

$$d_1 = \frac{\ln(F_{t,t+\delta}/K) + \frac{1}{2}\sigma_t^2\,t}{\sigma_t\sqrt{t}}, \qquad d_2 = d_1 - \sigma_t\sqrt{t}$$

- σ_t: **Black implied volatility** for caplet/floorlet expiring at t. Same role as equity IV: the single number that reproduces the market price when plugged into Black.
- $P(0, t+\delta)$: discount factor to the **payment** date, not to option expiry. Payoff is determined at t but cash settles at $t+\delta$, so dollars get pulled back from $t+\delta$.

Why discount to $t+\delta$ and not t: payoff is determined at t (rate fixes) but cash settles at $t+\delta$. The $L\delta$ scales rate-to-dollars, $P(0, t+\delta)$ moves dollars from payment date back to today.

SOFR backward-looking adjustment. Compounded SOFR keeps fixing through the accrual period, so in practice replace the option maturity t in d_1, d_2 with the midpoint $(t + (t+\delta))/2$ as a small convexity correction.

$\star$ **Caplet as a ZCB Put Option.** A caplet can be rewritten exactly as a European put on a zero-coupon bond. Critical because tree and short-rate models price ZCB options natively; this equivalence avoids needing a separate caplet pricer.

Algebraic rewrite. Discount the caplet payoff back from $t+\delta$ to t (rate R_t already known at t):

$$\frac{\delta L}{1+\delta R_t}\max(R_t - K,\, 0) \;=\; L(1+\delta K)\max\!\left(\frac{1}{1+\delta K} - \frac{1}{1+\delta R_t},\, 0\right)$$

Resulting European put bond option

- Notional: $L(1+\delta K)$
- Strike: $1/(1+\delta K)$, the price of a δ-period ZCB at rate K
- Option maturity: t
- Underlying: ZCB maturing at $t+\delta$, current price $1/(1+\delta R_t)$

When rates rise, ZCB prices fall, so a put on a ZCB pays off, exactly when a caplet pays off. Same exposure, different wrapping. This lets Hull-White, Vasicek, and tree models price caplets through their ZCB-option pricers without any caplet-specific code.

$\star$ **Cap and Floor.** A **cap** is a portfolio of caplets, a **floor** is a portfolio of floorlets, each strung together over consecutive accrual periods. Both are valued as the sum of individual caplet/floorlet values, since each piece is an independent option.

Standard market conventions.

- Quarterly payments on 3-month SOFR
- A T-year quarterly cap has $4T-1$ caplets: the $t=0$ fixing is already known at inception so it carries no optionality
- Each caplet uses the forward SOFR rate for its own accrual window

Pricing reduces to: price each caplet with Black, sum.

Cap Pricing Procedure.

1. For each caplet i, read off the forward rate F_i and discount factor $P(0,\, t_i+\delta)$ from the curve
2. Pull caplet vol σ_{t_i} from market quotes (or strip from cap vols)
3. Compute d_1, d_2 and apply Black to get one caplet value
4. **Sum: Cap = $\sum_i$ Caplet$_i$**

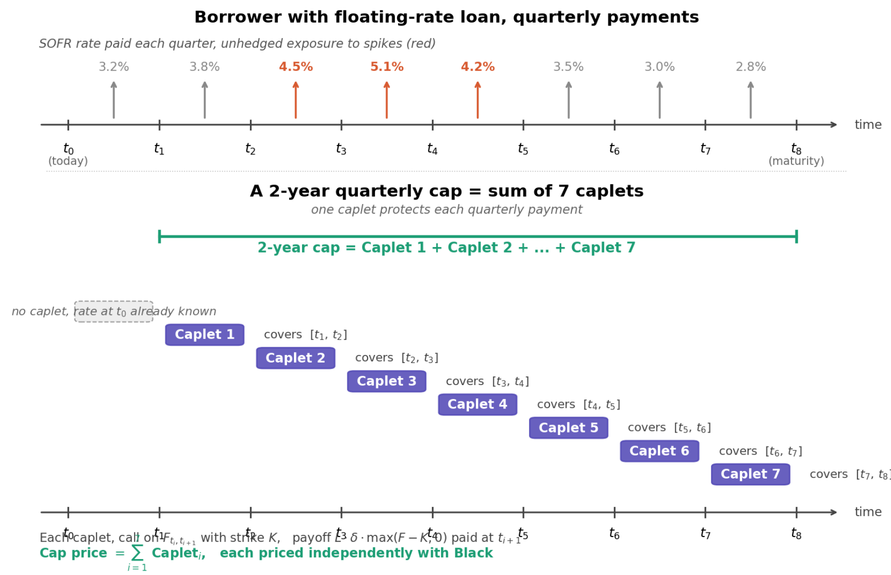

Top timeline: ***realized*** *SOFR as time unfolds, not forecasts at* t_0*. Today we only know the first rate (already fixed) plus a curve of* ***forward rates*** $F_{t_i,t_{i+1}}$ *for the rest. Each caplet is a call written today on one future* $F_{t_i,t_{i+1}}$*, observable at* t_i *and paid at* t_{i+1}*. The cap bundles all of them: one premium today, protection on every quarterly fixing.*

First period has no caplet because the rate for $[t_0, t_1]$ *is already fixed at* t_0*, nothing random to insure. Hence a* T*-year quarterly cap has* $4T - 1$ *caplets, not* $4T$*.*

⋆ **Cap Volatility vs Caplet Volatility.** Two quoting conventions for the same total cap price:

- **Caplet vol** σ_t: a different vol for each caplet, the term structure of caplet vols
- **Cap vol** $\bar{\sigma}_T$: a single **flat vol** that, applied uniformly to all caplets, reproduces the same cap price

$$\boxed{\text{Cap}(\bar{\sigma}_T) \;=\; \sum_{t=2}^{4T} \text{Caplet}(\sigma_t) \;=\; \sum_{t=2}^{4T} \text{Caplet}(\bar{\sigma}_T)}$$

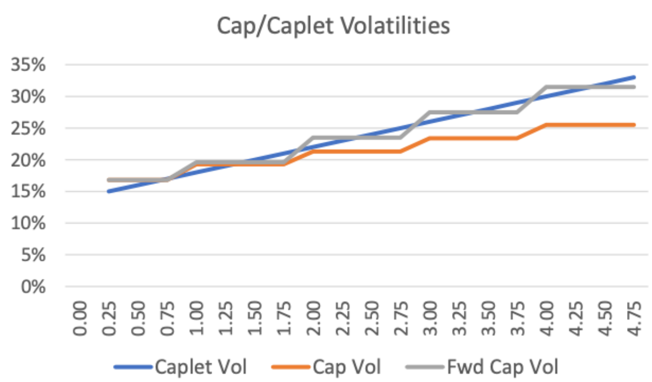

- **Cap vol**: convenient one-number quote per maturity, what Bloomberg surfaces show.
- **Caplet vol**: actual input to Black, must be **stripped** (bootstrapped) from cap vols for individual caplet pricing or exotic structures.
- **Forward cap vol**: vol of the marginal year between two consecutive cap maturities, recovered from differences of cap variances.

Same idea as zero rates vs forward rates: cap vol is the average over $[0, T]$*, caplet vol is the local rate, forward cap vol is the slope. Markets quote cap vol because it is one number; pricing exotics requires the full caplet vol curve.*

⋆ **Put-Call Parity:** Same-strike, same-maturity, same-payment-date caplet, floorlet, and FRA satisfy:

$$\boxed{\text{Caplet} - \text{Floorlet} \;=\; \text{FRA}}$$

FRA Exchanges the **realized SOFR for a fixed rate** K over the accrual window $[t, t+\delta]$, settled at $t + \delta$. Same payoff as a swap leg, but for one period only. $\text{FRA}(t) \;=\; L\,\delta\,(F_{t,t+\delta} - K)$

No optionality, the holder takes the rate move in either direction. Equivalent to a SOFR future on the same accrual window (ignoring convexity). This is exactly the linear payoff that sits between a caplet and a floorlet, which is why $\textit{Caplet} - \textit{Floorlet} = \textit{FRA}$*.*

Proof. Apply $\max(x, 0) - \max(-x, 0) = x$ to $x = \text{SOFR} - K$:

$$L\,\delta\,[\max(\text{SOFR} - K,\,0) - \max(K - \text{SOFR},\,0)] \;=\; L\,\delta\,(\text{SOFR} - K)$$

Aggregated version. $\boxed{\text{Cap} - \text{Floor} \;\approx\; \text{Swap}}$

ignoring the first reset which is already known at inception (no optionality there).

Swap (Interest Rate Swap). Exchanges **floating SOFR for a fixed rate** K on every accrual period over the swap's life, settled at the end of each period.
A strip of FRAs back-to-back. Swap $= \sum_{i=1}^{n} L\,\delta\,(F_{t_i,t_{i+1}} - K) \;=\; \sum_{i=1}^{n} \text{FRA}_i$

Implication: any two of cap, floor, swap pin the third by no-arbitrage. Traders use this to check vol surface consistency, an inconsistent cap/floor vol would imply free money against the swap rate.

5.1.5 Interest Rate Collar.

⋆ **Interest Rate Collar.** A borrower's hedging strategy that traps the effective borrowing rate inside a corridor $[K_F,\,K_C]$ by combining two opposite options. **Long cap** at strike K_C **plus short floor** at strike K_F, with $K_F < K_C$.

$$\boxed{\text{Collar} = \text{long cap}(K_C) + \text{short floor}(K_F)}$$

Position breakdown. Borrower paying floating SOFR rate S on each accrual:

- **Long cap at K_C**: pays the borrower whenever S exceeds K_C, neutralizing rate spikes above the upper strike. This is the actual hedge, the protective leg.
- **Short floor at K_F**: borrower owes money whenever S falls below K_F, giving up the windfall of unusually low rates. This is the financing leg, sold to pay for the cap.

Core idea: ***give up gains below K_F to fund protection above K_C****, trapping the effective borrowing rate inside the corridor $[K_F,\,K_C]$ regardless of where SOFR moves. A standalone cap costs an upfront premium smaller corporates often can't justify; selling a floor brings premium in to offset it. Pick K_F low enough that you don't expect rates to fall there and the cap is effectively "free", trading a wide range of outcomes for a narrow predictable one.*

Per-period cash flow breakdown.

Regime	Cap payoff	Floor payoff	Borrow	Net (Effective rate)
SOFR $< K_F$	0	$-L\,\delta\,(K_F - \text{SOFR})$	$-L\,\delta\,\text{SOFR}$	$-L\,\delta\,K_F$
$K_F \le \text{SOFR} \le K_C$	0	0	$-L\,\delta\,\text{SOFR}$	$-L\,\delta\,\text{SOFR}$
SOFR $> K_C$	$L\,\delta\,(\text{SOFR} - K_C)$	0	$-L\,\delta\,\text{SOFR}$	$-L\,\delta\,K_C$

Borrower trades upside hedging cost for giving up downside benefit. Selling the floor finances the cap. A ***zero-cost collar*** *chooses K_F, K_C such that the floor premium exactly pays for the cap, very common in corporate hedging.*

Effective borrowing rate. The borrower's net payment per period, clamped inside the corridor:

$$\boxed{\text{Effective rate} \;=\; \min\Big\{K_C,\; \max\{K_F,\; \text{SOFR}\}\Big\}}$$

Within $[K_F,\,K_C]$ the borrower simply pays SOFR; outside, the collar kicks in and pins the rate to whichever strike was breached.

5.1.6 Swaptions

⋆ **Swaption (Swap Option).** An option granting the holder the right to enter into an interest rate swap at a specified future date. **Payoff at maturity is the PV of the underlying swap, written as Annuity × (Swap rate – Strike).**

- **Payer swaption**: right to **pay** a specified fixed swap rate K and **receive SOFR compounded in arrears**. ITM when rates rise (paying K beats the new market rate).
- **Receiver swaption**: right to **receive** the specified fixed rate K and **pay SOFR compounded in arrears**. ITM when rates fall.

Same payer/receiver naming as in swaps. The "compounded in arrears" detail matters because the floating leg references actual realized SOFR (set at the end of each period), the same backward-looking design as 3-month SOFR futures.

Swaption vs Swap payoff at maturity. A **swaption is option-type**, payoff is $\max(\cdot, 0)$. A **swap is forward-type**, payoff linear in either direction. *Same call vs forward distinction as in equity. The exercise decision is one-shot at T_{Opt}: take the swap if it's profitable, walk if not.*

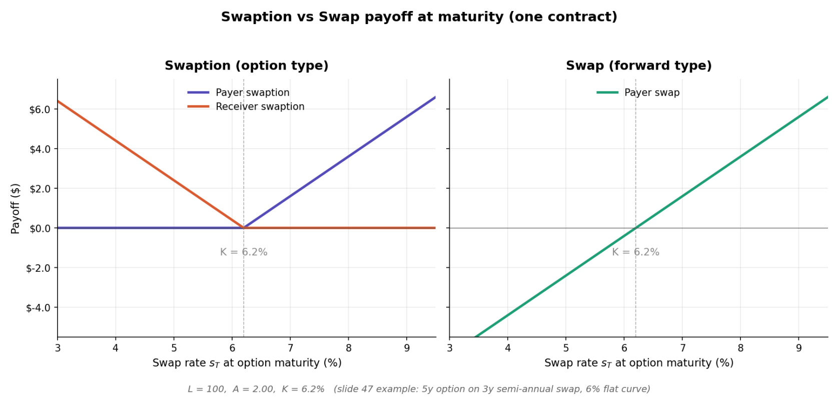

Swap: contract to exchange fixed-rate payments for floating-rate payments VS **Swaption**: option to exchange a fixed-rate bond for a floating-rate bond

Bond option approximations.

- **Pay-fixed (payer) swaption** ≈ **put on the bond** with strike at par. Rates up → bond down → put gains.
- **Receive-fixed (receiver) swaption** ≈ **call on the bond** with strike at par. Rates down → bond up → call gains.

Fixed leg = coupon bond, floating leg = par at every reset, so a swap is short-fixed-bond + long-par. The pay-fixed swaption is then short a coupon bond struck at par. This equivalence lets short-rate models price swaptions through their bond-option pricers.

Swaption Parameters.

- L: notional (principal of underlying swap), in dollars; sets the dollar scale, never changes hands.
- $T_{\mathbf{Opt}}$: option maturity, when the entry decision is made
- T_n: final swap payment date
- **Swap Tenor** $= T_n - T_{\mathbf{Opt}}$: the swap's own life, starting at option maturity
- m: frequency of swap payments per year (e.g., $m = 2$ for semi-annual). Determines how many payment dates the annuity A sums over: $m \cdot$ tenor total.
- **Swap payment dates**: $T_{\text{Opt}} + 1/m$, $T_{\text{Opt}} + 2/m$, …, T_n, all after T_{Opt}
- K: strike, the **fixed rate** written into the swaption (in %). The rate the holder pays (payer) or receives (receiver) if they exercise. Same role as a strike price, but quoted as a rate not a dollar level.
- $s_{T_{\mathbf{Opt}}}$: market swap rate observed at option expiry (random, in %). The par swap rate the market quotes at T_{Opt} for a swap of the underlying tenor; zeros out a fresh ATM swap at that moment.
- s_F: forward swap rate today (in %), for a swap starting at T_{Opt} and maturing at T_n. Today's

curve-implied prediction of $s_{T_{\text{Opt}}}$, the log-normal underlying in Black's swaption formula.

Two "swap rate" quantities to keep distinct: s_F is known today (curve quote), $s_{T_{\text{Opt}}}$ is random (revealed at expiry). The swaption pays off based on $s_{T_{\text{Opt}}}$, but Black prices it today using s_F as the log-normal underlying with vol σ.

And two distinct time scales: T_{Opt} when the option matures, T_n when the underlying swap matures. A "5-year option on a 3-year swap" means $T_{\text{Opt}} = 5$, swap tenor $= 3$, $T_n = 8$.

⋆ **Per-Date Swaption Payoff.** A swaption is **not a single lump-sum option**, nor is it a portfolio of independent options. Once exercised at T_{Opt}, the holder enters a real swap that pays out on **every** subsequent swap-payment date. On each such date, the holder receives:

$$\boxed{\text{Payer payoff} \;=\; \underbrace{\frac{L}{m}}_{\substack{\text{notional}\\\text{per period}}} \cdot \max\big(\underbrace{s_{T_{\text{Opt}}} - K}_{\substack{\text{rate excess}\\\text{over strike}}},\ 0\big)} \qquad \boxed{\text{Receiver payoff} \;=\; \underbrace{\frac{L}{m}}_{\substack{\text{notional}\\\text{per period}}} \cdot \max\big(\underbrace{K - s_{T_{\text{Opt}}}}_{\substack{\text{rate shortfall}\\\text{below strike}}},\ 0\big)}$$

Single decision, stream of payments. The exercise decision happens **once** at T_{Opt}. After that, the rate $s_{T_{\text{Opt}}}$ is locked in and the same number multiplies through every future payment date. Only one source of randomness, the rate observed at T_{Opt} with total swaption value sums the per-date payoffs across all swap-payment dates, discounted to today:

$$\boxed{\text{Swaption value} \;=\; L \cdot \underbrace{\frac{1}{m}\sum_i P(0, T_i)}_{\text{annuity factor } A} \cdot \mathbb{E}\Big[\max(s_{T_{\text{Opt}}} - K,\ 0)\Big]}$$

The annuity A bundles "stream of payments tied to one rate observation" into a single number, leaving Black's formula to handle the expected max. This is the technical device that lets a multi-payment instrument be priced as a single European option on the forward swap rate.

Annuity factor A. Present value of \$1 paid on every swap payment date:

$$\boxed{A \;=\; \frac{1}{m} \sum_{i=1}^{m\cdot(T_n - T_{\text{Opt}})} P\Big(0,\ T_{\text{Opt}} + \frac{i}{m}\Big)}$$

The annuity bundles every payment-date discount into one number. It plays the same role for swaptions that $L\,\delta\,P(0, t{+}\delta)$ played for caplets: a "swap-shaped" discount factor that converts the rate-move payoff into total dollar value across the swap's life.

⋆ **Swaption Valuation (Black's Model).** Treat s_F as the log-normal underlying. Standard Black with discount factor replaced by $L \cdot A$.

$$\boxed{\text{Payer Swaption} \;=\; L \cdot A \cdot \Big[s_F\, N(d_1) - K\, N(d_2)\Big]}$$

$$\boxed{\text{Receiver Swaption} \;=\; L \cdot A \cdot \Big[K\, N(-d_2) - s_F\, N(-d_1)\Big]}$$

$$d_1 = \frac{\ln(s_F/K) + \frac{1}{2}\sigma^2 T_{\text{Opt}}}{\sigma\sqrt{T_{\text{Opt}}}}, \qquad d_2 = d_1 - \sigma\sqrt{T_{\text{Opt}}}$$

- s_F: forward swap rate as of today
- σ: volatility of the forward swap rate (the **swaption vol**)

- T_{Opt}: option maturity (used in d_1, d_2, NOT swap tenor)

Swaption Pricing Procedure.

1. **Annuity factor**: $A = \frac{1}{m}\sum_i P(0, T_{\text{Opt}} + i/m)$ over every swap payment date. Discount-factor-weighted "swap horizon."
2. **Forward swap rate**: s_F from the curve, the par swap rate for a swap starting at T_{Opt} and maturing at T_n.
3. **Compute** d_1, d_2: standard Black with $T = T_{\text{Opt}}$ and the swaption vol σ.
4. **Apply Black**: $L \cdot A \cdot [\ldots]$ for payer or receiver. Bracket sign flips between flavors.

Scale anchor: a 3-year semi-annual swap at a 6% flat curve has annuity $A \approx 2.0$, so an ATM payer swaption with $L = 100$ and $\sigma\sqrt{T_{Opt}} \approx 0.45$ prices to roughly 2.19. The payer is worth about 2.2% of notional, the typical magnitude for ATM swaptions in mid-vol markets.

5.1.7 Negative Rates

⋆ **Why Black Breaks for Negative Rates.** Black assumes F is log-normal, which forces $F > 0$. Several major economies (Japan, Germany, Switzerland, not yet US) have observed negative policy rates, breaking the model directly.

Two standard fixes.

- **Normal (Bachelier) model**: forward rate is normal, F unbounded
- **Shifted Lognormal model**: $F + \alpha$ is log-normal, allows $F > -\alpha$

Both retain analytic tractability close to Black. Normal is the post-2008 default for vol quoting in low/negative rate markets. Shifted Lognormal is the smaller modification, keeps log-normal mechanics intact.

⋆ **Normal (Bachelier) Model. Forward rate follows arithmetic Brownian motion:**

$$\boxed{dF_k = \sigma_k^*\, dz}$$

σ_k^* is the **normal volatility**, quoted in absolute rate units (basis points), not percent.

Caplet/Floorlet pricing.

$$\boxed{\text{Caplet} = L\,\delta_k\,P(0, t_{k+1})\Big[(F_k - K)\,N(d) + \sigma_k^*\sqrt{t_k}\,N'(d)\Big]}$$

$$\boxed{\text{Floorlet} = L\,\delta_k\,P(0, t_{k+1})\Big[(K - F_k)\,N(-d) + \sigma_k^*\sqrt{t_k}\,N'(d)\Big]}$$

$$d = \frac{F_k - K}{\sigma_k^*\sqrt{t_k}}$$

Differences from Black.

- No $\ln(F/K)$ term: d depends on the **difference** $F - K$, not the ratio
- Adds an $N'(d)$ term: extra value contribution from the symmetric normal density
- σ^* has units of rate (e.g., 80 bp); Black's σ is dimensionless (e.g., 25%)

Normal model treats rate moves as additive (1 bp shift regardless of rate level), Black treats them as multiplicative (proportional to rate). Near zero rates, additive is more realistic since percent moves blow up. The $N'(d)$ term is the time-value bump from the normal density at the strike.

Black vs Normal Volatility Mapping. For any fixed (F, K, T), there is a one-to-one bijection between Black vol and Normal vol that produces the same option price. Markets often quote both:

- ATM approximation: $\sigma^* \approx \sigma_{\text{Black}} \cdot F$
- Used to translate between conventions across regions and products

The two vols are alternative coordinates on the same option price. Use Normal in negative-rate or near-zero environments; Black where rates are clearly positive.

⋆ **Shifted Lognormal Model.** Assume $F_k + \alpha$ is log-normal: $\boxed{d(F_k + \alpha) = \sigma\,(F_k + \alpha)\,dz}$

where $\alpha > 0$ is the **shift** or **displacement**, allowing $F_k > -\alpha$.

Black formula adjustment. Just shift both F and K by α:

$$\boxed{\text{Black price}(F_k,\,K) \;\longrightarrow\; \text{Black price}(F_k + \alpha,\,K + \alpha)}$$

Smallest possible modification of Black, no new functional form. Pick α to match the lower bound on rates (e.g., $\alpha = 1\%$ allows rates down to -1%). The shift acts as a translation of the entire log-normal world into negative-rate territory.

Choosing the shift α. α is a **deterministic parameter** fixed before pricing, defines the **left boundary** of the rate's universe, the model lets F_k drift toward $-\alpha$ but never reach it, so $-\alpha$ acts as a hard floor on rates.

- **Policy-driven**: pick $-\alpha$ slightly below the lowest rate considered policy-feasible. ECB rates bottomed near -0.5% so desks used $\alpha \approx 1\%$ to 2%; Swiss rates hit -0.75% with $\alpha \approx 1.5\%$.
- **Calibrated**: treat α as a free parameter, choose the value minimizing pricing error against a strip of market quotes:

$$\alpha^* = \arg\min_{\alpha} \sum_i \Big(\text{Market price}_i - \text{Shifted-Black price}_i(\alpha)\Big)^2$$

  Standard trading-desk practice, α becomes one more knob alongside σ to fit the smile.
- **Convention-driven**: some markets adopt a fixed value by convention so quotes are directly comparable across desks. EUR caps used $\alpha = 3\%$ during the negative-rate era.

5.1.8 Convexity Adjustment

⋆ **Convexity Adjustment.** A correction to the forward rate used when a contract's **payment timing breaks the standard structure**, fixing a small but real bias between the **expected yield** and the **expected bond price** that arises from the nonlinear yield-price relationship.

The standard structure (no adjustment needed). A **δ-maturity rate** observed at time t pays off at time $t + \delta$. Examples: floating legs of swaps, FRAs, caplets. This timing alignment lets us:

- Use the **forward rate** directly to project the cash flow
- Discount at the risk-free rate from the payment date

The non-standard structure (adjustment needed). A contract pays **at time** t the δ-maturity yield observed at t, with no δ-period delay. Example: a contract exchanging the 5-year yield for a fixed rate, settled the day the yield is observed.

The standard structure works because the rate's payment naturally lives inside the accrual window it describes. When payment timing skips that window, the forward rate is no longer the right unbiased estimate of the future rate, the convexity adjustment is the correction.

⋆ **Forward Bond Price vs Forward Yield.** The bond price-yield relationship $B = G(y)$ is **non-linear** (convex). Under the T-forward measure Q^T:

$$\boxed{\text{Forward bond price} \;=\; \mathbb{E}^T[B_T]} \qquad \text{but} \qquad \text{Forward yield} \neq \mathbb{E}^T[y_T]$$

Forward bond prices are tradeable assets and therefore equal expected future prices under Q^T. But yields are derived quantities, $y_T = G^{-1}(B_T)$, with G non-linear. By Jensen's inequality:

$$\mathbb{E}^T[G(y_T)] \;\neq\; G(\mathbb{E}^T[y_T])$$

So even though prices match expectations under Q^T, yields don't. The expected future yield must be **adjusted** from the time-0 forward yield, this is the convexity adjustment.

Equal bond price increments do NOT correspond to equal yield increments Equal yield drops at high yields produce small price gains; equal yield drops at low yields produce large price gains. This price-yield curvature is the bond's **convexity**, exactly the second-order term that contaminates yield expectations.

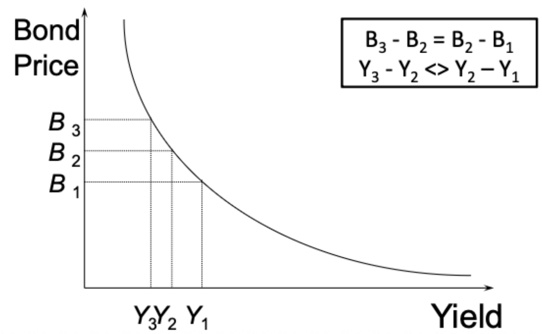

⋆ **Convexity-Adjusted Forward Yield.** Under the T-forward measure, the expected yield at time T is the forward yield plus a correction term involving the curvature of G:

$$\boxed{\mathbb{E}^T[y_T] \;=\; y_F \;-\; \underbrace{\tfrac{1}{2}\, y_F^2\, \sigma_y^2\, T\, \frac{G''(y_F)}{G'(y_F)}}_{\text{convexity adjustment}}}$$

- $G(y)$: bond price as a function of yield (non-linear, decreasing)
- y_F: forward bond yield at time 0
- σ_y: forward yield volatility
- T: time to the contract's reference observation
- $G''(y_F)/G'(y_F)$: ratio of second to first derivative of G, captures the curvature

The adjustment scales with T (longer horizon, more time for variance to matter), with σ_y^2 (more volatile yield, more Jensen's bias), and with the curvature ratio G''/G' (more convex price-yield, larger correction).

⋆ **Adjustment for a Single-Period Rate.** For a **contract paying the τ-period rate at time T,** the underlying is a τ-period ZCB with $G(y) = 1/(1 + y\tau)$:

$$G(y_F) = \frac{1}{1 + y_F\tau}, \quad G'(y_F) = -\frac{\tau}{(1 + y_F\tau)^2}, \quad G''(y_F) = \frac{2\tau^2}{(1 + y_F\tau)^3}$$

Negative G' flips the sign of the adjustment to positive:

$$\boxed{\mathbb{E}^T[y_T] \;=\; y_F \;+\; \frac{y_F^2\, \sigma_y^2\, T\, \tau}{1 + y_F\tau}}$$

The adjusted yield is **larger** than the forward yield, magnitude scaling in y_F^2, σ_y^2, T, and τ.

Contract value with Adj. A contract paying the τ-period rate at T on notional L, scaled to dollars:

$$\boxed{V_0 \;=\; P(0,T) \cdot L \cdot \tau \cdot \left[\, y_F + \frac{y_F^2\, \sigma_y^2\, T\, \tau}{1 + y_F\tau} \right]}$$

Same dollar scaling as a regular rate payoff (notional × accrual × rate), but evaluating the rate at the convexity-adjusted forward instead of the raw forward. The bracketed term IS the convexity-adjusted forward rate, that's the only piece that differs from the standard pricing.

Pricing Procedure.

1. Read y_F and σ_y from the curve and vol surface.
2. Compute the convexity-adjusted forward: $y_F + y_F^2 \sigma_y^2 T\tau/(1 + y_F\tau)$.
3. Multiply by $L \cdot \tau$ to convert to dollars.
4. Discount at $P(0, T)$.

Scale anchor: a 10-year contract on the 1-year rate at flat 4% with $\sigma_y = 20\%$ gives a convexity adjustment of only 6 bp (4% becomes 4.06%). The adjustment grows linearly in T and quadratically in σ_y, so 30-year contracts on volatile rates can see adjustments of tens of basis points.

5.2 Term Structure Models

⋆ **Interest Rate Modeling.** While Black's model is a direct pricing formula for one derivative at a time, telling you today's price and nothing else, a **term structure model** describes how the **entire interest rate curve evolves over time**. These models are the overarching framework for pricing anything interest-rate-sensitive in a consistent, full-curve way. Within that framework, **short-rate models** describe the curve by modeling just the instantaneous short rate $r(t)$ and deriving everything else; inside short-rate models sit the **equilibrium** and **no-arbitrage** sub-families.

Term structure models break down into several subcategories based on what you directly model.

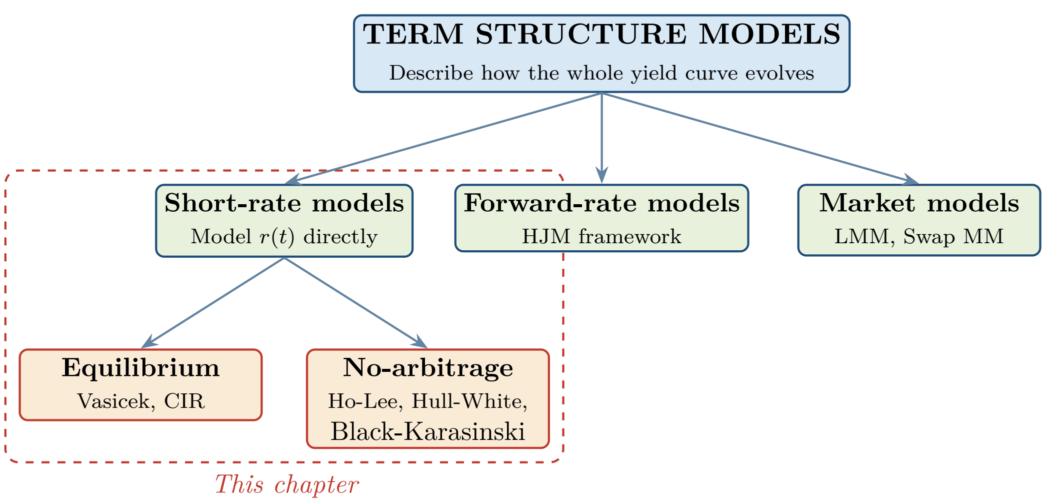

Modeling complexities.

- Rate volatility varies by tenor (**short rates more volatile than long rates**, plus skew effects similar to equity options)
- Payoffs and discounting are correlated, both depend on the same underlying rates
- Must accommodate negative rates and avoid unbounded explosions

Unlike equities where one stock has one price, an interest-rate model must describe an entire curve at every point in time. This is why every rate model carries a parameter set whose job is to keep different-maturity behavior consistent.

Short rate $r(t)$. The **instantaneous** risk-free rate, the rate earned over an infinitesimal interval $[t,\, t+dt]$:

$$r(t) = \lim_{\Delta t \to 0} \text{rate over } [t,\, t+\Delta t]$$

Zero tenor, the shortest possible rate. Real markets proxy it with overnight SOFR (or Fed Funds, pre-2023 LIBOR); models like Vasicek and Hull-White treat it as the true instantaneous rate.

The short rate is the ***modeling primitive*** *at the bottom of the rate hierarchy: every observable rate on the curve (1-month, 1-year, 30-year yields) is derived from how $r(t)$ is expected to evolve under the risk-neutral measure.*

⋆ **Two Modeling Philosophies.** Term structure models split into two camps based on what they prioritize:

Type	Starting point	What it fits
Equilibrium model	Economic assumptions, derive short rate dynamics	Derives a term structure, doesn't try to match today's curve
No-arbitrage model	Short rate / forward rate dynamics directly	Calibrated to exactly match today's bond prices and some option prices

Equilibrium models are theoretical frameworks, useful for understanding how rates behave but not for pricing real derivatives, since they generally won't reproduce the current curve. No-arbitrage models add a time-dependent drift that's calibrated to fit the curve exactly, which is what trading desks actually use.

⋆ **Discount Factor with Stochastic Rates.** The zero-coupon bond price in a stochastic-rate world is the risk-neutral expectation of accumulated discounting:

$$\boxed{P(t,T) \;=\; \mathbb{E}^Q\left[e^{-\int_t^T \tilde{r}(u)\,du}\right]}$$

- $\tilde{r}$: risk-neutral stochastic process for the instantaneous short rate, the modeling primitive
- Q: risk-neutral probability measure

Generalizes the deterministic $P(t,T) = e^{-r(T-t)}$. With r random, you cannot pull discounting outside the expectation, the integral inside the exponent itself is now stochastic. Every term structure model is essentially a recipe for evaluating this expectation in closed form or numerically.

What a risk-neutral short rate process must specify.

- **Number of factors** (one for single-factor, two or more for richer dynamics)
- **Functional form for volatility**, with mean reversion to keep rates bounded
- **Volatility structure** that gives different-tenor bonds different volatilities
- **Consistency with traded bond prices** (no-arbitrage class)
- Optional consistency with some option prices (no-arbitrage class)

Strategy: model one random variable, the short rate $r(t)$, and derive everything else (bond prices, yields, the yield curve) from that one model.

⋆ **Mean Reversion.** The drift of r is **negative when rates are high** and **positive when rates are low**, pulling r back toward a long-run level b. Encoded in SDEs as a drift term proportional to $(b-r)$.

- **High rates** ($r > b$): drift $(b-r) < 0$, rate pulled **down**
- **Low rates** ($r < b$): drift $(b-r) > 0$, rate pulled **up**

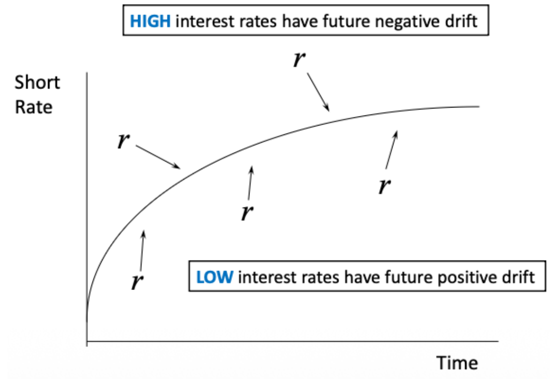

Three canonical term structure shapes. Zero rate $\tilde{r}$ as a function of maturity:

- **Upward sloping**: long rates higher than short rates (normal economy)
- **Inverted**: short rates above long rates (often a recession signal)
- **Hump-shaped**: long-end and short-end below the middle, transitional regime

Any plausible model must produce all three shapes. Equilibrium models with fixed parameters struggle here, especially with humps, one reason traders prefer no-arbitrage models.

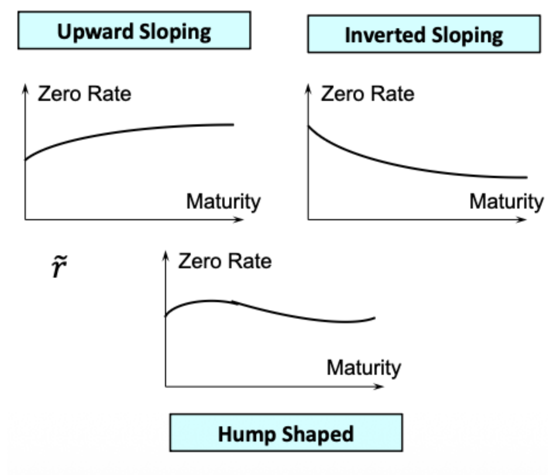

5.2.1 Equilibrium Models

Equilibrium Models. Take economic assumptions about how rates *should* behave (mean reversion, volatility) as **input**, and produce a full term structure curve as **output**, useful for understanding rate dynamics, but won't generally match today's observed market curve.

Mental picture: ***economic theory in*** *$\to$* ***curve shape out****, no calibration to live market prices.*

⋆ **Vasicek Model.** The simplest mean-reverting short-rate model: $r(t)$ **wanders randomly with constant volatility σ while being pulled back toward a long-run level b at speed a.** Allows negative rates (Gaussian).

Single-factor mean-reverting Gaussian short-rate model: $\boxed{dr \;=\; a(b-r)\,dt + \sigma\,dz}$

- a: mean reversion strength, how fast r pulls back to b
- b: mean reversion level, the long-run target rate
- σ: short rate volatility (constant, additive)

⋆ **Cox-Ingersoll-Ross (CIR) Model.** Same mean-reverting drift as Vasicek, but volatility scales with $\sqrt{r}$: vol shrinks as rates approach zero, so rates stay non-negative under the Feller condition.

Variant of Vasicek where volatility scales with $\sqrt{r}$: $\boxed{dr \;=\; a(b-r)\,dt + \sigma\sqrt{r}\,dz}$

Same drift as Vasicek, only the diffusion changes. The $\sqrt{r}$ does two things at once: it shrinks volatility to zero as $r \to 0$ (so r cannot cross into negative territory if $2ab \geq \sigma^2$, the ***Feller condition****), and it makes vol higher when rates are high. Vasicek's constant σ allows negative rates, CIR does not.*

⋆ **Affine Model of Bond Prices.** Both Vasicek and CIR yield closed-form bond prices in the **affine form**:

$$\boxed{P(t,T) \;=\; A(t,T)\, e^{-B(t,T)\, r}}$$

where $A(t,T)$ and $B(t,T)$ are deterministic functions of (a,b,σ), computed once and reused for every bond.

What "affine" means. Taking logs: $\ln P(t,T) \;=\; \underbrace{\ln A(t,T)}_{\text{intercept}} \;-\; \underbrace{B(t,T)}_{\text{slope}} \cdot r$

$\ln P$ is a **linear function of** r, slope $-B(t,T)$ plus intercept $\ln A(t,T)$. That's the literal meaning of "affine": linear plus a constant.

Why this is the gold-standard property: once you know r, every $P(t,T)$ on the entire curve is just a plug-and-play formula. No simulation, no numerical integration. The model "compresses" the whole term structure into one number r plus the deterministic functions A, B. Bond options, swaptions, and Greeks all inherit this tractability and reduce to closed-form (or near-closed-form) expressions.

Properties.

- **Mean-reverting**, so unlike equity models the rate is bounded in expectation
- **Few parameters**, so generally cannot fit current bond prices exactly without adjustments
- **Tractable**: log-linear-in-r structure cascades cleanly through every downstream pricing calculation

⋆ **$A(t,T)$ and $B(t,T)$ Formulas.** Closed-form expressions, useful only as reference:

Vasicek:

$$B_{\text{Vas}}(t,T) = \frac{1 - e^{-a(T-t)}}{a}$$

$$A_{\text{Vas}}(t,T) = \exp\left[\frac{(B_{\text{Vas}}(t,T) - (T-t))(a^2 b - \sigma^2/2)}{a^2} - \frac{\sigma^2 B_{\text{Vas}}(t,T)^2}{4a}\right]$$

CIR:

$$B_{\text{CIR}}(t,T) = \frac{2(e^{\gamma(T-t)} - 1)}{(\gamma + a)(e^{\gamma(T-t)} - 1) + 2\gamma}$$

$$A_{\text{CIR}}(t,T) = \left[\frac{2\gamma\, e^{(a+\gamma)(T-t)/2}}{(\gamma + a)(e^{\gamma(T-t)} - 1) + 2\gamma}\right]^{2ab/\sigma^2}, \qquad \gamma = \sqrt{a^2 + 2\sigma^2}$$

- $B(t,T)$: how sensitive the bond price is to a move in the short rate, the model's *duration* for a T-maturity bond. Larger B means rate changes hit the price harder.
- $A(t,T)$: the leftover "convexity correction" piece, captures everything in the bond price that isn't a direct function of r, including time decay and the variance contribution from σ.

B controls rate sensitivity, A packages everything else, so that $P = A \cdot e^{-Br}$ closes the loop.

Important caveat. Neither Vasicek nor CIR fits actual bond prices without adjustment. The model produces a term structure from the parameter set (a,b,σ), not the other way around. This is the classic limitation that motivates the no-arbitrage models in slide 75 onward.

⋆ **Mean and Variance of r_t.** Both models share the same mean (the drift is identical), but the variance differs because of the $\sqrt{r}$ in CIR's diffusion.

Mean (both Vasicek and CIR): Given r_0 at $t = 0$: $\boxed{\mathbb{E}[r_t] \;=\; b + (r_0 - b)\, e^{-at}}$

Mean reversion toward b at rate a. Starts at r_0, decays exponentially toward b with half-life $\ln 2/a$.

Variance.

$$\textbf{Vasicek:} \quad \mathrm{Var}(r_t) = \frac{\sigma^2}{2a}\left(1 - e^{-2at}\right)$$

$$\textbf{CIR:} \quad \mathrm{Var}(r_t) = \frac{\sigma^2}{a}\left[r_0\, e^{-at}(1 - e^{-at}) + \frac{b}{2}(1 - e^{-at})^2\right]$$

Both variances saturate (approach a finite limit) as $t \to \infty$. Vasicek's long-run variance is $\sigma^2/(2a)$, the spring's resistance limits how far r can wander. CIR's variance also depends on r_0 and b, since the diffusion $\sigma\sqrt{r}$ amplifies as the rate level rises.

Vasicek term structures. With fixed parameters (a, b, σ), Vasicek can produce **upward sloping**, **inverted**, and **hump-shaped** curves depending on the initial r_0 relative to b, but only weakly. The hump shape is especially hard to generate cleanly.

This is the core limitation: a small parameter set can sketch the right qualitative shapes but cannot match every observed market curve exactly. For pricing exotic derivatives where curve fit matters, you need the no-arbitrage extension (Hull-White) introduced after this section.

⋆ **Interest Rate Sensitivity in Vasicek.** Defined relative to the **short rate** r (the model's one factor), not parallel curve shifts:

$$\boxed{\text{Duration} \;=\; -\frac{1}{P(t,T)}\frac{\partial P(t,T)}{\partial r} \;=\; B(t,T)}$$

$$\boxed{\text{Convexity} \;=\; \frac{1}{P(t,T)}\frac{\partial^2 P(t,T)}{\partial r^2} \;=\; B(t,T)^2}$$

Key consequences.

- Different-maturity bonds have different sensitivities to r
- All maturities are **perfectly correlated** since they share the same single factor
- Different from traditional parallel-shift duration (which equals maturity for a ZCB)

In a one-factor model the entire curve is a deterministic function of one number r, so any rate move is a coordinated move along that function. Real markets show imperfect correlation across maturities, which is exactly the motivation for two-factor and HJM models later.

⋆ **Risk-Neutral Bond Price Dynamics.**
Applying Itô's lemma to the affine bond price gives a clean SDE for $P(t, T)$:

$$\boxed{\text{Vasicek:} \quad dP(t,T) \;=\; r\,P(t,T)\,dt - \sigma B_{\mathrm{Vas}}(t,T)\,P(t,T)\,dz}$$

$$\boxed{\text{CIR:} \quad dP(t,T) \;=\; r\,P(t,T)\,dt - \sigma\sqrt{r(t)}\,B_{\mathrm{CIR}}(t,T)\,P(t,T)\,dz}$$

- **Drift** $=$ r: bond is a tradeable asset, must have risk-neutral drift r
- **Volatility** $= \sigma B(t,T)$ **(Vasicek) or** $\sigma\sqrt{r}\,B(t,T)$ **(CIR)**: bond price vol depends on the underlying rate model and the maturity-dependent $B(t,T)$

Consistent with general risk-neutral theory, every traded asset drifts at r under Q. The bond's volatility shrinks as $T \to t$ since $B(t,t) = 0$, the bond's price approaches its known maturity value of \$1.

⋆ **Real-World Conversion via Market Price of Risk.**

The risk-neutral process $dr = a(b - r)dt + \sigma\, dz$ converts to the real-world process by adjusting the drift through the **market price of interest rate risk λ**:

$$\boxed{dr \;=\; a(b^* - r)\, dt + \sigma\, dz, \quad b^* = b + \frac{\lambda\sigma}{a}}$$

- λ: market price of interest rate risk, the extra return per unit of rate-risk exposure the market demands. **Negative** for bonds since rates and bond prices move opposite ways. Defined by the drift gap: $\lambda = \frac{\text{real-world drift} - \text{risk-neutral drift}}{\sigma}$ a Sharpe-ratio-like quantity, estimated from historical rate data, observed bond price spreads vs risk-neutral pricing, or bond-return Sharpe ratios.
- Same volatility σ, only the long-run mean b shifts

Real-world rates are theoretically lower than risk-neutral rates ($\lambda < 0 \Rightarrow b^* < b$), so real-world bond prices are higher than risk-neutral. Pricing always uses the risk-neutral process; real-world is only relevant for forecasting and risk management. The market price of risk is the bridge between the two.

⋆ **General Interest Rate Derivative Valuation.** Any derivative $f(r,t)$ with no intermediate cash flows, written under a generic risk-neutral short-rate process, satisfies a Black-Scholes-type PDE.

Setup. The short rate evolves under risk-neutral drift $m(r,t)$ and volatility $s(r,t)$, generic enough to contain Vasicek, CIR, Hull-White, etc. as special cases:

$$dr \;=\; m(r,t)\, dt \;+\; s(r,t)\, dz$$

The valuation PDE. Applying Itô's lemma to $f(r,t)$ and standard arbitrage arguments yield:

$$\boxed{\frac{\partial f}{\partial t} \;+\; m\,\frac{\partial f}{\partial r} \;+\; \tfrac{1}{2}\, s^2\,\frac{\partial^2 f}{\partial r^2} \;=\; r\, f}$$

- $\partial f/\partial t$: time decay, how f changes purely from clock-time passing
- $m\,\partial f/\partial r$: drift-weighted delta, expected change in f from the rate's expected drift
- $\frac{1}{2}s^2\,\partial^2 f/\partial r^2$: variance term, the convexity contribution from rate volatility
- $r\, f$: discounting at the current short rate, no-arbitrage forces the riskless portfolio to grow at rate r

1. Apply Itô's lemma to $f(r,t)$ to get df
2. Construct a riskless portfolio combining f with a hedging instrument
3. No-arbitrage requires the portfolio's drift to equal r times its value
4. Rearrange to get the PDE above

Identical structure to Black-Scholes: time decay + drift-weighted delta + half-vol-squared gamma = risk-free-rate-times-value. Only difference: the underlying is r itself, so the right-hand side discounts at r which IS the underlying. This self-referential structure is what makes interest rate derivatives more complicated than equity ones, the discount rate and the state variable are the same object.

Special case: zero-coupon bond. Plug $f = P(t,T)$ with terminal condition $P(T,T) = 1$ into the general PDE. This produces the master equation whose solution yields $A(t,T)$ and $B(t,T)$ for any affine model.

Vasicek, plug in: $\quad m = a(b - r), \qquad s = \sigma$

CIR, plug in: $\quad m = a(b-r), \qquad s = \sigma\sqrt{r}$

Solving the PDE under each model's (m, s) produces the closed-form $A(t,T)$, $B(t,T)$.

Deterministic case as sanity check. If rates are constant ($m = 0$, $s = 0$) then short rate r constant, the PDE collapses to:

$$\frac{\partial P(t,T)}{\partial t} \;=\; r\,P(t,T), \qquad P(T,T) = 1$$

$$\boxed{P(t,T) \;=\; e^{-r(T-t)}}$$

Recovers the standard deterministic ZCB price. Turning on stochastic dynamics adds the $m \cdot \partial P/\partial r$ and $\frac{1}{2}s^2 \cdot \partial^2 P/\partial r^2$ terms, capturing how the random short rate affects the bond's value through drift expectation and variance contribution.

5.2.2 No Arbitrage Models

Equilibrium models like Vasicek and CIR **cannot match today's bond curve** because their parameter set is too small. No-arbitrage models fix this by adding a **time-dependent function** $\theta(t)$ (or equivalently a deterministic shift $\alpha(t)$) calibrated to reproduce observed bond prices exactly.

Core Idea: A no-arbitrage model decomposes the short rate into a stochastic piece and a deterministic time-dependent piece $\theta(t)$ that is chosen so the model-implied initial bond prices exactly match the observed market curve:

⋆ **Risk-Neutral Bond Pricing with Deterministic Shift.** The general no-arbitrage framework decomposes the short rate into a **stochastic core** plus a **deterministic shift** chosen to fit the curve.

Zero-Coupon Bond Price: At time t maturing at time T with stochastic interest rates

$$\boxed{P(t,T) \;=\; \mathbb{E}^{\mathbb{Q}}\left[\exp\left(-\int_t^T (\tilde{r} + \alpha(t))\,dt\right)\right]}$$

- $\tilde{r}$: zero-drift Gaussian stochastic component, the "engine" of the model
- $\alpha(t)$: deterministic function chosen so model-implied bond prices match the observable initial term structure
- Q: risk-neutral probability measure

Different notation (same formula): $\quad P(t,T) \;=\; \mathbb{E}^{\mathbb{Q}}\big[\exp\big(-\int_t^T (\tilde{r}_s + \alpha(s))\,ds\big)\big]$

- t: **fixed** pricing time ("today"). A specific number, doesn't change inside the integral.
- T: **fixed** maturity time. The bond's expiry, also a specific number.
- s: **running** dummy variable. Sweeps every instant from t to T as the integral accumulates the discount factor.

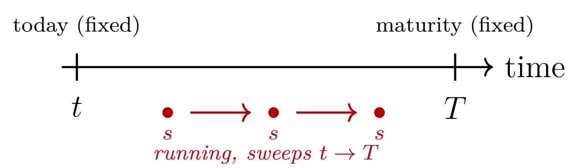

Two equivalent ways to calibrate to the curve. The same model can be written as a **level shift** or as a **drift adjustment**:

Form	Where it lives	Role
$\alpha(t)$	Inside the discount integral, as part of $r(t) = \tilde{r}(t) + \alpha(t)$	**Level shift**, adjusts where the rate sits
$\theta(t)$	Inside the SDE: $dr = [\theta(t) - ar]\,dt + \sigma\,dz$	**Drift adjustment**, adjusts the rate's expected motion

Linking equation. Differentiating $r = \tilde{r} + \alpha(t)$ and matching SDE coefficients gives:

$$\boxed{\theta(t) \;=\; \alpha'(t) + a\,\alpha(t)}$$

A simple first-order ODE: θ is fully determined by α (and vice versa once boundary data is fixed). They are **not** the same object, α is a level, θ is a drift coefficient, but they encode the same calibration information.

⋆ **Ho-Lee Model.**

The simplest no-arbitrage short-rate model: a **driftless Gaussian** process plus a time-varying drift $\theta(t)$ that calibrates to the curve.

$$\boxed{dr \;=\; \theta(t)\,dt + \sigma\,dW_t}$$

$$\theta(t) \;=\; \frac{\partial f(0,t)}{\partial t} + \sigma^2 t, \qquad f(0,t) \;=\; -\frac{\partial \ln P(0,t)}{\partial t}$$

- $f(0,t)$: instantaneous t-forward rate observed at time 0 (derived from the market curve)
- σ: short rate volatility (constant)
- $\theta(t)$: time-varying drift, fit to the initial term structure

*Compared to Vasicek: Ho-Lee has **no mean reversion** ($a = 0$). The drift comes entirely from $\theta(t)$, which is calibrated daily from the curve. Trade-off: perfect curve fit, but rates can drift unboundedly far from any "natural" level. Short rate is Gaussian so negative rates are allowed.*

Ho-Lee bond prices. Affine form, with $A(t,T)$ pinned by the initial market curve:

$$\boxed{P(t,T) \;=\; A(t,T)\,e^{-r(t)(T-t)}} \quad \textit{with} \quad \boxed{\ln A(t,T) \;=\; \underbrace{\ln\frac{P(0,T)}{P(0,t)}}_{\text{curve-matching}} \;+\; \underbrace{\tfrac{1}{2}\,\sigma^2\,t\,(T-t)^2}_{\text{convexity correction}}}$$

- **Curve-matching** $\ln(P(0,T)/P(0,t))$: ratio of today's market bond prices, anchors the model so $P(0,T)$ matches the observed market curve at $t = 0$. Without it, the model would produce some theoretical curve unrelated to what's traded.
- **Convexity correction** Jensen's-inequality adjustment from stochastic rates. Each factor has a clean meaning:
  - $\frac{1}{2}$: standard Itô / Gaussian-MGF coefficient
  - σ^2: variance of the short rate per unit time (Ho-Lee's constant vol, squared because it's variance not std)
  - t: accumulated variance up to time t (Brownian variance grows linearly in t)
  - $(T-t)^2$: squared time-to-maturity, the bond's quadratic sensitivity to rate volatility over the discount horizon

Where it comes from: $P(t,T) = \mathbb{E}^{Q}[e^{-\int_t^T r(s)\,ds}]$ with r Gaussian under Ho-Lee. The Gaussian MGF gives $\mathbb{E}[e^{-X}] = e^{-\mathbb{E}[X] + \frac{1}{2}Var[X]}$. The mean produces the $-r(t)(T-t)$ exponent; the variance $\sigma^2 t(T-t)^2$ produces the convexity term with the $\frac{1}{2}$ factor in front. Without this correction the model would systematically underprice bonds, since Jensen's inequality on e^{-X} implies $\mathbb{E}[e^{-X}] > e^{-\mathbb{E}[X]}$.

Duration and Convexity.

$$\text{Duration} \;=\; -\frac{1}{P(t,T)}\frac{\partial P}{\partial r} \;=\; (T-t), \qquad \text{Convexity} \;=\; \frac{1}{P(t,T)}\frac{\partial^2 P}{\partial r^2} \;=\; (T-t)^2$$

Duration just equals time to maturity, identical to a deterministic ZCB. The model is so simple that bond price sensitivity to r matches the parallel-shift intuition exactly.

⋆ **Hull-White 1-Factor Model.**

Vasicek's mean-reverting structure plus Ho-Lee's $\theta(t)$ for curve fitting. Combines the best of both.

$$\boxed{dr \;=\; [\theta(t) - ar]\,dt + \sigma\,dz}$$

$$\theta(t) \;=\; \frac{dF^M(0,t)}{dt} + a\,F^M(0,t) + \frac{\sigma^2}{2a^2}\left(1 - e^{-2at}\right)$$

$$F^M(0,t) \;=\; -\frac{\partial \ln P^M(0,t)}{\partial t}$$

- $\theta(t)$: deterministic function, fits the initial term structure exactly
- a, σ: mean-reversion and volatility parameters, calibrated to match market option prices
- $F^M(0,t)$: market-observed time-t forward rate at time 0
- $P^M(0,t)$: market-observed t-maturity ZCB price at time 0

- **Mean-reverting** (unlike Ho-Lee): rates don't run off to infinity
- **Curve-matching** (unlike Vasicek): exactly fits today's bond prices via $\theta(t)$
- **Gaussian short rate**: negative rates allowed
- **Closed-form bond and option prices**: tractable for daily desk use

⋆ **Hull-White Bond Prices.** Affine form with $A(t,T)$ pinned to the market curve:

$$\boxed{P(t,T) \;=\; A(t,T)\,e^{-B(t,T)\,r(t)}}$$

$$B(t,T) \;=\; \frac{1 - e^{-a(T-t)}}{a}$$

$$\ln A(t,T) \;=\; \ln\frac{P^M(0,T)}{P^M(0,t)} + B(t,T)\,F^M(0,t) - \frac{\sigma^2}{4a^3}\left(e^{-aT} - e^{-at}\right)^2\left(e^{2at} - 1\right)$$

- $\ln(P^M(0,T)/P^M(0,t)) + B(t,T)F^M(0,t)$: the **curve-matching** component, ensures $P(0,T) = P^M(0,T)$ at $t = 0$
- $-\frac{\sigma^2}{4a^3}(\cdots)$: the **convexity correction** from stochastic dynamics

The $B(t,T)$ formula is identical to Vasicek's, mean reversion enters identically. The new piece is $A(t,T)$, which encodes the entire market curve via P^M and F^M plus a small variance-driven adjustment. Memorize B as $(1 - e^{-a(T-t)})/a$, treat A as reference.

⋆ **Hull-White Properties.**

$\mathbb{E}[r(t)] \;=\; r(0)\,e^{-at} + \alpha(t) - \alpha(0)\,e^{-at}$	Mean drifts toward the long-run level.
$\alpha(t) \;=\; F^M(0,t) + \frac{\sigma^2}{2a^2}\left(1 - e^{-at}\right)^2$	Curve-fit shift plus a convexity bump.
$\mathbb{E}[r(\infty)] \;=\; F^M(0,\infty) + \frac{\sigma^2}{2a^2}$	Long-run forward rate plus convexity correction.
$\mathrm{Var}(r(t)) \;=\; \frac{\sigma^2}{2a}\left(1 - e^{-2at}\right)$	Variance grows from 0, saturates over time.
$\mathrm{Var}(r(\infty)) \;=\; \frac{\sigma^2}{2a}$	Long-run variance cap, identical to Vasicek.
$\text{Duration} \;=\; -\frac{1}{P(t,T)}\frac{\partial P}{\partial r} \;=\; B(t,T)$	Positive, same maturity-dependent form as Vasicek.

Same $\sigma^2/(2a)$ long-run variance as Vasicek, $\theta(t)$ doesn't affect variance. The mean drifts toward $F^M(0,\infty) + \sigma^2/(2a^2)$, the long-run market forward rate plus a small convexity bump. Duration $B(t,T)$ is again the maturity-dependent rate sensitivity, positive and matching the Vasicek formula.

⋆ **Hull-White ZCB Option Price (Jamshidian Formula).** Closed-form for European call/put on a zero-coupon bond:

$$\boxed{\text{Call} \;=\; L\,P(0,T_B)\,N(h) - K\,P(0,T)\,N(h-\sigma_P)}$$

$$\boxed{\text{Put} \;=\; K\,P(0,T)\,N(-h+\sigma_P) - L\,P(0,T_B)\,N(-h)}$$

$$h \;=\; \frac{1}{\sigma_P}\ln\frac{L\,P(0,T_B)}{P(0,T)\,K} + \frac{\sigma_P}{2}, \qquad \sigma_P \;=\; \frac{\sigma}{a}\left(1-e^{-a(T_B-T)}\right)\sqrt{\frac{1-e^{-2aT}}{2a}}$$

- L: notional of the underlying bond
- K: strike price (in dollars, since the underlying is a bond price not a rate)
- T: option maturity, when exercise is decided
- T_B: bond maturity ($T_B > T$), the bond keeps living after the option expires
- h: the d_1 analog, controls $N(h)$ which acts like a delta term
- $h-\sigma_P$: the d_2 analog, $N(h-\sigma_P)$ is the **risk-neutral probability of ending in-the-money**
- σ_P: **integrated bond price volatility**, the total standard deviation of the log forward bond price between today and option maturity. The analog of $\sigma\sqrt{T}$ in Black-Scholes

Identical structure, only the underlying and total-vol quantities change. σ_P bundles short-rate vol σ, mean reversion a, option maturity T, and bond maturity T_B into one number.

Breaking down σ_P: $$\sigma_P \;=\; \underbrace{\frac{\sigma}{a}\left(1-e^{-a(T_B-T)}\right)}_{a\cdot B(T,T_B)\text{: bond duration at expiry}} \cdot \underbrace{\sqrt{\frac{1-e^{-2aT}}{2a}}}_{\text{std of integrated short rate up to } T}$$

- First factor: how sensitive the bond's price is to a unit move in the short rate at option maturity, this is just $\sigma \cdot B(T,T_B)$
- Second factor: the standard deviation of how much the short rate can have moved by time T
- Product: total bond price volatility accumulated from today to option maturity

⋆ **Black-Karasinski 1-Factor Model.**

A **lognormal** short-rate model: take logs of r and apply mean-reverting Gaussian dynamics.

$$\boxed{d\ln r \;=\; [\theta'(t) - a(t)\,\ln r]\,dt + \sigma(t)\,dz}$$

- Future r is **lognormal** (since $\ln r$ is Gaussian)
- **Precludes negative rates** (lognormal $r > 0$ always)
- **Very little analytic tractability**, no closed-form bond or option prices, requires trees or simulation

Shifted version. Generalization that allows rates down to $-\alpha$:

$$d\ln(r+\alpha) \;=\; [\theta'(t) - a(t)\,\ln(r+\alpha)]\,dt + \sigma(t)\,dz$$

Bounds rates at $-\alpha$ instead of zero, useful when the strict-positive constraint of Black-Karasinski is too tight (e.g., negative-rate environments).

Trade-off vs Hull-White: Black-Karasinski guarantees positive rates but loses the closed-form machinery. In practice this matters less now since Hull-White's negative-rate possibility is acceptable in most markets, and tree-based pricing handles BK fine.

Black-Karasinski VS Hull-White. Hull-White is Gaussian on r directly; Black-Karasinski is Gaussian on $\ln r$. This single change cascades into several differences:

Property	Hull-White	Black-Karasinski
Distribution of r	Gaussian	Lognormal
Negative rates?	Allowed	Precluded ($r > 0$ always)
Bond price formula	Closed-form (affine)	No closed form
Option pricing	Closed-form (Jamshidian)	Trees or Monte Carlo only
Volatility behavior	Constant σ on r	Constant σ on $\ln r$, so vol of r scales with r

⋆ **Hull-White 2-Factor Model.**
Adds a second stochastic factor u (a mean-reverting "long-rate" component) to Hull-White 1-Factor, producing richer curve dynamics that single-factor models miss.

- **r, level factor**: the instantaneous short rate, drives **parallel curve shifts** (entire curve moving up or down together). Same role as in the 1-factor model.
- **u, slope factor**: an auxiliary mean-reverting component, drives **curve shape changes** (steepening, flattening, twisting) independent of overall level moves.

$$\boxed{dr \;=\; [\theta(t) + u - ar]\,dt + \sigma_1\,dz_1} \quad \text{with} \quad \boxed{du \;=\; -bu\,dt + \sigma_2\,dz_2}$$

$$\rho\,dt \;=\; \mathbb{E}[dz_1\,dz_2]$$

Bond price form.

$$\boxed{P(t,T) \;=\; A(t,T)\,e^{-B(t,T)\,r - C(t,T)\,u}}$$

$$\frac{1}{P(t,T)}\frac{\partial P}{\partial r} \;=\; -B(t,T), \qquad \frac{1}{P(t,T)}\frac{\partial P}{\partial u} \;=\; -C(t,T)$$

Two-factor structure breaks single-factor's perfect-correlation flaw: now different-maturity bonds can have imperfect correlation, controlled by ρ and the relative magnitudes of σ_1, σ_2. Functional forms for A, B, C are in Hull's Technical Note No. 14, treat them as reference.

⋆ **Hull-White 2-Factor ZCB option price.** Same Black-shape formula as the 1-factor case, but σ_p is computed differently to account for the second stochastic factor:

$$\boxed{\text{Call} \;=\; L\,P(0,T_B)\,N(h) - K\,P(0,T)\,N(h - \sigma_p)}$$

$$\boxed{\text{Put} \;=\; K\,P(0,T)\,N(-h + \sigma_p) - L\,P(0,T_B)\,N(-h)}$$

$$h \;=\; \frac{1}{\sigma_p}\ln\frac{L\,P(0,T_B)}{P(0,T)\,K} + \frac{\sigma_p}{2}$$

The σ_p here is NOT the 1-factor formula. The 2-factor model has two Brownians, so bond price volatility has three contributions:

- Variance from σ_1 (short-rate factor)
- Variance from σ_2 (long-rate factor u)
- Covariance term $\rho\sigma_1\sigma_2$ between the two

Schematically:

$$\sigma_p^{2\text{F}} = \sqrt{\text{Var}_{\sigma_1} + \text{Var}_{\sigma_2} + 2\rho \cdot \text{cross-term}}$$

The full closed-form is given in Hull's Technical Note No. 23. *Same modular structure as the 1-factor version: the only difference is what σ_p contains. Once you have σ_p, the rest of the formula is mechanical, just plug into the Black-Scholes-shaped expressions for call and put.*

⋆ **Parameter Estimation for No-Arbitrage Models.**

Parameter	**Data source**	**Practical procedure**
$\theta(t)$	Today's bootstrapped bond/swap curve	Strip ZCB prices $P^M(0,t)$ from market quotes, compute forward curve $F^M(0,t)$, evaluate $\theta(t)$ in closed form. Re-done daily.
a, σ	Market prices of liquid options (caps, swaptions)	Solve $\min_{a,\sigma} \sum_i (V_i^{\text{model}} - U_i^{\text{market}})^2$ over a basket of ATM instruments. Re-calibrated when option prices shift materially.
λ *(real-world only)*	Historical short-rate time series	Compare real-world drift (from regression) to risk-neutral drift (from curve), take $\lambda \approx (\mu^{\text{real}} - \mu^Q)/\sigma$. Treasury values typically $\lambda \in [-0.4, -0.1]$.

Clean separation of concerns: $\theta(t)$ is curve-driven, (a, σ) are option-driven, λ is data-driven. This separability is what makes Hull-White practically useful, you re-fit θ daily without disturbing the calibrated (a, σ), and never need λ unless doing real-world risk work.

5.3 Interest Rate Trees

⋆ **Why Trees for Interest Rates.** A **trinomial tree** is the practical computational tool when closed-form pricing breaks down: **path-dependent payoffs** (Bermudan swaptions, callable bonds) or any structure with early exercise. Equity-style binomial trees don't directly transfer because rates have unique features that demand a different design.

Three properties unique to interest rate trees.

- **Discounting is node-and-time dependent**: every node has its own short rate r, used to discount one step back. Equity trees use a single risk-free rate; rate trees use a varying r.
- **Rates must be bounded**: real rates don't go to $\pm\infty$, the tree limits the range with mean-reversion. Stock prices are unbounded above; rates are not.
- **Payoff and discounting are correlated**: both depend on the same r, so a high-rate node has both a large potential payoff (e.g., for a caplet) AND a large discount factor pulling its present value down.

Risk-neutral rates aren't tied to real-world rate expectations, the tree's job is to be consistent with bond and option market prices, not to forecast actual rate paths. The "rates" in the tree are pricing-measure rates, not the rates you'd see in a Federal Reserve forecast.

5.3.1 Hull-White Tree

⋆ **Hull-White Tree Building: 5-Step Procedure.** Given the SDE $\boxed{dr = [\theta(t) - ar]\,dt + \sigma\,dz}$ the tree is constructed in two stages: build a no-drift skeleton, then shift to fit the curve.

1. **Build no-drift tree** R: $dR = -aR\,dt + \sigma\,dz$, with $\theta(t) = 0$, $R(0) = 0$. **R is an auxiliary process**, NOT the actual short rate, just a symmetric scaffold centered at zero. The dynamics depend only on a (mean reversion) and σ (volatility), so this step is independent of the market curve.
2. **Match mean and standard deviation at each node**: solve three constraints (mean drift $-aR\Delta t$, variance $\sigma^2\Delta t$, probabilities sum to 1) for the three branching probabilities (p_u, p_m, p_d). Different formulas at normal vs upper/lower nodes.
3. **Shift to fit the curve**: determine $\alpha(t)$ iteratively so the model reproduces today's bond prices $P^M(0,t)$. The actual short rate is $r(t) = R(t) + \alpha(t)$. **One α_i per time step**, each pinned by one bond price.
4. **Calibrate (a, σ) to option prices**: minimize pricing errors against liquid caps/swaptions. Recalibrated less frequently than α (weekly, not daily) since vol parameters are stable.
5. **Price interest rate derivatives**: backward induction on the calibrated tree, payoffs at terminal nodes, discount step-by-step using each node's local rate.

*Two-stage philosophy: (a, σ) control the **dynamics** (Steps 1, 2, 4), $\alpha(t)$ controls the **curve fit** (Step 3). Building R first instead of r directly lets you separate these: rebuild the curve-fit shift daily without redoing the tree skeleton. Same separation of concerns as in continuous Hull-White, just discretized.*

⋆ **Three Branching Processes.** To keep R bounded, the tree switches branching style based on how far the node is from the center. Distance measured in units of ΔR, with j being the node index.

- **Normal (a)**: standard $\nearrow, \rightarrow, \searrow$, used for $|j| < j_{\max}$. Three branches go up, mid, and down by ΔR.
- **Upper (c)**: node is at $+j_{\max}$ (top of the tree), branches go $\rightarrow, \searrow, \searrow\searrow$. **No "up" branch**, prevents rates from drifting unboundedly high.
- **Lower (b)**: node is at $-j_{\max}$ (bottom), branches go $\nearrow\nearrow, \nearrow, \rightarrow$. **No "down" branch**, prevents unboundedly low rates.

The branching switch is the discrete analog of mean reversion. Continuous Hull-White pulls rates back via $-ar$ in the drift; the tree pulls rates back by changing branch direction at the boundaries. Same effect, different mechanism.

Node indexing in the tree. Each node has two coordinates that locate it on the time-vs-state grid:

- i: time step (column index), $i = 0, 1, 2, \ldots$
- j: vertical position (row index), counts how many ΔR steps the node sits above ($j > 0$) or below ($j < 0$) the central row

R-value at any node. $\boxed{R(i,j) \;=\; j \cdot \Delta R}$ Central row is $j = 0$ (where $R = 0$). Outer rows $j = \pm j_{\max}$ use upper/lower branching to keep rates bounded; all rows in between use normal branching.

The R-value depends only on j since the central column is anchored at $R = 0$ at every time. The shift α_i is what makes the actual rate r time-dependent: $r(i,j) = R(i,j) + \alpha_i = j\Delta R + \alpha_i$.

⋆ **Step 1: The R Tree (no-drift skeleton).** Built around $R = 0$ with vertical node spacing:

$$\boxed{\Delta R \;=\; \sigma\sqrt{3\,\Delta t}}$$

$$\boxed{j_{\max} \;=\; \text{smallest integer greater than } 0.184/(a\,\Delta t)}$$

Why $\Delta R = \sigma\sqrt{3\Delta t}$? This is the spacing that minimizes interpolation error and matches the variance $\sigma^2 \Delta t$ of R over a time step.

Why $j_{\max} \approx 0.184/(a\Delta t)$? Hull-White showed this is the threshold beyond which middle-branching probabilities would go negative. Switching to upper/lower branching at $\pm j_{\max}$ keeps probabilities valid.

Probability constraints at every node.

- **Mean of change in** $R = -aR\Delta t = -aj\Delta R \cdot \Delta t$
- **Std of change in** $R = \sigma\sqrt{\Delta t}$
- **Probabilities sum to 1**

Three constraints (mean, variance, sum) determine the three probabilities (p_u, p_m, p_d). Different formulas apply at normal vs upper vs lower branching nodes since the geometry is different.

⋆ **Step 1: Branching Probabilities.** Solving the three matching constraints gives different probabilities for each branching type.

Normal branching (used for $|j| < j_{\max}$):

$$p_u = \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - aj\Delta t), \qquad p_m = \frac{2}{3} - a^2 j^2 \Delta t^2, \qquad p_d = \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + aj\Delta t)$$

Upper branching (at $j = +j_{\max}$):

$$p_u = \frac{7}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - 3aj\Delta t), \quad p_m = -\frac{1}{3} - a^2 j^2 \Delta t^2 + 2aj\Delta t, \quad p_d = \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 - aj\Delta t)$$

Lower branching (at $j = -j_{\max}$):

$$p_u = \frac{1}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + aj\Delta t), \quad p_m = -\frac{1}{3} - a^2 j^2 \Delta t^2 - 2aj\Delta t, \quad p_d = \frac{7}{6} + \frac{1}{2}(a^2 j^2 \Delta t^2 + 3aj\Delta t)$$

Two key properties: (i) at $j = 0$, normal branching simplifies to $p_u = p_d = 1/6$, $p_m = 2/3$ (the standard "central" trinomial weights), (ii) probabilities vary with j, since farther-from-center nodes have stronger mean reversion pulling them back.

- **Each node** sits at a (i, j) coordinate: time step (column) and vertical row.
- **Vertical spacing** between adjacent rows is ΔR, identical at every time step.
- **Outer rows** (top and bottom) hit $\pm j_{\max}$, where branching switches to upper/lower mode to bend rates back toward the center.

The tree extends rightward in time but its vertical span is capped at $\pm j_{\max} \cdot \Delta R$.

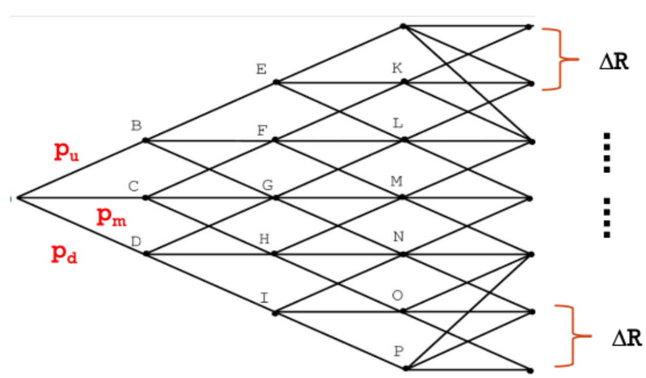

General definition of α_i. The shift α_i is the level adjustment applied to all nodes at time step i, chosen so the model reproduces the $(i+1)$-year market bond price. Determined iteratively, one α per time step.

Given $Q(i, j)$ already computed, α_i solves: $\boxed{P^M(0,\, i+1) \;=\; \sum_j Q(i,j) \cdot e^{-(\alpha_i + j\Delta R)\Delta t}}$

Equivalently, isolating α_i: $\boxed{\alpha_i \;=\; \dfrac{\ln \sum_j Q(i,j)\, e^{-j\Delta R\,\Delta t} \;-\; \ln P^M(0,\, i+1)}{\Delta t}}$

⋆ **Worked Example: Hull-White Tree Calibration.** Build a 3-step trinomial tree from given parameters and a market curve.

Setup parameters.

- $\sigma = 0.01$, $a = 0.1$, $\Delta t = 1$ year
- Market zero rates: 1y = 3.824%, 2y = 4.512%, 3y = 5.086%

Compute tree spacing.

$$\Delta R = \sigma\sqrt{3\Delta t} = 0.01\sqrt{3} = 1.732\%$$

$$j_{\max} = \text{ceil}(0.184/(0.1 \cdot 1)) = \text{ceil}(1.84) = 2$$

So nodes at $j \in \{-2, -1, 0, +1, +2\}$, with ± 2 using upper/lower branching.

Step 1: R-tree values (after computing probabilities). Nine nodes A-I:

Node	A	B	C	D	E	F	G	H	I
R	0.000%	1.732%	0.000%	-1.732%	3.464%	1.732%	0.000%	-1.732%	-3.464%
p_u	0.1667	0.1217	0.1667	0.2217	0.8867	0.1217	0.1667	0.2217	0.0867
p_m	0.6666	0.6566	0.6666	0.6566	0.0266	0.6566	0.6666	0.6566	0.0266
p_d	0.1667	0.2217	0.1667	0.1217	0.0867	0.2217	0.1667	0.1217	0.8867

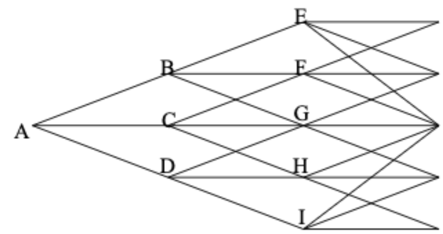

Notice nodes E (top, $j = +2$) and I (bottom, $j = -2$) use upper/lower branching, hence their lopsided probabilities (E has $p_u = 0.89$ for "stay near top", I has $p_d = 0.89$). Middle nodes have the symmetric weights.

Step 2: Shift Nodes by α_i to Fit the Curve. Each time step gets its own shift α_i, chosen so the model reproduces the observed $(i+1)$-year bond price. Worked forward through the tree.

Arrow-Debreu prices $Q(i, j)$. Define $Q(i, j)$ as today's value of a security paying \$1 at node j at time $i\Delta t$, and zero elsewhere. These are the building blocks for solving for α_i.

Time step 0 (α_0). At $t = 0$ there's only the root node ($R = 0$), so the sum collapses to one term and α_0 has a closed form:

$$\alpha_0 \;=\; -\frac{\ln P^M(0,1)}{\Delta t} \;=\; y_1 \;=\; 3.824\%$$

This is the only step where α_i equals a market rate directly. Propagate forward with $Q(0, 0) = 1$:

$$Q(1, 1) = p_u^{\text{root}} \cdot e^{-\alpha_0 \Delta t} = 0.1667 \cdot e^{-0.03824} = 0.1604$$

$$Q(1, 0) = p_m \cdot e^{-\alpha_0} = 0.6666 \cdot 0.9625 = 0.6417, \qquad Q(1, -1) = 0.1604$$

Time step 1 (α_1). The 2-year ZCB price from market is $P^M(0, 2) = e^{-0.04512 \cdot 2} = 0.9137$. Plug into the general formula with three nodes at time 1:

$$0.9137 \;=\; Q(1,1)\, e^{-(\alpha_1 + 0.01732)} + Q(1,0)\, e^{-\alpha_1} + Q(1,-1)\, e^{-(\alpha_1 - 0.01732)}$$

Solve numerically: $\alpha_1 = 5.205\%$. No closed form exists since the equation has multiple exponential terms.

Year-1 rates after shifting.

- Node B: $\alpha_1 + 1 \cdot \Delta R = 5.205\% + 1.732\% = 6.937\%$
- Node C: $\alpha_1 + 0 \cdot \Delta R = 5.205\%$
- Node D: $\alpha_1 - 1 \cdot \Delta R = 5.205\% - 1.732\% = 3.473\%$

Time step 2 (α_2). First compute $Q(2,j)$ for all five terminal nodes via forward induction.

Example: $Q(2,1)$ for node F. Node F (at $j = +1$, time 2) is reachable from B and C:

$$Q(2,1) \;=\; \underbrace{Q(1,1)\cdot p_m^B \cdot e^{-r_B \Delta t}}_{\text{from B via mid branch}} + \underbrace{Q(1,0)\cdot p_u^C \cdot e^{-r_C \Delta t}}_{\text{from C via up branch}}$$

$$= 0.1604 \cdot 0.6566 \cdot e^{-0.06937} + 0.6417 \cdot 0.1667 \cdot e^{-0.05205} = 0.1998$$

All Q(2,j) values: $Q(2,2) = 0.0182$, $Q(2,1) = 0.1998$, $Q(2,0) = 0.4736$, $Q(2,-1) = 0.2033$, $Q(2,-2) = 0.0189$.

Solve for α_2. Market 3-year ZCB price: $P^M(0,3) = e^{-0.05086\cdot 3} = 0.8585$. Set:

$$0.8585 \;=\; \sum_{j=-2}^{+2} Q(2,j)\cdot e^{-(\alpha_2 + j\Delta R)\Delta t} \qquad \mathit{Solve}: \alpha_2 = 6.25\%.$$

Each α_i is determined by ONE bond price: the $(i+1)$-year zero. This is why the procedure is sequential, you must complete time step i before solving for α_{i+1}. After all α_i are found, the tree exactly matches every observed market bond price, by construction.

Final calibrated tree. After all shifts, every node has its specific short rate:

Node	A	B	C	D	E	F	G	H	I
r	3.824%	6.937%	5.205%	3.473%	9.716%	7.984%	6.252%	4.520%	2.788%

This tree is now ready for pricing any interest rate derivative. Backward induction works the same as in the motivating example: start at terminal payoffs, apply expected-discounted-value at each node going back.

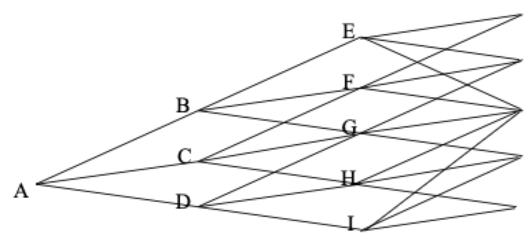

$\star$ **Worked Example: 2-Year Caplet on the Tree.** Strike 11%, observed at year 2, **paid at year 3** (settlement in arrears), notional 100. Payoff $= \max(r - 11\%, 0)\cdot 100$.

Why the timing matters. Caplet pays at $t+\delta$ based on rate observed at t. Observation at year 2 + payment at year 3 = **one extra discount factor** compared to a vanilla year-2 option.

Backward induction (rolled into one chain).

$$\text{Year 2 payoffs:} \quad \{3,\ 1,\ 0,\ 0,\ 0\} \text{ for } r \in \{14\%, 12\%, 10\%, 8\%, 6\%\}$$

$$\text{Year 1 upper } (r = 12\%): \quad (0.25\cdot 3 + 0.5\cdot 1)\cdot e^{-(0.12+0.14)} = 1.25\cdot e^{-0.26} = 0.973$$

$$\text{Year 1 mid } (r = 10\%): \quad 0.25\cdot 1\cdot e^{-(0.10+0.12)} = 0.205$$

$$\text{Year 1 lower } (r = 8\%): \quad 0$$

$$\boxed{\text{Year 0 (root, } r = 10\%): \quad (0.25\cdot 0.973 + 0.5\cdot 0.205)\cdot e^{-0.10} = 0.317}$$

Each backward step uses TWO discount factors instead of one. The exponent $-(r_i + r_{i+1})\Delta t$ reflects: rate observed at year 2 times discount-from-year-3-back-to-year-1. Same settlement-in-arrears convention as in the Black caplet, just discretized on the tree.

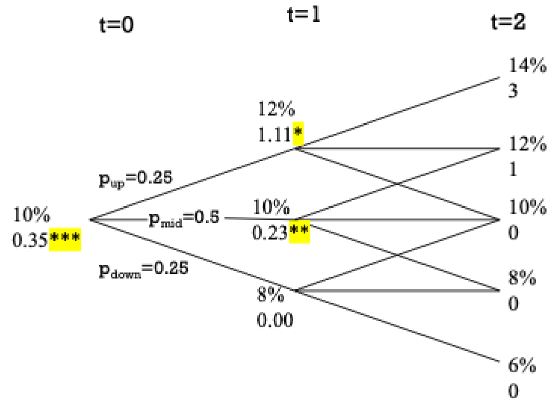

⋆ **Analytic Implementation: Forward Induction for α_m and $Q(m, j)$.** The algorithm in compact form, given $Q(m, j)$ already computed:

Solve for α_m from the $(m + 1)$-year bond price.

$$\boxed{P_{m+1} \;=\; \sum_{j=-n_m}^{+n_m} Q(m, j) \cdot e^{-(\alpha_m + j\Delta R)\Delta t}}$$

where n_m is the number of nodes on each side at time m. Solve for α_m:

$$\boxed{\alpha_m \;=\; \frac{\ln \sum_{j=-n_m}^{+n_m} Q(m, j)\, e^{-j\Delta R\Delta t} - \ln P_{m+1}}{\Delta t}}$$

Solve for $Q(m + 1, j)$ via forward induction.

$$\boxed{Q(m + 1, j) \;=\; \sum_k Q(m, k) \cdot q_{k,j} \cdot e^{-(\alpha_m + k\Delta R)\Delta t}}$$

where $q_{k,j}$ is the probability of moving from node (m, k) to node $(m + 1, j)$.

This is an inductive recipe: given $Q(m, \cdot)$, you can compute α_m from the $(m + 1)$-year bond, then compute $Q(m + 1, \cdot)$ from the transition probabilities. Repeat until the tree is fully built. The algorithm is $O(N^2)$ in tree depth.

Recursive procedure (worked through manually).

Initialization (time 0):

- $Q(0, 0) = 1$
- $P_1 = Q(0, 0) \cdot e^{-\alpha_0 \Delta t} \Rightarrow \alpha_0 = -\ln P_1/\Delta t$

Time 1 nodes B, C, D:

$$Q(1, +1) = Q(0, 0) \cdot p_u \cdot e^{-\alpha_0\Delta t}, \quad Q(1, 0) = Q(0, 0) \cdot p_m \cdot e^{-\alpha_0\Delta t}, \quad Q(1, -1) = Q(0, 0) \cdot p_d \cdot e^{-\alpha_0\Delta t}$$

$$P_2 = \sum_{j=-1}^{+1} Q(1, j) \cdot e^{-(\alpha_1 + j\Delta R)\Delta t} \Rightarrow \alpha_1 = \frac{1}{\Delta t} \ln \sum_{j=-1}^{+1} Q(1, j)\, e^{-j\Delta R\Delta t} - \frac{\ln P_2}{\Delta t}$$

Time 2 nodes E, F, G, H, I: similar pattern, summing over all parent paths into each $(m + 1, j)$ node.

Extensions to other models. Generalize the SDE to:

$$df(r) = [\theta(t) - af(r)]\,dt + \sigma\,dz$$

- $f(r) = r$: Hull-White (the case we built above)
- $f(r) = \ln r$: Black-Karasinski (lognormal rates, no negative rates allowed)

Same tree-building procedure works: build the no-drift skeleton on $f(r)$, then shift to fit the curve. For BK, after the tree is built you exponentiate to get $r = e^{f(r)}$ at each node.

*The 5-step procedure is essentially **model-agnostic**, change $f(r)$ and you change the model. This is why trees are the practical workhorse for any short-rate model that lacks closed-form solutions.*

⋆ **Calibrating a and σ to Market Option Prices.** The volatility parameters are chosen to match observed prices of liquid traded instruments, typically caps, floors, and swaptions.

Optimization problem.

$$\boxed{\min_{a,\sigma} \;\sum_{i=1}^{n} (U_i - V_i)^2 + P}$$

- U_i: market price of the i-th calibrating instrument
- V_i: model price of the i-th calibrating instrument (computed from the tree given a, σ)
- P: penalty function discouraging large changes in a, σ between successive calibration runs (smoothness over time)

Two-stage workflow on a desk: (a, σ) are calibrated less frequently (e.g., weekly) using a basket of liquid options. Once fixed, $\theta(t)$ is re-fit daily to today's curve via the Step 2 shifting procedure. The penalty term P keeps day-to-day vol parameters stable so risk metrics (Greeks) don't jump around.

5.3.2 Joint Equity Interest Rate Model

Standard Black-Scholes assumes constant rates. For long-dated equity options, callable structured notes, or products mixing equity and rate exposure, that approximation breaks down. This section builds a joint model where **both the stock and the short rate are stochastic and correlated**, then derives a Black-Scholes-like option formula that accounts for the rate uncertainty.

⋆ **Joint Equity-Rate Dynamics (Risk-Neutral).** Equity follows lognormal dynamics with stochastic discount rate; short rate follows Hull-White; the two Brownian motions are correlated.

$$\boxed{\frac{dS}{S} \;=\; (r(t) - q)\,dt + \sigma_S\,dZ} \quad s.t. \quad \boxed{dr(t) \;=\; [\theta(t) - ar(t)]\,dt + \sigma_r\,dW}$$

$$\boxed{dZ\,dW \;=\; \rho\,dt} \qquad \theta(t) = \frac{\partial f^M_{0,t}}{\partial t} + af^M_{0,t} + \frac{\sigma_r^2}{2a}\left(1 - e^{-2at}\right)$$

- S: equity price (lognormal under risk-neutral)
- $r(t)$: stochastic short rate (Hull-White)
- q: continuous dividend yield
- σ_S: equity volatility (constant, no stochastic vol assumed)
- σ_r: short-rate volatility
- ρ: correlation between equity and rate shocks
- $f^M_{0,t}$: initial market forward curve, used to calibrate $\theta(t)$

*Same equity SDE as Black-Scholes, but the drift now contains the **stochastic** short rate $r(t)$ instead of a constant. The short-rate process is the standard Hull-White from earlier, unchanged. New ingredient:*

the correlation ρ ties them together, capturing how equity and rates co-move (typically negative for stocks since rate hikes are bad for equity).

$\star$ **Equity Option Price with Stochastic Rates.** Same Black-Scholes shape, but the discount factor and total variance are upgraded.

$$\boxed{C(t,T,K) \;=\; S(t)\,e^{-q(T-t)}\,N(d_1) - K\,P(t,T)\,N(d_2)}$$

$$\boxed{P(t,T,K) \;=\; K\,P(t,T)\,N(-d_2) - S(t)\,e^{-q(T-t)}\,N(-d_1)}$$

$$d_1 = \frac{\ln\dfrac{S(t)}{K\,P(t,T)} - q(T-t) + \frac{1}{2}\sigma_{T,t}^2}{\sigma_{T,t}}, \qquad d_2 = d_1 - \sqrt{\sigma_{T,t}^2(T-t)}$$

Two key changes from standard Black-Scholes.

- **Discount factor**: $e^{-r(T-t)}$ becomes $P(t,T)$, the actual stochastic-rate ZCB price. Already curve-aware.
- **Variance**: $\sigma_S^2(T-t)$ becomes $\sigma_{T,t}^2$, the **total integrated variance** that bundles equity vol, rate vol, and their covariance.

Assumptions.

- Joint Gaussian distribution of equity and rate shocks
- Non-stochastic equity volatility (σ_S constant)
- Need to estimate the total variance $\sigma_{T,t}^2$ (next slide)

The formula reduces to standard Black-Scholes when $\sigma_r = 0$ (rates deterministic), so this is a strict generalization. For short-dated options, the rate-vol contribution is negligible and Black-Scholes is fine; for long-dated options, ignoring rate vol systematically misprices.

$\star$ **Total Variance Decomposition.** The total variance $\sigma_{T,t}^2$ splits into three contributions: rate, equity, and a covariance cross-term.

Rate variance. Variance of the integrated discount factor under Hull-White:

$$V(t,T) \;=\; \frac{\sigma_r^2}{a^2}\left[(T-t) + \frac{2}{a}e^{-a(T-t)} - \frac{1}{2a}e^{-2a(T-t)} - \frac{3}{2a}\right]$$

Equity variance. Standard Brownian variance:

$$\sigma_S^2(T-t)$$

Total variance: $$\boxed{\sigma_{T,t}^2 \;=\; \underbrace{V(t,T)}_{\text{rate}} + \underbrace{\sigma_S^2(T-t)}_{\text{equity}} + \underbrace{\frac{2\rho\sigma_r\sigma_S}{a}\left[(T-t) - \frac{1}{a}\left(1 - e^{-a(T-t)}\right)\right]}_{\text{covariance}}}$$

Schematic: **Total Variance = Rate Variance + Equity Variance + Covariance**.

The covariance term carries the sign of ρ, with negative ρ (typical for equity-rate correlation) the cross-term is negative and total variance is below the simple sum. For long-dated options this decomposition matters because the rate-variance contribution grows faster than the equity-variance contribution.

Calibration principle.

- Market quotes **implied volatility** as a single number (model-free)
- Implied vol is implicitly a **total volatility** (covers all sources of uncertainty)
- The model splits this total into equity, rate, and covariance components based on its specific structure

- Decomposition is **model-specific**; different models will allocate the total variance differently

⋆ **Calibration Strategies.** Two practical approaches depending on what the product is sensitive to.

Strategy 1: Joint calibration. Calibrate to BOTH equity options AND interest rate options (caps, swaptions) simultaneously.

- Used for derivatives with strong sensitivity to BOTH equity and rates (e.g., callable convertibles, equity-linked notes with rate embedded options)
- Solves for $\sigma_S, \sigma_r, a, \rho$ all together
- Computationally expensive but most accurate for hybrid products

Strategy 2: Fix rates, calibrate equity. Take rate parameters (σ_r, a) and correlation ρ as given (from rate market), then solve for σ_S to match equity options.

- Used when the product is primarily equity-sensitive with rate optionality secondary
- Standard $\theta(t)$ still calibrates to the initial term structure regardless
- Faster, decouples the equity calibration from the rate calibration

Practical rule: long-dated equity options on indices typically use Strategy 2 (rate parameters from the rate desk's calibration). Truly hybrid products (callable convertibles, rate-linked equity notes) use Strategy 1. The choice is about where the dominant risk lives.

5.3.3 Forward Rate Models

⋆ **Forward Rate Models.** A different approach to term-structure modeling: instead of modeling the single instantaneous short rate $r(t)$ and **deriving** forward rates, model the **entire forward curve** directly and let $r(t) = f(t,t)$ fall out as a special case (the diagonal).

- **Heath-Jarrow-Morton (HJM)**: framework for continuous instantaneous forward rates $f(t,T)$. **Infinite-dimensional** unless restricted by separable volatility (which makes it collapse to a known short-rate model like Hull-White).
- **Brace-Gatarek-Musiela (BGM)**: models a discrete set of forward rates (e.g., 3-month, 6-month tenors), **finite-dimensional** arbitrage-free restriction of HJM, easier to calibrate. Also called the **LIBOR Market Model (LMM)**, the workhorse for cap and swaption pricing on real desks.

Short-rate models compress the curve into one number $r(t)$; forward rate models track every forward simultaneously. Trade-off: forward rate models capture richer dynamics (curve twists, tenor-specific vol) but are harder to handle analytically. HJM in full generality requires Monte Carlo or specialized restrictions; BGM is the practical industry standard.

HJM Notation.

- $P(t,T)$: discount bond price at time t for principal \$1 maturing at T
- $f(t,T)$: instantaneous forward rate at time t for maturity T, defined as $f(t,T) = -\partial \ln P(t,T)/\partial T$
- $r(t)$: short-term risk-free rate at t, equals $f(t,t)$ (the diagonal of the forward curve)
- $dW(t)$: Brownian motion driving term structure movements
- $\Sigma_p(t,T)$: volatility of bond price $P(t,T)$
- $\mu(t,T)$: drift of the instantaneous forward rate
- $\sigma_f(t,T)$: volatility of the instantaneous forward rate

⋆ **HJM Bond Price Dynamics.** Under the risk-neutral measure, every bond price follows an Itô process with drift r and tenor-dependent volatility:

$$\boxed{dP(t,T) \;=\; r(t)\,P(t,T)\,dt + \Sigma_p(t,T)\,P(t,T)\,dz(t)}$$

with boundary conditions $P(t,t) = 1$ and $\Sigma_p(t,t) = 0$.

Forward rate from finite-difference limit.

$$f(t,T_1,T_2) \;=\; \frac{\ln P(t,T_1) - \ln P(t,T_2)}{T_2 - T_1}, \qquad f(t,T) = \lim_{T_2 \to T_1} f(t,T_1,T_2)$$

Bond drift is r because bonds are tradeable assets, must drift at risk-free rate under Q. The volatility $\Sigma_p(t,T)$ vanishes at maturity since $P(T,T) = 1$ exactly. Forward rates are derived from bond prices via differences in log-prices.

⋆ **HJM Forward Rate Process.** Applying Itô's lemma to the forward rate definition $f = -\partial \ln P/\partial T$:

$$\boxed{df(t,T) \;=\; \mu(t,T)\,dt + \sigma_f(t,T)\,dW(t)}$$

HJM Drift Condition (one-factor case).

$$\boxed{\mu(t,T) \;=\; \sigma_f(t,T)\int_t^T \sigma_f(t,\tau)\,d\tau} \quad s.t \quad \sigma_f(t,T) \;=\; -\frac{\partial}{\partial T}\Sigma_p(t,T)$$

The HJM Result. The drift of the instantaneous forward rate is fully determined by its volatility.

This is the central HJM theorem: ***specify the volatility, and no-arbitrage forces the drift****. You don't need to model rates directly, you only need to model the forward-rate vol surface, and the drift is implied. This is why HJM is called a "framework", it's a recipe for constructing arbitrage-free models given any volatility specification.*

HJM Result, proof sketch. From Itô's lemma applied to $\ln P$:

$$d\ln P(t,T_1) = \left[r(t) - \frac{\Sigma_p(t,T_1)^2}{2}\right] dt + \Sigma_p(t,T_1)\,dW(t)$$

$$d\ln P(t,T_2) = \left[r(t) - \frac{\Sigma_p(t,T_2)^2}{2}\right] dt + \Sigma_p(t,T_2)\,dW(t)$$

Subtract, divide by $T_2 - T_1$, take limit $T_2 \to T_1$, gives the forward rate process with drift $\mu(t,T)$ pinned by no-arbitrage.

In HJM, the short rate is the diagonal of the forward curve: $r(t) = f(t,t)$. Its dynamics are model-implied and generally path-dependent (depends on the entire forward curve history), unlike in short-rate models where $r(t)$ is the modeling primitive.

⋆ **Separable Volatility.** A specific class of HJM volatilities that yield tractable closed forms. The vol is the product of a time-of-trading function and a time-to-maturity function:

$$\boxed{\sigma_f(t,T) \;=\; \sigma(t)\,g(T-t), \qquad g(0) = 1}$$

Drift under separable form.

$$\mu(t,T) \;=\; \sigma_f(t,T)\int_t^T \sigma_f(t,\tau)\,d\tau \;=\; \sigma(t)^2\,g(T-t)\int_0^{T-t} g(s)\,ds$$

Exponential kernel $g(x) = e^{-ax}$ for $a > 0$:

$$\int_0^{T-t} e^{-as}\, ds \;=\; \frac{1 - e^{-a(T-t)}}{a}$$

Separable structure is what makes HJM tractable. The exponential kernel is the standard choice because it's the form that produces Hull-White as a special case (next slide). Other kernels give other short-rate models.

⋆ **Separable HJM Reduces to Hull-White.** With $\sigma_f(t,T) = \sigma(t)e^{-a(T-t)}$ and short rate $r(t) = f(t,t)$, taking the limit $T \to t$:

$$\boxed{dr(t) \;=\; [\theta(t) - ar(t)]\, dt + \sigma(t)\, dW_t}$$

This is exactly Hull-White, with:

$$\theta(t) \;=\; \frac{dF^M(0,t)}{dt} + aF^M(0,t) + \frac{\sigma^2}{2a^2}\left(1 - e^{-2at}\right) \quad \text{[function of } t \text{ only, not } T\text{]}$$

Critical structural insight: ***Hull-White is a special case of HJM*** *with exponential separable volatility. So everything from slides 79-82 (Hull-White formulas, θ calibration, ZCB option price) can be derived from the HJM framework. This is why the broader HJM framework matters even if you mostly use Hull-White: it tells you exactly which restriction you're imposing.*

⋆ **BGM (LIBOR Market Model).** Models discretely compounded forward rates $F_k(t)$ associated with accrual periods $[t_k, t_{k+1}]$ rather than continuous instantaneous forwards $f(t,T)$. **Finite-dimensional, arbitrage-free, easier to calibrate**.

Setup.

- **Each forward rate is lognormal** under its own forward measure
- Forward rates typically set at **3-month tenors** (LIBOR/SOFR convention)
- A 10-year model with 3-month tenors needs **40 separate forward rates**
- Forward-rate volatilities are implied from cap, floor, and swaption market prices

Why BGM is the industry standard: instantaneous forward rates are abstract constructs (no instrument trades on them), but the discrete forwards F_k ARE the LIBOR/SOFR rates that markets actually quote and trade. BGM models the rates that traders actually see, with a vol per forward calibrated directly from cap/swaption quotes.

- t_k: k-th reset date
- $F_k(t)$: forward rate at time t for the period $[t_k, t_{k+1}]$
- $m(t)$: index of the next reset date, $m(t) = \min\{k : t < t_{k+1}\}$
- $\delta = t_{k+1} - t_k$: equal accrual spacing (e.g., 0.25 for 3-month)
- $\Lambda_k(t)$: volatility of $F_k(t)$

Stationary Volatility Structure. The vol of $F_k(t)$ depends only on the number of accrual periods between the next reset date and t_k, NOT on the absolute calendar time:

$$\boxed{\Lambda_k(t) \;=\; \Lambda_{k-m(t)}}$$

Stationarity says "a forward rate's vol depends on how far in the future it is, not when we're observing it." This dramatically reduces the number of parameters and matches market behavior: a 3-month-forward rate has a similar vol regardless of whether observed today or a year from now, as long as it's still 3 months away.

Forward rate evolution matrix. Each forward rate F_k evolves through time, but only "exists" until its reset date t_k. Beyond that, it becomes a fixed/realized rate. The triangular structure of the matrix reflects this: $F_k(t_j)$ is defined only for $j \leq k$.

⋆ **Rolling Forward Measure.** Numeraire choice that makes BGM tractable for Monte Carlo.

Rolling Forward Numeraire. The bond maturing at the next reset date. Under this measure, you can discount from t_{i+1} to t_i at the rate observed at time t_i (the "obvious" choice).

F_k process under rolling forward measure.

$$\boxed{\frac{dF_k(t)}{F_k(t)} = \sum_{i=m(t)}^{k} \frac{\delta\, F_i(t)\, \Lambda_{i-m(t)}\, \Lambda_{k-m(t)}}{1+\delta\, F_i(t)}\, dt + \Lambda_{k-m(t)}\, dz}$$

Two key observations.

- Under the $T_{m(t)+1}$-forward measure, the drift of F_k depends on **all intervening forwards** $F_{m(t)}, F_{m(t)+1}, \ldots, F_k$. This is because the numeraire change introduces drift terms tied to every shorter forward.
- In the limit $\delta \to 0$ (continuous compounding), BGM with rolling risk-neutrality becomes the HJM model in the traditional risk-neutral world.

The drift summation is the price you pay for using the rolling numeraire: it's not a clean lognormal SDE, the drift involves the very forwards you're modeling. This makes Monte Carlo non-trivial but still much faster than infinite-dimensional HJM.

BGM Monte Carlo (Euler discretization).

$$F_k(t_{j+1}) = F_k(t_j)\exp\left[\sum_{i=j+1}^{k} \frac{\delta\, F_i(t_j)\, \Lambda_{i-j-1}\, \Lambda_{k-j-1}}{1+\delta\, F_i(t_j)}\,\delta - \frac{\Lambda_{k-j-1}^2}{2}\,\delta + \Lambda_{k-j-1}\,\varepsilon\sqrt{\delta}\right]$$

- **Drift approximation**: evaluated at t_j and held constant over $[t_j, t_{j+1}]$ ("drift freezing")
- **Lognormal Euler discretization**: standard log-space update with the Itô correction $-\Lambda^2/2$
- ε: standard normal random draw

Drift freezing is the practical compromise: the true drift depends on path-dependent $F_i(t)$ values, but using $F_i(t_j)$ as an approximation for the entire interval $[t_j, t_{j+1}]$ is accurate enough for typical $\delta = 0.25$ and dramatically simpler to simulate.

⋆ **Volatility Structure Calibration.** Connect BGM forward vols Λ_i to market caplet vols σ_k.

Setup.

- σ_k: caplet volatility for the period $[t_k, t_{k+1}]$ (Black/market convention)
- Λ_i: forward volatility for the period $[t_i, t_{i+1}]$ (BGM model)

Cap-pricing consistency condition: $\boxed{\sigma_k^2\, t_k = \sum_{i=1}^{k} \Lambda_{k-i}^2\, \delta}$

Worked example. Given Black caplet vols across three periods, recover the forward vols Λ_i:

Period	1	2	3
$\sigma_{\textbf{Black}}$	24%	22%	20%
Λ	24.0%	19.8%	15.2%

1. $\Lambda_0 = \sigma_1 = 24.0\%$ (only one term in the sum)
2. $\Lambda_1 = \sqrt{2 \cdot 0.22^2 - 1 \cdot 0.24^2} = 19.8\%$
3. $\Lambda_2 = \sqrt{3 \cdot 0.20^2 - 2 \cdot 0.22^2} = 15.2\%$

The formula $\sigma_k^2 t_k = \sum \Lambda_{k-i}^2 \delta$ says: a caplet covering k periods accumulates variance from each of the k forward-rate volatilities. Solving sequentially: each new Λ is determined by inverting the cap-vol-to-forward-vol relation. ***Assumes 1-factor model with independent increments and flat accrual spacing.***

⋆ **BGM Calibration.** Two flavors depending on which products you're pricing.

Cap calibration.

- BGM can be **exactly calibrated** to caplet prices in theory (one Λ per caplet)
- Caps require decomposition into caplets (a cap is a strip of caplets)

Swaption calibration.

- Analytic approximation exists for European swaptions in BGM
- Includes the joint correlation of forward rates (which a strip of caplets cannot capture alone)

Calibration as optimization.

$$\boxed{\min \sum_{i=1}^{n} (U_i - V_i)^2 + P}$$

- U_i: market price of the i-th calibrating instrument (cap, swaption, etc.)
- V_i: BGM model price of the i-th instrument
- P: penalty function discouraging large changes or excess curvature in vols

Practical workflow: caps give you the marginal vol structure (one Λ per caplet via the recursive formula), swaptions give you the correlation structure (joint dynamics across forwards). A full BGM calibration combines both, balancing fit quality across the cap and swaption surfaces with smoothness in the vol parameters.

Chapter 6

Credit Risk

6.1 Credit Risk Measures

⋆ **Credit Risk.** Inability of an asset to pay a contractual obligation. **Risk decreases with collateralization**, fully collateralized derivatives (e.g., futures margins) carry virtually none.

⋆ **Credit Spread.** The additional return above the **risk-free rate** that compensates investors for bearing default risk.
The spread is the market's price for default risk: it embeds both the expected loss from defaults and an additional risk premium (e.g., for illiquidity).

Credit Risk Types:
- Default Risk. Risk contingent on the underlying asset actually defaulting.
Instruments exposed to this risk include the **Credit Default Swap** (buys/sells protection against a default event) and the **Basket Default Swap** (payoff contingent on one or more defaults in a basket).

- Spread Risk. Risk arising from **credit spreads changing**, independent of whether a default occurs. Instruments exposed include credit contingent loans and total return bond funds.

Default risk is a discrete event; spread risk is continuous: a position loses value any time the market reprices an issuer's creditworthiness, even if no default ever happens.

⋆ **Credit Discount Factor**, Present value of \$1 promised by a defaultable bond, discounting at the risk-free rate *plus* the issuer's credit spread:

$$\boxed{\tilde{P}(0,T) = e^{-(R_C + Spd)\,T}}$$

- R_C = continuously compounded risk-free rate
- Spd = firm's **credit spread** over the risk-free rate; compensation for actual defaults plus a risk premium (e.g., illiquidity)
- $R_C + Spd$ = simply the risky bond's total yield
- Spread changes cause bond price volatility alongside changes in R_C

Adding Spd to the discount rate lowers every promised cash flow's present value: a wider spread means the market demands more return, so the bond is worth less today.

⋆ **Risky Bond Price.** Sum of all cash flows discounted with the **credit discount factor**:

$$\boxed{B_0 = \sum_{n=1}^{Tm} \frac{C}{m}\,\tilde{P}(0,t_n) + 100\,\tilde{P}(0,T)} \qquad C = r_f + Spd \quad (coupon)$$

Identical in structure to a risk-free bond price, but every discount factor $\tilde{P}$ carries the spread, so the risky bond always prices below an equivalent risk-free bond, with the gap widening as the spread rises.

Credit spreads are time-varying and spike sharply during periods of market stress, reflecting sudden increases in perceived default risk and illiquidity risk premiums.

Credit Risk Historical Measures

⋆ **Credit Ratings.** Ordinal scores assigned by rating agencies (S&P, Moody's, Fitch, AM Best, Kroll) to measure a firm's future default likelihood. Ratings are **ordinal**: they rank creditworthiness but are not direct numeric default probabilities. They update only when material new information changes the long-run default outlook, known as **through-the-cycle** methodology.

Because ratings are through-the-cycle, they lag point-in-time market signals like spreads: a firm's spread can widen significantly before any rating action occurs.

The scale splits into two categories:

- **Investment Grade** (AAA/Aaa to BBB/Baa): preferred or required by institutional investors such as pensions and insurers
- **Non-Investment Grade** (BB/Ba and below): speculative / high-yield / junk; historically far higher and more cyclical default rates, spiking sharply during recessions

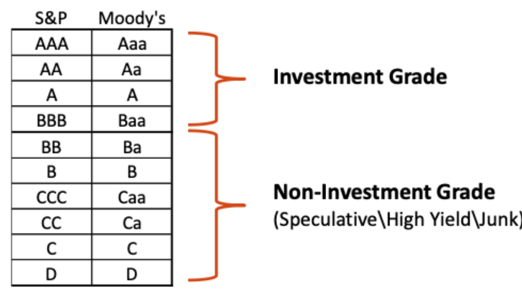

S&P	Moody's
AAA	Aaa
AA	Aa
A	A
BBB	Baa
BB	Ba
B	B
CCC	Caa
CC	Ca
C	C
D	D

Default Rate Measures

Probability of Default = [PD]

⋆ **Cumulative PD: $PD(0,T)$.** Probability of defaulting at any point between now and horizon T. Always increasing in T: the longer the window, the more chances to default.
$PD(0,T)$ has no closed-form formula; it is read directly from historical cumulative default rate tables (e.g. Moody's), where each entry gives the empirically observed fraction of issuers of a given rating that defaulted within T years.

⋆ **Annualized PD: $\overline{PD}(0,T)$.** Per-year default rate implied by the cumulative PD, as if risk were spread evenly across each year:

$$\boxed{\overline{PD}(0,T) = 1 - \Big(1 - PD(0,T)\Big)^{1/T}}$$

The geometric annualization of the survival probability: the annual rate that, compounded over T years, reproduces the observed cumulative default rate.

⋆ **Unconditional PD: $PD_U(t-1,\, t)$.** Probability of defaulting in the window $(t-1,\, t)$, measured from time 0 with no conditioning on prior survival:

$$\boxed{PD_U(t-1,\, t) = PD(0,t) - PD(0,t-1)}$$

⋆ **Conditional PD: $PD_C(t-1,\, t)$.** Probability of defaulting between $t-1$ and t, *given* the firm survived to $t-1$. This is the discrete-time **hazard rate**:

$$\boxed{PD_C(t-1,\, t) = \frac{PD(0,t) - PD(0,t-1)}{1 - PD(0,t-1)}}$$

The removes already-defaulted firms: conditioning on survival shrinks the pool, making PD_C the more relevant measure for pricing default risk over a specific interval.

The shape of PD_C over time differs by rating class:

- **Investment Grade:** Hazard rates are flat to mildly increasing: a high-quality firm has little room to improve but can always deteriorate over time
- **Non-Investment Grade:** Hazard rates *decline* with survival time due to a **selection effect**: the weakest firms default early, leaving a progressively stronger survivor pool. A speculative-grade firm that survives several years has implicitly demonstrated resilience relative to its original cohort

Recovery Rate & Loss Given Default

⋆ **Recovery Rate (RR).** The market price of a bond immediately after default, expressed as a fraction of face value: what creditors can expect to get back. Estimated empirically from observed post-default trading prices, and varies by seniority of the claim.
Idea: $RR = 40\%$ means the bond trades at \$40 per \$100 face value immediately after default.

Loss Given Default: $\boxed{LGD = 1 - RR}$

Idea: How much you lose given default has already happened. It's a loss conditional on default occurring, expressed per dollar of face value.

RR is strongly driven by seniority: senior secured lenders recover first from the asset pool, leaving less for junior claimants. The standard market convention for senior unsecured bonds is $RR \approx 40\%$, implying $LGD \approx 60\%$.

Expected credit loss combines both dimensions of risk: $\boxed{\text{Loss} = PD \times LGD}$

Your expected loss before knowing if default happens. It weights the loss-if-default by the probability that default actually occurs.

Idea: A bond can have a high PD but low LGD (assets are valuable in default), or a low PD but high LGD (total wipeout if default occurs); both matter for expected loss.

6.1.1 Rating Transition Matrix

⋆ **Transition Matrix.** A matrix of **one-year probabilities of migrating from one rating to another.** Each row represents the starting rating, each column the ending rating; rows must sum to 1. The diagonal entries (staying in the same rating) are the largest, reflecting rating stability.

The transition matrix captures not just the risk of default but the full distribution of credit quality changes: a position can lose value simply by being downgraded, even without defaulting.

	Aaa	Aa	A	Baa	Ba	B	Caa	Ca-C	Def	Sum
Aaa	91.19%	8.09%	0.62%	0.07%	0.02%	0.00%	0.00%	0.00%	0.00%	100.0%
Aa	0.83%	89.70%	8.90%	0.44%	0.06%	0.04%	0.02%	0.00%	0.02%	100.0%
A	0.05%	2.61%	91.21%	5.44%	0.48%	0.10%	0.04%	0.00%	0.05%	100.0%
Baa	0.03%	0.14%	4.23%	90.79%	3.80%	0.68%	0.15%	0.02%	0.16%	100.0%
Ba	0.01%	0.05%	0.43%	6.64%	83.49%	7.60%	0.76%	0.12%	0.91%	100.0%
B	0.01%	0.03%	0.14%	0.47%	5.45%	82.61%	7.32%	0.56%	3.41%	100.0%
Caa	0.00%	0.01%	0.02%	0.08%	0.37%	7.37%	80.58%	3.13%	8.44%	100.0%
Ca-C	0.00%	0.00%	0.05%	0.00%	0.66%	2.79%	12.24%	48.96%	35.30%	100.0%
D	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	100.00%	100.0%

The raw matrix includes a **Withdrawn Rating (WR)** column for firms whose ratings were withdrawn during the year (e.g., due to debt repayment or acquisition). Since WR is not a true credit state, the matrix can be **rescaled**.

⋆ **Transition Matrix Evolution.** To price multi-period credit risk, we need an n-year transition matrix $X_{0,n}$, built by chaining together annual matrices:

$$\boxed{X_{0,n} = X_{0,1}\, X_{1,2} \cdots X_{n-1,n}}$$

In practice, forward annual matrices $X_{t,t+1}$ are not directly observable, so they are approximated by the empirically observed $X_{0,1}$. Critically, this means: $X_{0,n} \neq (X_{0,1})^n$

Simply raising the 1-year matrix to the n-th power assumes credit migration is Markovian and stationary; empirically this understates multi-year default risk because rating drift and path-dependence matter over longer horizons.

6.2 Default Survival Modeling

Two main frameworks for modeling when and whether a firm defaults
The key distinction: in reduced form, default is a surprise with no advance warning; in structural models, default is predictable: you can see the firm's value approaching the barrier.

6.2.1 Reduced Form Model

⋆ **Reduced Form Model.** Default time τ is treated as a surprise, modeled as a **Cox process** with a time-varying intensity λ_t. The survival probability (probability of no default up to time t) is:

$$\boxed{P(\tau > t) = \exp\left(-\int_0^t \lambda(s)\,ds\right)}$$

- $P(\tau > t)$ = probability the firm has not defaulted by time t (survival probability).
- τ = random default time, so $\tau > t$ means survival.
- $\lambda(s)$ = instantaneous hazard rate at time s: how intensely the firm is "trying to default" at that instant.
- $\int_0^t \lambda(s)\,ds$ = accumulates total default intensity from 0 to t: a running tally of default pressure over time.
- λ_t = the instantaneous default rate: the higher it is, the faster survival probability decays. Default cannot be predicted in advance; it arrives as a jump.
The negative sign flips it into a survival measure, and the exponential converts accumulated intensity into a probability bounded between 0 and 1. A higher or longer-lasting λ means more accumulated intensity, a smaller exponent, and therefore a lower survival probability.

⋆ **Hazard Rate $\lambda(t)$.** The instantaneous conditional default rate: the fraction of the surviving population that defaults per unit time. Let $P(\tau > t) = V(t)$ **be the survival probability** (probability of surviving to t). The hazard rate governs how fast $V(t)$ decays:

$$V(t+\Delta t) - V(t) = -\lambda(t)\,V(t)\,\Delta t \quad\Rightarrow\quad \frac{dV(t)}{dt} = -\lambda(t)\,V(t)$$

Solving this ODE gives the **survival function**: $\boxed{P(\tau > t) = V(t) = e^{\int_0^t -\lambda(\tau)\,d\tau}}$

$\lambda(t)$ acts like a decay rate: the larger it is, the faster the pool of survivors shrinks. The integral accumulates total default intensity over time, and exponentiation converts it into a survival probability.

Cumulative default probability $Q(0,t)$ is simply the complement:

$$\boxed{Q(0,t) = 1 - V(t) = 1 - e^{\int_0^t -\lambda(\tau)\,d\tau}}$$

⋆ **Credit Spreads and PDs.** The risky discount factor can be decomposed into two scenarios: the firm survives, or it defaults and pays back only RR:

$$\tilde{P}(t,T) = P(t,T)\left[\underbrace{(1 - Q(t,T))}_{\text{P survive: receive 1}} \cdot 1 + \underbrace{Q(t,T)}_{\text{P default: receive } RR} \cdot RR\right]$$

Solving for the **RN default probability $Q(t,T)$:** (constant hazard rate case)

$$\boxed{Q(t,T) = 1 - e^{\lambda(T-t)} = \frac{1 - \tilde{P}(t,T)/P(t,T)}{1 - RR} = \frac{1 - e^{-Spd(T-t)}}{1 - RR}}$$

The spread does the work here: a wider credit spread implies a higher implied default probability, scaled down by $LGD = 1 - RR$ because bondholders recover something even in default.

Under a constant hazard rate, this simplifies to a clean approximation valid for short horizons or low spreads:

$$\boxed{\frac{Q(t,T)}{T-t} \approx \lambda \approx \frac{Spd}{1 - RR}}$$

Intuitively: the spread you earn per year is approximately the annual default rate times what you lose if default happens. If $LGD = 60\%$ and the spread is $120bp$, the implied $\lambda \approx 2\%$ per year.
A hazard rate of $\lambda = 2\%$ means the firm defaults at a rate of 2 per year: out of 100 similar firms, approximately 2 would default each year.

⋆ **Risk-Neutral PDs.** Implied from market prices of traded instruments (bond prices, CDS, equity). Forward-looking by construction. **Used for pricing credit derivatives.**

⋆ **Real-World PDs.** Estimated from historical default frequencies (e.g. Moody's data). Backward-looking; not forward-looking. Contain **risk premiums**. **Used for risk management, VaR, and scenario analysis.**

Risk-neutral PDs are always higher than real-world PDs: the gap is the credit risk premium that investors demand for bearing default risk. For AAA bonds the ratio is ∼44x; even for B-rated bonds it is ∼2x, reflecting that high-grade bonds carry disproportionately large premiums relative to their tiny actual default rates.

The credit spread decomposes as:

$$\underbrace{Spd}_{\text{total spread}} = \underbrace{\text{Historic Default Loss}}_{\text{actual expected loss}} + \underbrace{\text{Extra Risk Premium}}_{\text{compensation beyond expected loss}}$$

- Historic Default Loss $= \overline{PD} \times (1 - RR)$.
- Risk premium dominates the spread, especially for IG bonds where expected losses are tiny but spreads are substantial.

Extra risk premium exists for three reasons:

- **Illiquidity**: corporate bonds trade infrequently; investors demand extra yield to hold an asset they cannot easily exit
- **Systematic risk**: defaults cluster in recessions and cannot be diversified away; this non-diversifiable risk commands a premium
- **Skewed return distribution**: bond payoffs are highly asymmetric (capped upside, large downside on default), requiring additional compensation for bearing skew risk

6.2.2 Structural Model

Firm value V_t is modeled explicitly. Default occurs when V_t hits a barrier B_t, it is predictable given the value process, with no surprise jump to default.

- **Merton:** default only at maturity: $\Pr\{V_T \leq B_T\}$
- **Black–Cox:** default at first passage time: $\tau = \inf\{t : V_t \leq B_t\}$

⋆ **Merton Model of the Firm**

Corporate securities modeled as options on the firm's Total Value.

Every corporate security is a claim on the firm's total asset value V, where $V = E + B$ (equity plus market value of debt). Debt holders have first claim; equity holders get whatever is left, making equity a **call option** on firm assets with strike equal to the face value of debt D.

Think of shareholders as having bought the firm's assets and financed part of it with debt. At maturity, if $V > D$ they pay off the debt and keep the rest; if $V < D$ they hand the keys to creditors and walk away: exactly the payoff of a call option.

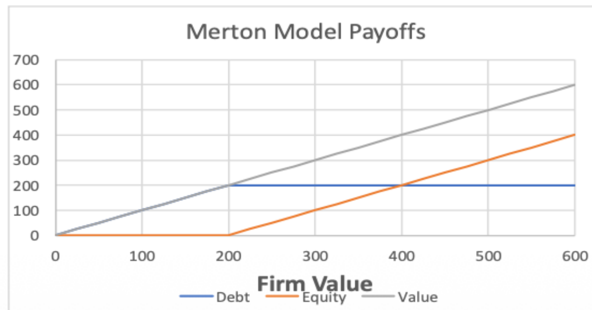

Firm asset value follows risk-neutral GBM, the standard Black-Scholes setup but applied to the whole firm rather than a stock price:

RN dynamics for firm's value (Equity + Debt): $\boxed{dV = rV\,dt + \sigma_V V\,dz}$

where V is total asset value, σ_V is asset volatility, r is the risk-free rate, and D is the face value of a single zero-coupon bond maturing at T.

Equity (E): Market capitalization of the firm = share price × shares outstanding. Represents the market value of shareholders' residual claim on firm assets after debt is repaid: $E = V - B$.

Payoff Structure at Maturity:

- **Equity**: residual claim, Call on V = $\boxed{E = \max(V - D,\, 0)}$ = Payoff : shareholders get nothing if firm value falls below debt face value
- **Debt**: first claim, payoff = $\min(V,\, D)$: creditors get paid in full if $V \geq D$, or recover whatever V is left if $V < D$

Via put-call parity $C - P = V - De^{-rT}$, the **market value of debt** is:

$$\boxed{B = V - E = De^{-rT} - P}$$

Debt is equivalent to a risk-free bond minus a put option that the firm implicitly sold to shareholders: shareholders have the right to "put" the assets back to creditors in default. The put value is the cost of default risk embedded in the debt.

Merton Model: Equations (1) & (2)

Since equity is a call on V with strike D, its value is given directly by Black-Scholes:

$$\boxed{E = VN(d_1) - De^{-rT}N(d_2)} \tag{1}$$

$$d_1 = \frac{\ln(V/D) + (r + \sigma_V^2/2)T}{\sigma_V\sqrt{T}}, \qquad d_2 = d_1 - \sigma_V\sqrt{T}$$

Here V plays the role of the stock price, D is the strike, and σ_V is asset volatility, all unobservable. What *is* observable is equity E and equity volatility σ_E. These are linked via the option delta relationship:

$$\boxed{\sigma_E = \frac{V}{E}\frac{\partial E}{\partial V}\,\sigma_V = N(d_1)\,\sigma_V\,\frac{V}{E}} \tag{2}$$

We observe σ_E and E (from stock market data) but cannot directly observe V or σ_V (firm assets don't trade). Equations **(1) and (2) are therefore solved simultaneously for the two unknowns V and σ_V.**

Deep in-the-money ($V \gg D$) $N(d_1) \approx 1$ and equity moves dollar-for-dollar with assets; **near default $N(d_1) \to 0$ and equity becomes insensitive to asset value since it is nearly worthless regardless.**

Once V and σ_V are recovered, all credit quantities follow PD, Credit Spread, Expected Loss.

Risk-neutral PD $\boxed{PD_{RN} = N(-d_2)}$

$N(-d_2)$ is the risk-neutral probability that $V_T < D$, i.e. firm value ends below the debt barrier at maturity.

Implied Credit Spread [s]**.** Backed out from the risky bond price:

$$\boxed{B = De^{-(r+s)T} \quad \Longrightarrow \quad s = -\frac{1}{T}\ln\left(\frac{B}{D}\right) - r}$$

The spread s is the extra yield above risk-free that makes the bond price equal to its market value $B = V - E$; it is fully determined by the firm's leverage and asset volatility, not assumed externally.

Real-world PD uses the actual asset drift μ instead of r:

$$\boxed{PD_{RW} = N(-d_{2,RW})}, \qquad d_{2,RW} = \frac{\ln(V/D) + (\mu - \sigma_V^2/2)T}{\sigma_V\sqrt{T}}$$

The two are linked by the market price of risk (Sharpe ratio of the firm's assets) $\boxed{\lambda_{V,M} = \frac{\mu - r}{\sigma_V}}$:

$$\boxed{PD_{RN} = \Phi\Big[\Phi^{-1}PD_{RW} + \lambda_{V,M}\sqrt{T}\Big]}$$

- $\Phi^{-1}PD_{RW}$ converts the real-world default probability into a standard normal z-score: how many standard deviations below zero firm asset value must fall for default to occur under the real-world measure.

- $\lambda_{V,M}\sqrt{T}$ adding this applies the Girsanov adjustment, shifting the drift from μ to r: since $r < \mu$ the firm grows more slowly under the risk-neutral measure, making default more likely, pushing the z-score toward default.

- **Outer** $\Phi[\cdot]$ converts back to a probability. Since $\lambda_{V,M} > 0$ always shifts the z-score toward default, $PD_{RN} > PD_{RW}$ always; the gap widens with a higher risk premium or longer horizon T.

The gap between PD_{RN} and PD_{RW} is exactly the credit risk premium.
The entire difference is driven by $\lambda_{V,M} = (\mu - r)/\sigma_V$, the risk premium investors demand for bearing asset risk, which is precisely why bond spreads always exceed pure expected loss.

Merton Model: Implementation Workflow

In practice V and σ_V are never directly observed; they are inferred from what is *observable: equity market cap and equity volatility. The workflow translates these market observables into a theoretical default probability.*

1. **Set horizon** T, typically 1 year; aggregate all debt obligations due by T to form D
2. **Observe inputs** E (equity market cap) and σ_E (equity volatility from options or historical returns)
3. **Solve equations (1) and (2) simultaneously** to recover unobservable asset value V and asset volatility σ_V
4. **Compute** $PD_{RN} = N(-d_2)$: this gives a *ranking* of default likelihood across firms; smooth σ_V estimates if noisy
5. **Map to real-world PD** using historical data and a market price of risk adjustment: the model's RN output is calibrated against observed default frequencies to produce an actionable real-world default probability

6.3 Correlated Defaults

⋆ **Default Correlation.** A measure of the tendency for firms to default around the same time. It matters in two contexts: portfolio credit risk management, where losses are driven by how many firms default together, and pricing of credit derivatives whose payoffs depend on multiple credits simultaneously.

A portfolio of 100 uncorrelated bonds with 1% PD each has very predictable losses: by the law of large numbers roughly 1 defaults. With high correlation, most periods see zero defaults but occasionally many default at once: the distribution has a fat tail. This tail risk is what correlation captures.

⋆ **Dependency Structure.** Defines how firm asset values co-vary. In the Merton structural framework, firm asset returns are correlated through a **common systematic factor**: the economy-wide shock that hits all firms simultaneously. The standard approach is the **Gaussian Copula**, which models joint defaults by assuming asset values follow a multivariate normal distribution.

⋆ **Gaussian Copula: Mapping to Normal Space**

Copula performs a percentile-to-percentile mapping: it takes each firm's marginal default probability $Q_i(t_i)$ and maps it to a standard normal random variable x_i:

$$\boxed{Q_i(t_i) = \Phi(x_i) \quad \iff \quad x_i = \Phi^{-1}[Q_i(t_i)]}$$

Each firm has its own default probability Q_i which can come from any distribution; the copula transforms it into a normal variable x_i so we can use the well-understood normal correlation structure to model joint defaults. The mapping preserves the marginal default probabilities exactly while imposing a normal dependency structure.

Critical Point $x_i = \Phi^{-1}(Q_i(t))$ is the threshold in normal space below which the firm defaults by time t.

Example: Default in year 1 corresponds to $x_i < -2.33$ (i.e. the bottom 1% of the normal distribution); default in year 2 but not year 1 corresponds to $-2.33 \leq x_i < -1.88$, and so on: each year's band maps exactly to the incremental default probability for that period.

$\star$ **Modeling Dependency: The One-Factor Model.** Each firm's latent normal variable x_i is decomposed into a shared systematic shock and a firm-specific idiosyncratic shock:

$$\boxed{x_i = \sqrt{\rho}\, F + \sqrt{1-\rho}\, Z_i \;\sim\; \Phi(0,1)}$$

- F: common **systematic factor**, standard normal; represents economy-wide shocks (recession, credit cycle) affecting all firms
- Z_i: **idiosyncratic shock** to firm i, standard normal, independent for each firms
- ρ: loading on the common factor; controls how much of each firm's variation is driven by the shared factor

$\sqrt{\rho}$ scales how much each firm is exposed to the common factor: a high ρ means firms move together (recession hits everyone); a low ρ means defaults are mostly idiosyncratic and independent. The weights $\sqrt{\rho}$ and $\sqrt{1-\rho}$ ensure x_i remains standard normal regardless of ρ.

Correlation between any two firms i and j follows directly:

$$\boxed{E[x_i x_j] = \sqrt{\rho}\sqrt{\rho}\,\mathrm{Var}(F) = \rho}$$

since F is the only shared component: idiosyncratic shocks Z_i and Z_j are independent and drop out. **So ρ is simultaneously the factor loading squared *and* the pairwise default correlation between any two firms.**

An alternative to the one-factor decomposition is a full Cholesky decomposition of the entire correlation matrix, which accommodates different pairwise correlations across firms.

Effect of Correlation on the Default Distribution Correlation fundamentally changes the *shape* of the loss distribution (number of defaults or equivalently the dollar amount lost across the portfolio over the horizon), not its mean.

With 10 firms each having $PD = 10\%$, the expected number of defaults is always 1 regardless of correlation. But:

- **Zero correlation ($\rho = 0$):** defaults are independent; distribution is tight around the mean, probability of 0 defaults is 35%, mass concentrated at 0 and 1
- **High correlation ($\rho = 100\%$):** all firms effectively share one fate: either all survive (90% of the time) or all default together (10% of the time), creating extreme bimodality
- **Intermediate correlation:** fat tails emerge: more mass at both zero defaults and many defaults simultaneously, less mass in the middle

This is the core insight of correlation risk: a credit portfolio manager who ignores correlation will severely underestimate tail losses. High correlation does not increase expected loss but dramatically increases the probability of catastrophic joint default scenarios: exactly the risk that materialised in 2008.

6.4 Credit Default Swaps (CDS)

⋆ **Credit Default Swap (CDS).** A contract where the **protection buyer** pays a periodic spread s (in bps per year) to the **protection seller** in exchange for a one-time payment of $100(1-RR)$ per \$100 notional if the reference entity defaults. The buyer is effectively buying insurance against default; the seller is taking on the credit risk.

A CDS lets you trade credit risk cleanly without touching the underlying bond, which is illiquid, hard to short, and may not be available. You can buy or sell protection on any credit name for any tenor and notional.

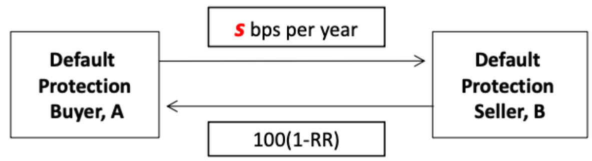

CDS Contract Specifications. The following are fixed and agreed at inception:

- **Reference Entity**: the firm whose default triggers the contract
- **Notional** N: the face value of protection being bought/sold (e.g. \$1 million); default payment $= N(1-RR)$
- **Contracted Spread** s^*: the fixed periodic fee the buyer pays per year on notional, set at inception and never changes; premium payment $= N \cdot s^* \cdot \Delta t$
- **Maturity** T: when the contract expires if no default occurs
- **Payment Frequency** Δt: length of each premium period (e.g. $\Delta t = 0.25$ for quarterly, $\Delta t = 1$ for annual)
- **Recovery Rate** RR: assumed recovery upon default; determines default payment $= N(1-RR)$

Everything else (the market value of the CDS, the implied Q, the par spread) fluctuates with market conditions. The contracted spread s^* is the only thing that stays fixed, which is why an existing CDS can gain or lose value after inception as credit conditions change.

Par CDS spread s is the spread s that makes the contract worth zero at inception, analogous to a par swap rate. Premiums are paid periodically until default or maturity; if default occurs mid-period, the buyer also pays an **accrual**: the pro-rated premium for the fraction of the period elapsed. Settlement is either cash (seller pays $100(1-RR)$) or physical delivery of the defaulted bond.

CDS curve (spreads across multiple maturities) is used to imply the full **term structure of risk-neutral default probabilities**, analogous to how the yield curve implies forward rates.

⋆ **CDS Valuation** From the protection buyer's perspective, the CDS value has three components:

$$\boxed{\text{CDS Value} = \text{Premiums} + \text{Accruals} - \text{Default Loss}}$$

- **Premiums**: PV of periodic spread payments s made while the reference entity survives. Each payment is weighted by the survival probability at that date
- **Accruals**: PV of the partial premium payment owed for the fraction of the period that elapses before default. Default is assumed to occur halfway through a period, so the accrual is $0.5 \times s \times PD_t$
- **Default Loss**: PV of the protection payment $LGD = (1-RR)$ made by the seller upon default, weighted by the unconditional default probability at each period

The CDS is zero-value at inception by construction: the spread s is set so that the PV of what the buyer pays (premiums + accruals) exactly equals the PV of what the seller pays (default loss). This is identical in spirit to an interest rate swap where fixed leg equals floating leg at inception. The general CDS valuation formulas applicable to any payment frequency, default probability structure, and default timing assumption are:

$$\boxed{\text{Premium Leg} = s\sum_t V(t)\cdot P(t)} \qquad \boxed{\text{Accrual Leg} = s\sum_t PD(t)\cdot \alpha \cdot P(t^*)}$$

$$\boxed{\text{Default Leg} = \sum_t PD(t)\cdot (1-RR)\cdot P(t^*)}$$

- $V(t)$: survival probability to payment date t
- $PD(t)$: unconditional default probability in period t
- $P(t)$: risk-free discount factor to payment date t
- $P(t^*)$: risk-free discount factor to assumed default date within the period
- α: fraction of period elapsed before default (e.g. 0.5 if midpoint assumption)

The breakeven spread solves premiums + accruals = default loss:

$$\boxed{s = \frac{\sum_t PD(t)\cdot (1-RR)\cdot P(t^*)}{\sum_t V(t)\cdot P(t) + \sum_t PD(t)\cdot \alpha \cdot P(t^*)}}$$

Note: $\boxed{PD(t) = Q_t \times V(t-1)}$ Q_t is the conditional default probability in period t given survival to $t-1$, while $V(t-1)$ is the probability of surviving to the start of that period. Their product gives the unconditional probability of defaulting in period t as seen from time 0.

Example: Each leg is computed by summing over all periods. For premiums (paid at end of year t if survived): $\text{Premium}_t = s \times \underbrace{(1-Q)^t}_{\text{survival prob}} \times \underbrace{e^{-rt}}_{\text{discount}}$

For default loss (default assumed at midpoint $t-0.5$): $\text{Loss}_t = \underbrace{PD_t}_{\text{uncond. default P}} \times (1-RR) \times e^{-r(t-0.5)}$

For accrual (half period's premium owed upon default): $\text{Accrual}_t = 0.5 \times s \times PD_t \times e^{-r(t-0.5)}$

Valuing an Existing CDS:

1. **Compute the premium + accrual annuity factor** using current market default probabilities and discount factors: this gives the PV of \$1 of spread per year (e.g. 4.113)
2. **Compute the default leg** using current market Q, RR, and discount factors: this gives the PV of expected default payments (e.g. 0.0511)
3. **Plug in the contracted spread** s^* (not the par spread) into the premium + accrual leg
4. **CDS Value** = (annuity factor $\times s^*$) $-$ default leg

The contracted spread s^* is fixed at inception and does not change. What changes over time is the par spread implied by current market conditions. If $s^* > s_{par}$, the buyer is paying more than the market currently charges for protection: the contract is in-the-money to the buyer. If $s^* < s_{par}$, the buyer is underpaying: the contract is in-the-money to the seller.

6.4.1 Alternative CDS's

⋆ **Binary CDS.** Same structure as a vanilla CDS but the **default payment is a fixed cash amount regardless of the recovery rate, $RR = 0$ by convention in the default leg,** so the payout is always \$1 per \$1 notional. The premium leg is identical to a vanilla CDS.

Since the default payoff is higher ($LGD = 1$ vs $1 - RR$), the breakeven spread is higher.

The binary CDS is more sensitive to the recovery rate assumption than a vanilla CDS. In a vanilla CDS, the $(1-RR)$ payoff and the $(1-RR)$ denominator in $Q \approx s/(1-RR)$ partially offset: changing RR moves both the calibrated Q and the payoff in opposite directions. In a binary CDS the payoff is fixed, so a change in RR only affects Q, making the spread much more sensitive.

Recovery Rate Sensitivity: CDS vs Binary CDS. In both contracts, the implied default intensity is calibrated from the spread:

$$Q \approx \frac{Spd}{1 - RR}$$

However the two contracts respond very differently to changes in RR:

- **Binary CDS**: highly sensitive to RR. The default payoff is a fixed \$1 regardless of RR, so changing RR only affects Q through calibration, with no offsetting effect on the payoff
- **Vanilla CDS**: relatively insensitive to RR near par. The payoff scales as $(1 - RR)$ while the calibrated $Q \sim Spd/(1-RR)$ scales inversely: the two effects partially cancel, leaving the vanilla CDS value approximately unchanged

⋆ **Forward Starting CDS.** A CDS that begins at a future date T_1 and runs until T_2. It is replicated by a long-short combination of standard CDS:

$$\boxed{\text{Forward CDS}(T_1, T_2) = \text{CDS}(0, T_2) - \text{CDS}(0, T_1)}$$

You are long protection from 0 to T_2 and short protection from 0 to T_1: the near-term legs cancel, leaving only the forward exposure from T_1 to T_2.

Key properties: the contract is cancelled if default occurs before T_1; it is more sensitive to the recovery rate assumption than a spot CDS (since default timing uncertainty is compressed into the forward window); and it is less liquid, carrying a liquidity premium in quoted spreads.

Uses include forward credit hedging (locking in protection cost for a future bond purchase) and relative value trades on the slope of the CDS term structure.

⋆ **CDS-Implied Risk-Neutral Default Probabilities** In practice, risk-neutral default probabilities Q are not directly observable; they must be **implied from observed CDS spreads**, exactly as implied volatility is backed out from option prices.

The CDS spread is the market's price for default risk. Inverting the pricing formula gives you the Q consistent with that price = RN default probability the market is embedding in the CDS.

Flat Q (single tenor): Assume a constant conditional default probability Q per year. The unconditional default and survival probabilities build recursively: to default in year t the firm must survive all prior years, so each period's default probability is Q scaled by the cumulative survival to that point.

Year t	Default Prob	Survival Prob
1	Q	$(1-Q)$
2	$Q(1-Q)$	$(1-Q)^2$
3	$Q(1-Q)^2$	$(1-Q)^3$
$\vdots$	$\vdots$	$\vdots$
t	$Q(1-Q)^{t-1}$	$(1-Q)^t$

Logic: Set premiums + accruals = default loss (all expressed as functions of Q and s) and solve numerically for the Q that matches the observed CDS spread s.
The result is the **implied risk-neutral default probability**: an analogue to implied vol, forward-looking and market-consistent.

Term Structure of Q (bootstrapping): To extract a full term structure of default probabilities, use CDS spreads at multiple maturities ($s_1, s_2, \ldots, s_5$ for 1Y through 5Y tenors) and allow a different Q_t for each period. Solved sequentially: Q_1 from the 1Y CDS, then Q_2 from the 2Y CDS holding Q_1 fixed, and so on, analogous to bootstrapping a discount curve from swap rates.

Year t	Default Prob	Survival Prob
1	Q_1	$(1-Q_1)$
2	$Q_2(1-Q_1)$	$(1-Q_1)(1-Q_2)$
3	$Q_3(1-Q_1)(1-Q_2)$	$(1-Q_1)(1-Q_2)(1-Q_3)$
$\vdots$	$\vdots$	$\vdots$

Logic: With 5 CDS tenors there are 5 unknown Q_t's and 5 equations, solved sequentially (bootstrap): Q_1 from the 1Y CDS, then Q_2 from the 2Y CDS holding Q_1 fixed, and so on. This gives the full **risk-neutral default probability term structure**, used as inputs for pricing more complex credit derivatives, directly analogous to how swap rates are bootstrapped into a discount curve.

⋆ **CDS Option.** The right to enter a **forward starting CDS** at a specified future date T at a pre-agreed strike spread K, provided no default has occurred before T. It expires worthless if the reference entity defaults before the option expiry.

- **Call**: right to *buy* protection at strike K;
  in-the-money when forward CDS spread $CDSF_{T,T_1} > K$ (credit has deteriorated)
- **Put**: right to *sell* protection at strike K;
  in-the-money when $CDSF_{T,T_1} < K$ (credit has improved)

Payoffs at expiry (if exercised):

$$\text{Call: } \max\{A(T,T_1)(CDSF_{T,T_1} - K),\, 0\}$$

$$\text{Put: } \max\{A(T,T_1)(K - CDSF_{T,T_1}),\, 0\}$$

$A(T,T_1)$ = **risky annuity factor**; the PV of \$1 per year on the forward CDS payment dates, discounted back to option expiry T. It converts a spread difference into a dollar payoff, analogous to a duration in interest rate swaptions.

$CDSF_{T,T_1}$ = **the forward CDS spread**; the breakeven spread agreed today on a protection contract starting at future date T and terminating at T_1. Determined by today's term structure of default probabilities, analogous to a forward interest rate.

The payoff structure is identical to an interest rate swaption: you are buying the right to enter a swap (here a CDS) at a fixed rate (here K). The risky annuity $A(T,T_1)$ plays the role of the swap's PV01.

Assuming the forward CDS spread is log-normally distributed, pricing follows a Black-Scholes-style formula:

$$\boxed{c = g_0\,[F_0\,N(d_1) - K\,N(d_2)]} \qquad \boxed{p = g_0\,[K\,N(-d_2) - F_0\,N(-d_1)]}$$

$$d_1 = \frac{\ln(F_0/K) + (r + \sigma^2/2)T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T}$$

- F_0: current forward CDS spread from T to T_1

- g_0: risky annuity value of \$1 per year on the $CDSF_T$ contract dates (today's value of $A(T, T_1)$)
- K: strike CDS spread
- σ: volatility of the forward CDS spread

This is exactly Black's model applied to CDS spreads: F_0 plays the role of the forward price, g_0 is the numeraire (like the annuity in swaption pricing), and σ is the implied spread vol. The log-normal assumption means spreads cannot go negative; alternative distributional assumptions (e.g. normal or shifted log-normal) are used when negative spreads are possible.

⋆ **Multi-Name (Basket) Credit Derivatives**

⋆ **First-to-Default (FTD).** A contract on a basket of N firms that pays off when the **first** entity defaults: payoff equals the loss on that first defaulted name, i.e. $N_i(1 - RR_i)$. Once triggered the contract terminates.

⋆ **nth-to-Default (NTD).** Generalisation of FTD: payoff triggers only when the nth firm in a basket of K names defaults ($n \leq K$), based on the loss of that nth defaulting entity. FTD is the special case $n = 1$.

FTD is the riskiest basket product: any single default wipes out the protection seller. NTD products allow investors to take on progressively more senior default risk, absorbing losses only after $n - 1$ names have already defaulted.

Valuation uses the same CDS framework but requires two additional inputs beyond individual Q's: the **default correlation** ρ between names (unobservable, must be implied from market prices) and the joint default distribution generated via Gaussian copula simulation.

6.5 Convertible / Callable Bond

⋆ **Convertible/Callable Bond.** A fixed income instrument with embedded options for both parties, a **fixed income investment with equity characteristics**:

- **Investor option (Conversion):** Right to convert the bond into a predetermined number of equity shares: the **conversion ratio**. Gives the investor upside participation if the stock rises
- **Issuer option (Call):** Right to redeem the bond at a predetermined **call price** Q_2. Acts as a mechanism to force conversion: if the issuer calls and conversion value exceeds the call price, the investor will convert rather than accept the call
- **Default:** Stock falls to zero; bondholder receives recovery $= RR \times$ face value

The conversion feature gives the investor upside in exchange for accepting a lower coupon than a straight bond. The call feature lets the issuer cap the investor's upside by forcing conversion once the stock has risen sufficiently.

⋆ **Equity Tree with Default** Valuation uses a trinomial tree where each node has **three possible states**: up, down, and default. The stock falls to zero upon default; the bond recovers $RR \times$ face value.

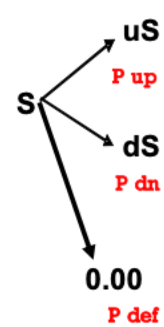

Three Branch Probabilities are calibrated to match the risk-free rate and equity volatility under the risk-neutral measure, with the default probability determined by the hazard rate λ:

$$\boxed{p_u = \frac{a - d\,e^{-\lambda \Delta t}}{u - d}} \qquad \boxed{p_d = \frac{u\,e^{-\lambda \Delta t} - a}{u - d}} \qquad \boxed{p_{def} = 1 - e^{-\lambda \Delta t}}$$

$$a = e^{r\Delta t}; \qquad u = e^{\sigma\sqrt{\Delta t}}; \qquad d = \frac{1}{u}$$

- a = risk-neutral growth factor over Δt; stock must grow at r on average in risk-neutral world
- u, d = up and down stock move factors driven by volatility σ
- p_{def} = probability of jumping to default, equal to the discrete hazard rate over Δt
- λ = default rate intensity (hazard rate)

The $e^{-\lambda\Delta t}$ factor scales the up and down probabilities: they do not sum to 1 on their own because the remaining probability mass $p_{def} = 1 - e^{-\lambda\Delta t}$ goes to the default branch. The standard CRR binomial tree is recovered when $\lambda = 0$.

⋆ **Optimal Exercise Decision**, At each node define three values:

- Q_1: continuation value: hold the bond (PV of future cash flows discounted back one step)
- Q_2: call value: what the issuer pays if they call the bond (the call price, e.g. \$113)
- Q_3: conversion value: what the investor receives if they convert (= conversion ratio × S)

The optimal exercise rules, both parties act in their own best interest:

$$\boxed{\text{Issuer calls if } Q_2 < Q_1} \qquad \boxed{\text{Investor converts if } Q_3 > \min(Q_1, Q_2)}$$

The issuer calls when it is cheaper to redeem at the call price than to let the bond continue, forcing the investor out at Q_2. The investor converts when equity value exceeds the best they could get by either holding or being called; they take the shares. This is analogous to the American option exercise decision, solved by backward induction through the tree.

Continuation value at each interior node: $\boxed{Q_1 = \left(p_u \cdot V_u + p_d \cdot V_d + p_{def} \cdot RR \cdot \text{Face}\right) e^{-r\Delta t}}$

where V_u and V_d are the bond values at the up and down child nodes (already reflecting optimal exercise at those nodes), and $RR \times$ Face is the recovery in the default state.

6.5.1 Collateralized Debt Obligation (CDO)

⋆ **CDO.** A structured credit product that takes a reference portfolio of credit assets and redistributes its losses across **tranches** with different attachment and detachment points: each tranche absorbs losses only after all junior tranches below it are wiped out, creating securities with distinct risk-return profiles from a single pool.

- **Cash CDO**: underlying assets are actual bonds or loans; cash flows from assets pay tranche coupons
- **CLO (Collateralized Loan Obligation)**: underlying assets are senior secured leveraged bank loans of BB/B speculative grade rating
- **Synthetic CDO** : no physical assets; the structure writes CDS contracts to gain exposure to default risk synthetically

The key innovation of a CDO is tranching: the same pool of risky assets can simultaneously produce a AAA-rated senior tranche (very unlikely to be hit by losses) and a high-yielding equity tranche (first to absorb any loss). Investors self-select based on their risk appetite.

⋆ **CDO Structure**

- **Attachment Point:** Portfolio loss level at which the tranche starts absorbing losses
  - **Lower Attachment Point L:** Cumulative portfolio loss level at which the tranche *begins* absorbing losses
  - **Upper Attachment Point H:** Loss level at which tranche principal is fully written down
- **Detachment Point:** Loss level at which the tranche is fully wiped out

Example: A mezzanine tranche with $L = 5\%$, $H = 15\%$ absorbs losses between 5% and 15% of portfolio notional. It is unaffected until cumulative losses exceed 5%, and wiped out once they reach 15%.

Losses flow through a waterfall: the equity tranche absorbs all losses first; once it is exhausted, losses hit the mezzanine; the senior tranche is only impaired after all junior tranches are wiped out.

Assets are transferred into a **Special Purpose Vehicle (SPV)**: a bankruptcy-remote legal entity that holds the collateral and issues the tranches. The SPV issues three classes of notes against a $100m portfolio:

Tranche	Typical Rating	Loss Absorption	Return
Equity (Jr)	Unrated	First losses (0% to ∼5%)	Residual (highest)
Mezzanine	A to BBB	Middle losses (∼5% to ∼25%)	SOFR + 250bp
Senior	AAA	Last losses (∼25% to 100%)	SOFR + 60bp

Ratings are assigned based on the structure and each tranche's loss potential, not the underlying asset ratings. The equity tranche is unrated and bears the most risk; the senior tranche is rated AAA despite the underlying assets being speculative grade.

6.5.2 Synthetic CDO

⋆ **Synthetic CDO.** Instead of holding physical bonds, the trust writes N CDS contracts as protection seller, gaining credit exposure synthetically without owning any assets. Investors provide cash which is posted as collateral (typically Treasuries); they earn SOFR plus the tranche spread. If defaults occur, collateral is drawn down to fund protection payments.

The key advantage over a cash CDO is flexibility and speed: you can create any desired credit exposure without needing to source actual bonds in the market. The trust sells protection on a reference portfolio of names, and tranches are defined on the losses that flow through.

Setup proceeds in three steps:
(1) trust writes N short CDS contracts targeting a desired rating, maturity, and correlation profile;
(2) tranche structure is defined iteratively: sizes, spreads, and attachment points are designed to achieve target ratings;
(3) sell tranches to investors, who fund their notes; proceeds are posted as collateral (Treasuries) to support CDS payments upon default

Synthetic CDO Definitions & Valuation

Tranche Definitions:

- **Lower Attachment Point α_L:** cumulative portfolio loss fraction at which the tranche begins absorbing losses
- **Upper Attachment Point α_H:** cumulative portfolio loss fraction at which the tranche is fully written down (zero principal remaining)

- N_j: tranche principal at payment date τ_j; $v(t)$: risk-free discount factor; s: tranche running spread
- n: number of reference entities; m: maturity index associated with τ_m

Tranche value has three components:

$$\boxed{\text{Tranche Value} = s(A+B) - C}$$

- A = PV of scheduled spread / premium payments
- B = PV of accrual payments on loss
- C = PV of the default losses to the tranche
- s = tranche running spread

⋆ **Dependency Modeling: Conditional Default Probability** Default correlation is modeled via the same one-factor Gaussian copula as before:

$$\boxed{x_i = \sqrt{\rho}\,F + \sqrt{1-\rho}\,Z_i \;\sim\; N(0,1)}$$

Firm i defaults by time T when $x_i < N^{-1}[Q_i(T)]$.
Conditioning on the common factor F, substituting x_i and solving for Z_i:

$$\boxed{Q_i(T|F) = N\left(\frac{N^{-1}[Q_i(T)] - \sqrt{\rho}\,F}{\sqrt{1-\rho}}\right)}$$

Given F, defaults across firms become independent: each firm's conditional default probability depends only on its own Q_i and the realized systematic factor. This is the key tractability result: conditioning on F breaks the correlation and allows the joint default distribution to be computed as a product of independent binomials.

⋆ **Tranche Probabilities & Losses:** Conditional on F, the probability of exactly k defaults from n names by time T follows a binomial (since defaults are independent given F, with identical marginal probability $Q(T|F)$):

$$\boxed{Pr(k,\, t\,|\,F) = \binom{n}{k}\, Q(t|F)^k\, [1 - Q(t|F)]^{n-k}}$$

The number of defaults needed to hit the lower and upper attachment points (assuming equal notional and constant RR):

$$\boxed{k_L = \frac{n\,\alpha_L}{1-RR}, \qquad k_H = \frac{n\,\alpha_H}{1-RR}}$$

k_L is the highest number of defaults for which the tranche suffers no loss; k_H is the lowest number for which the tranche is fully wiped out. Between k_L and k_H the tranche is partially impaired.

Expected tranche principal at time τ_j conditional on F is:

$$\boxed{E[N_j(F)] = \underbrace{\sum_{k=0}^{k_L-1} Pr(k,\tau_j|F)}_{\substack{\text{prob of } <k_L \text{ defaults} \\ \text{tranche fully intact, principal=1}}} + \underbrace{\sum_{k=k_L}^{k_H-1} Pr(k,\tau_j|F)}_{\substack{\text{prob of partial impairment} \\ k_L \le k < k_H \text{ defaults}}} \cdot \underbrace{\frac{\alpha_H - k(1-RR)}{\alpha_H - \alpha_L}}_{\substack{\text{fraction of principal} \\ \text{surviving after } k \text{ defaults}}}}$$

First sum covers scenarios where fewer than k_L defaults occur, tranche is fully intact, principal $= 1$. The second sum covers partial impairment scenarios ($k_L \le k < k_H$) where the fraction $\frac{\alpha_H - k(1-RR)}{\alpha_H - \alpha_L}$ of principal survives.

Tranche Valuation Formulas

$A(F)$: PV of scheduled spread payments:

$$\boxed{A(F) = \sum_{j=1}^{m} (\tau_j - \tau_{j-1})\, N_j(F)\, \upsilon(\tau_j)}$$

$B(F)$: PV of accrual payment on losses (default assumed at period midpoint):

$$\boxed{B(F) = \sum_{j=1}^{m} 0.5(\tau_j - \tau_{j-1})\, (N_{j-1}(F) - N_j(F))\, \upsilon(0.5\tau_{j-1} + 0.5\tau_j)}$$

$C(F)$: PV of default losses paid by the tranche:

$$\boxed{C(F) = \sum_{j=1}^{m} [N_{j-1}(F) - N_j(F)]\, \upsilon(0.5\tau_{j-1} + 0.5\tau_j)}$$

$N_{j-1}(F) - N_j(F)$ is the loss of tranche principal between consecutive payment dates: it is positive only when defaults hit the tranche in that period. B and C are discounted at the midpoint of each period, consistent with the assumption that losses occur at the midpoint.

The **breakeven tranche spread** sets value to zero and solves:

$$\boxed{s = \frac{E[C(F)]}{E[A(F)] + E[B(F)]}}$$

Since all quantities are conditional on F, the expectation is taken over the distribution of F: either by discretizing F and using Gaussian quadrature, or via Monte Carlo for complex structures.

Implied Tranche Correlation

Correlation ρ is the key unobservable input: it cannot be directly observed and must be implied from traded tranche prices, exactly as implied volatility is backed out from option prices.

⋆ **Compound Correlation.** The ρ implied from the market price of a single individual tranche using the one-factor Gaussian copula. Different tranches imply different correlations, and in practice compound correlation is not stable across tranches and is not arbitrage-free, which limits its usefulness.

⋆ **Base Correlation.** The ρ that correctly prices the $[0, X\%]$ tranche (always anchored at zero loss) consistently with the market, where $X\%$ is the detachment point. Any tranche $[\alpha_L, \alpha_H]$ is priced as the difference between two base tranches:

$$\text{Tranche}[\alpha_L, \alpha_H] = \text{Base}[0, \alpha_H] - \text{Base}[0, \alpha_L]$$

Base correlation is monotonically increasing in the detachment point: more senior tranches require higher implied correlation to match market prices. It is more stable and internally consistent than compound correlation, making it the market standard.

The correlation smile in CDO tranches is analogous to the volatility smile in equity options: the market assigns different implied correlations to different tranches, reflecting that the Gaussian copula is an imperfect model. Just as traders use implied vol as a quoting convention rather than a belief about the true distribution, they use implied correlation as a way to quote and compare tranche prices.

Chapter 7

Applications

7.1 Principal Component Analysis (PCA)

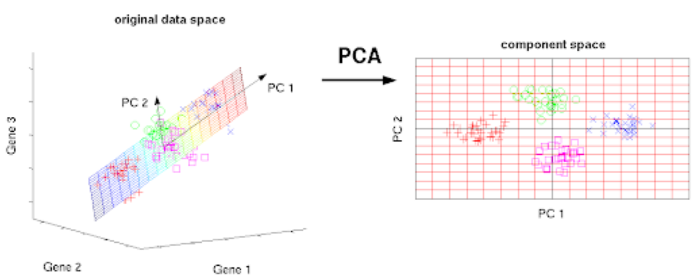

⋆ **Core Idea.** Data in K correlated variables often lies on a much lower-dimensional **subspace**, a flat plane inside that space. **PCA** finds a new coordinate system aligned with the directions the data actually varies in, re-expresses every point in those coordinates, and discards directions with negligible variance:

$$\underbrace{K\text{-dimensional correlated variables}}_{\text{original space}} \xrightarrow{\Sigma=\Gamma\Delta\Gamma'} \underbrace{m\text{-dimensional uncorrelated PCs}}_{\text{component space}}, \quad m \ll K$$

For yield curves: instead of modeling 10 tenors separately, PCA finds that almost all movement is explained by 3 modes, level shift, slope tilt, and curvature bend. PCA is a ***risk management and scenario design tool****, not an arbitrage-free pricing model.*

⋆ **Step 1: Data Transformation and Standardization.** Raw rates $X(t)$ must first be converted into changes $Z(t)$. Three conventions:

- **Absolute:** $Z(t) = X(t) - X(t-1)$, shocks in rate-point units, symmetric confidence intervals
- **Log:** $Z(t) = \ln(X(t)/X(t-1))$, proportional shocks, requires strictly positive rates
- **Displaced log:** $Z(t) = \ln((X(t)+\delta)/(X(t-1)+\delta))$, $\delta > 0$ controls how negative rates can go

Then **standardize** tenor by tenor: $\boxed{\tilde{Z}(t) = (Z(t) - \mu_T)/\sigma_T}$ where $\mu_T = \frac{1}{T}\sum_t Z(t)$ is the **sample mean** and $\sigma_T = \sqrt{\frac{1}{T-1}\sum_t (Z(t) - \mu_T)^2}$ is the **sample standard deviation** over historical window T.

Standardization puts all tenors on equal footing. Without it, noisier short-end tenors dominate PC1 purely due to scale, not economic importance. Absolute differences give symmetric up/down shocks; log-based transforms produce asymmetric bounds since a proportional shock is larger when rates are high.

⋆ **Step 2: Eigendecomposition of Σ.** Let Σ be the **covariance matrix** of $\tilde{Z}(t)$. Large off-diagonal entries mean tenors co-move, making individual tenor modeling redundant. Decompose:

$$\boxed{\Sigma = \Gamma\Delta\Gamma'} \qquad \text{equivalently:} \quad \Sigma\,\gamma_i = \lambda_i\,\gamma_i$$

- $\Gamma = [\gamma_1 \mid \gamma_2 \mid \cdots \mid \gamma_K]$: matrix of **eigenvectors**, orthonormal ($\Gamma\Gamma' = I$, $\gamma_i'\gamma_i = 1$). Each column γ_i is the **loading pattern** for PCi, i.e. how each tenor participates in that mode.
- Δ: diagonal matrix of **eigenvalues** $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_K \geq 0$, measuring variance along each axis.
- γ_1: direction of **maximum variance**. γ_2: maximum remaining variance subject to $\gamma_2 \perp \gamma_1$. Each subsequent γ_i captures the next independent direction.

$$\boxed{V = \sum_{i=1}^{K} \lambda_i} \quad \boxed{\text{Percent explained by PC}(i) = \frac{\lambda_i}{V}}$$

For a **correlation matrix**: $V = K$, so % explained $= \lambda_i/K$ directly.

Orthogonality is the key property: each PC captures a completely independent mode of variation. In interest rates, PC1 (level), PC2 (slope), PC3 (curvature) do not overlap, so you can shock them independently without double-counting risk. The eigendecomposition is solved numerically (e.g. via SVD) for $K > 2$, there is no closed form.

⋆ **Step 3: Scores, Reconstruction, and Stress Scenarios.**

There are two distinct uses of PCA once Γ and Δ are obtained: **(A) backward-looking reconstruction** from observed data, and **(B) forward-looking stress scenario generation.**

(A) PC Score and Reconstruction (historical/backward). Project each day's standardized rate changes onto the PC axes to get a **score**, one scalar per PC per day: $\boxed{\text{score}_{i,t} = \gamma_i'\tilde{Z}(t) = \sum_{k=1}^{K} \gamma_{ik}\,\tilde{Z}_k(t)}$

The score tells you how hard factor i fired on day t. Loadings γ_{ik} are fixed from the eigendecomposition; scores change every observation. Large $|score_{1,t}|$: mostly a level day. Large $|score_{2,t}|$: mostly a slope day. Typically all three are non-zero simultaneously.

Reconstruct the original rate change at tenor k by summing contributions from all m PCs:

$$\boxed{Z_k(t) \approx \mu_T + \sigma_T \sum_{i=1}^{m} \text{score}_{i,t} \cdot \gamma_{ik}}$$

All three PCs fire simultaneously, each contributing an independent piece of the move:

$$Z_k(t) \approx \underbrace{\text{score}_{1,t} \cdot \gamma_{1k}}_{\text{level shift}} + \underbrace{\text{score}_{2,t} \cdot \gamma_{2k}}_{\text{slope change}} + \underbrace{\text{score}_{3,t} \cdot \gamma_{3k}}_{\text{curvature change}}$$

With $m = K$ reconstruction is exact. With $m = 3$ it captures $> 99\%$ of the move. The $\mu_T + \sigma_T(\cdots)$ simply unstandardizes back to rate-point units.

(B) Stress Scenario Projection (forward-looking). Replace the observed score with its α-percentile under normality. Since $\text{score}_{1,t} \sim N(0, \lambda_1\,dt)$, its standard deviation is $\sqrt{\lambda_1\,dt}$, so the α-percentile score is $\pm\sqrt{\lambda_1}\,\Phi^{-1}(\alpha)\sqrt{dt}$.

The **standardized PC1 shock** applied across all tenors is: $\boxed{\pm\,\gamma_{1k}\sqrt{\lambda_1}\,\Phi^{-1}(\alpha)\,\sqrt{dt}}$

Translating back to the original rate level at tenor k: $\boxed{X_t^k = X_{t-1}^k + \mu_T + \sigma_T\,\gamma_{1k}\underbrace{\sqrt{\lambda_1}\,\Phi^{-1}(\alpha)\,\sqrt{dt}}_{\alpha\text{-percentile PC1 score}}}$

This is the reconstruction formula with $m = 1$ and the score fixed at its confidence-level value, not a separate formula. Setting PC2/PC3 scores to zero isolates a pure level-shift stress. Since PCs are orthogonal, PC2 and PC3 stresses can be added on top independently without any interaction.

⋆ **Example: US Treasury PCA (1953–2024), $K = 10$ tenors.**

	3 mo	6 mo	1 yr	2 yr	3 yr	5 yr	7 yr	10 yr	20 yr	30 yr
3 mo	1.00	0.99	0.96	0.88	0.83	0.75	0.71	0.69	0.65	0.66
6 mo	0.99	1.00	0.99	0.94	0.89	0.80	0.75	0.72	0.67	0.68
1 yr	0.96	0.99	1.00	0.97	0.93	0.85	0.79	0.75	0.69	0.69
2 yr	0.88	0.94	0.97	1.00	0.99	0.92	0.86	0.81	0.74	0.72
3 yr	0.83	0.89	0.93	0.99	1.00	0.97	0.91	0.86	0.79	0.76
5 yr	0.75	0.80	0.85	0.92	0.97	1.00	0.99	0.96	0.90	0.87
7 yr	0.71	0.75	0.79	0.86	0.91	0.99	1.00	0.99	0.96	0.93
10 yr	0.69	0.72	0.75	0.81	0.86	0.96	0.99	1.00	0.99	0.97
20 yr	0.65	0.67	0.69	0.74	0.79	0.90	0.96	0.99	1.00	0.99
30 yr	0.66	0.68	0.69	0.72	0.76	0.87	0.93	0.97	0.99	1.00
Mean	0.0039	0.0038	0.0035	0.0030	0.0026	0.0019	0.0014	0.0009	0.0004	0.0001
Std	0.0130	0.0129	0.0125	0.0113	0.0106	0.0099	0.0096	0.0092	0.0089	0.0082

	PC1	PC2	PC3	PC4	PC5	PC6	PC7	PC8	PC9	PC10
3 mo	0.30	0.40	0.48	0.47	-0.38	-0.03	0.29	-0.26	0.01	-0.02
6 mo	0.31	0.39	0.26	0.00	0.29	-0.23	-0.38	0.61	0.18	0.03
1 yr	0.32	0.36	0.01	-0.31	0.47	0.30	-0.05	-0.41	-0.45	-0.04
2 yr	0.32	0.24	-0.33	-0.39	-0.20	0.23	0.24	-0.04	0.61	0.24
3 yr	0.33	0.11	-0.45	-0.12	-0.37	-0.32	0.10	0.21	-0.37	-0.48
5 yr	0.33	-0.12	-0.36	0.33	-0.02	-0.28	-0.26	-0.17	-0.18	0.66
7 yr	0.33	-0.25	-0.18	0.41	0.23	0.17	-0.29	-0.24	0.39	-0.51
10 yr	0.32	-0.33	0.02	0.21	0.08	0.55	0.35	0.49	-0.24	0.13
20 yr	0.31	-0.39	0.24	-0.22	0.34	-0.53	0.46	-0.11	0.12	-0.01
30 yr	0.30	-0.38	0.41	-0.39	-0.45	0.13	-0.46	-0.08	-0.07	0.01
Eigenvalue	8.576	1.047	0.273	0.030	0.005	0.002	0.002	0.001	0.000	0.000
% Expained	86%	11%	3%	0%	0%	0%	0%	0%	0%	0%

The correlation matrix (left) shows all tenors highly correlated (adjacent ≈ 0.99, 3M vs 30Y $= 0.66$) with short-end volatility higher (Std 0.0130 at 3M vs 0.0082 at 30Y), confirming standardization is needed. The eigenvector table (right) **is** Γ: rows are tenors, columns are PCs, entries are loadings γ_{ik}.

- **PC1 (Level):** $\gamma_{1k} \approx +0.31$ uniformly across all tenors. Score > 0: all rates rise (**parallel shift up**); score < 0: shift down. Explains 86% ($\lambda_1 = 8.576$). Dominant because all rates mostly move together.
- **PC2 (Slope):** γ_{2k} decreases monotonically from $+0.40$ (3M) to -0.39 (30Y). Score > 0: short up, long down (**flattening**); score < 0: **steepening**. Explains 11% ($\lambda_2 = 1.047$).
- **PC3 (Curvature):** γ_{3k} positive at wings, negative at belly. Score > 0: wings rise, belly falls (**butterfly**); score < 0: belly rises relative to wings. Explains 3% ($\lambda_3 = 0.273$).

These three modes correspond exactly to parallel shift, twist, and butterfly used in stress testing, recovered purely from data without imposing them. Together $> 99\%$ of variance is explained, confirming 3 factors are sufficient for term structure modeling.

PC1 Confidence Bounds. The confidence bounds plot shows **stressed rate levels** at each tenor (y-axis in decimal, e.g. $0.08 = 8\%$): base $\approx 4.7\%$ current rates, up/down shocks from the PC1 projection formula. Under **absolute difference** transformation: shock magnitude is identical up and down $\Rightarrow$ **symmetric bounds**. Under **log** or **displaced log**: shock scales with rate level $\Rightarrow$ **asymmetric bounds**. "'

7.2 Equity Volatility Modeling

⋆ **Equity Volatility.** Three key empirical observations motivate all models below:

- **Non-constant:** σ varies substantially over time, Black-Scholes constant vol assumption fails in practice.
- **State-dependent:** volatility spikes in distressed economic environments (1987 crash $\approx 300\%$ annualized, 2008–2009 $\approx 180\%$, COVID $\approx 200\%$) and is low in calm periods ($\approx$ 10–20%).
- **Volatility clustering:** large moves tend to follow large moves, small moves follow small moves. Today's vol is informative about tomorrow's vol.

Volatility σ measures the **magnitude** of price swings, not their direction; it is a standard deviation, always positive. A 2% up move and a 2% down move contribute equally to σ. However, empirically vol reacts asymmetrically: **downside moves cause larger vol spikes than upside moves of the same size (leverage effect). Falling prices $\Rightarrow$ higher firm leverage $\Rightarrow$ equity riskier $\Rightarrow$ σ jumps. Fear amplifies vol in sell-offs more than euphoria does in rallies.**

Four modeling approaches, in increasing sophistication:

- **Equal weighted:** standard historical volatility, all past observations get the same weight
- **Time-decay weighted (EWMA):** recent observations get more weight, old ones decay exponentially
- **Conditional on past volatility (ARCH/GARCH):** today's variance depends explicitly on past squared returns and past variances
- **Stochastic volatility:** volatility is itself a random process (e.g. Heston)

Time scaling. Variance is additive over independent periods, so volatility scales with $\sqrt{n}$:

$$\boxed{\sigma_{n\text{-Day}} = \sqrt{n}\,\sigma_{1\text{-Day}}} \qquad \boxed{\sigma_{\text{Annual}} = \sqrt{252}\,\sigma_{\text{Daily}}}$$

252 trading days assumed per year. This applies under i.i.d. returns, if returns are autocorrelated (as in vol clustering) the $\sqrt{n}$ rule is an approximation.

Note: Variance additivity $\sigma_n^2 = n\sigma_1^2$ holds exactly only under **i.i.d. returns**: independent, identically distributed with no autocorrelation. Under **vol clustering**, daily variances are correlated so the simple sum over/understates true n-period variance depending on the structure. In that case the correct annual variance is: $\boxed{\sigma^2_{\text{Annual}} = \sum_{t=1}^{252} \sigma_t^2}$

a sum of **time-varying** daily variances, not 252× a single number. The $\sqrt{252}\,\sigma_{\text{Daily}}$ shortcut then approximates this by assuming today's vol σ_t persists unchanged for the full year.

⋆ **Stock Returns.** Two conventions for the return u_t at time t:

$$\boxed{u_t = \frac{S_t}{S_{t-1}} - 1 \quad \text{(discrete)}} \qquad \boxed{u_t = \ln\left(\frac{S_t}{S_{t-1}}\right) \quad \text{(continuous/log)}}$$

For daily data the two are approximately equal.

The underlying process assumed is: $\boxed{u_t = \mu\,dt + \sigma_t\,dz}$

where μ is the **mean return**, σ is the **return volatility**, and dz is a **random innovation** ($dz \sim N(0, dt)$).

⋆ **Standard (Equal-Weighted) Volatility.**

$\sigma_{n,m}$ = volatility estimate at time n using m past observations.

$$\boxed{\sigma^2_{n,m} = \frac{1}{m-1}\sum_{i=1}^{m}(u_{n-i} - \bar{u})^2} \qquad \bar{u} = \frac{1}{m}\sum_{i=1}^{m} u_{n-i}$$

This is the **standard sample variance** formula applied to returns. Every past observation receives weight $1/m$ regardless of age. **Key choice: the lookback window m, longer m gives smoother but slower-reacting estimates.**

By the **Law of Large Numbers**, as $m \to \infty$ the equal-weighted estimator converges to the true **unconditional long-run variance:** $\frac{1}{m}\sum_{i=1}^{m} u_{n-i}^2 \xrightarrow{m\to\infty} E[u_t^2] = \sigma^2$

This is the **only estimator with this property**, equal weights are the only weighting scheme that samples the full unconditional distribution of returns fairly. EWMA and GARCH capture **local, regime-specific** variance If you want the true long-run average variance across all market regimes, equal-weighted historical vol over a long window is the only statistically honest answer. The main weakness is the **ghost feature**: a large return contributes equally for m days then drops out of the window suddenly, causing an artificial jump in the estimate.

⋆ **Non-Constant (Weighted) Volatility.** Generalize by assigning a weight w_i to each past observation:

$$\boxed{\sigma_n^2 = \sum_{i=1}^{m} w_i\, u_{n-i}^2} \qquad \text{subject to: } \sum_{i=1}^{m} w_i = 1, \quad w_i \geq 0$$

Note the mean return $\bar{u}$ is dropped (assumed zero for daily data, drift $\mu\, dt$ is negligible relative to $\sigma\, dz$ over short horizons). The equal-weighted model is the special case $w_i = 1/m\ \forall\, i$.

⋆ **EWMA (Exponentially Weighted Moving Average).** Assign weights that decline geometrically at rate $\lambda \in (0, 1)$ as you move back in time:

$$\boxed{\sigma_n^2 = \lambda\, \sigma_{n-1}^2 + (1 - \lambda)\, u_{n-1}^2}$$

Although the formula appears to use only one day, σ_{n-1}^2 already contains all past history compressed inside it. Recursively substituting:

$$\sigma_n^2 = (1-\lambda)u_{n-1}^2 + \lambda(1-\lambda)u_{n-2}^2 + \lambda^2(1-\lambda)u_{n-3}^2 + \cdots + \lambda^{m-1}(1-\lambda)u_{n-m}^2 + \lambda^m \sigma_{n-m}^2 \Rightarrow$$

$$\Rightarrow \quad \sigma_n^2 = (1-\lambda)\sum_{i=1}^{\infty} \lambda^{i-1} u_{n-i}^2 \qquad \text{where} \quad \sum_{i=1}^{\infty}(1-\lambda)\lambda^{i-1} = 1$$

so observation i periods ago gets weight $(1-\lambda)\lambda^{i-1}$, declining exponentially. Think of σ_{n-1}^2 as a **running balance**: all previous squared returns are already baked in with their exponential weights, you just add today's new information on top.

- **Infinite memory:** all past observations enter, weights never reach exactly zero
- **Recency:** recent returns weighted more heavily than old ones
- **One parameter:** λ fully controls the decay speed
- **No mean reversion:** if vol spikes it stays high until new small returns drag it down, there is no long-run anchor pulling variance back

Today's variance is a blend of yesterday's variance estimate (λ weight) and yesterday's squared return ($(1-\lambda)$ weight). Large $\lambda \Rightarrow$ slow decay $\Rightarrow$ long memory. Small $\lambda \Rightarrow$ fast decay $\Rightarrow$ quick reaction to new shocks.

EWMA Half-Life. The **half-life** HL is the number of periods until an observation has half the weight of the most recent observation, solving $\frac{1}{2} = \lambda^{HL}$: $\boxed{HL = -\frac{\ln(2)}{\ln(\lambda)}}$

λ	HL (periods)
0.900	7
0.950	14
0.980	34
0.990	69
0.995	138

*Half-life is a **calibration device**: λ alone is abstract, half-life translates it into something economically interpretable, how many days before a past shock loses half its influence on today's variance estimate. In practice you choose HL first based on your view of how long vol regimes persist in your asset class, then back out $\lambda = e^{-\ln(2)/HL}$. For example: equity vol mean reverts quickly so short $HL \approx$ 14–22 days ($\lambda \approx$ 0.95–0.97); credit spreads persist longer so $HL \approx$ 69–138 days ($\lambda \approx$ 0.99–0.995).*
$\lambda = 0.90$ implies a half-life of 7 days: a shock from one week ago has half the weight of yesterday's shock in today's variance estimate.

Estimating λ:

- **Hardcode (most common):** RiskMetrics standard: $\lambda = 0.94$ for daily, $\lambda = 0.97$ for monthly. Requires zero estimation and works well empirically. Most practitioners use this.
- **Match target half-life:** if you have an economic view that vol shocks should decay over HL periods: $\lambda = e^{-\ln(2)/HL}$ e.g. 1-month half-life ($HL = 22$ days) $\Rightarrow \lambda = 0.969$. You are encoding intuition about regime persistence, not fitting data.
- **MLE / minimize forecast error:** choose λ to minimize out-of-sample forecast error over historical data: $\hat{\lambda} = \arg\min_\lambda \sum_t (u_t^2 - \sigma_t^2(\lambda))^2$ Most statistically rigorous but improvement over $\lambda = 0.94$ is often small in practice.

⋆ **ARCH(p): Auto-Regressive Conditionally Heteroscedastic.** Generalize EWMA to include a **long-run variance anchor** V_L plus p past squared returns with finite memory:

$$\boxed{\sigma_n^2 = \gamma V_L + \sum_{i=1}^{p} \alpha_i u_{n-i}^2} \qquad \text{subject to: } \gamma + \sum_{i=1}^{p} \alpha_i = 1$$

- V_L: long-run variance rate (fixed level vol mean-reverts toward)
- α_i: weight on squared return i periods ago
- γ: weight pulling variance back to V_L
- Finite lag p (vs EWMA's infinite memory)

ARCH adds two things EWMA lacks: a ***long-run anchor*** *V_L that variance is always pulled back toward, and* ***finite memory*** *(only p recent shocks enter directly). Without the σ_{n-j}^2 recursion term, ARCH needs many lags to capture realistic vol persistence, which makes estimation hard, this is what GARCH fixes.*

⋆ **GARCH(p,q): Generalized ARCH.** Add lagged conditional variances back into the recursion:

$$\boxed{\sigma_n^2 = \gamma V_L + \sum_{i=1}^{p} \alpha_i u_{n-i}^2 + \sum_{j=1}^{q} \beta_j \sigma_{n-j}^2} \qquad \gamma + \sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j = 1$$

Three components of today's variance:

- σ_{n-1}^2: **previous conditional variance** (smoothing/persistence)
- u_{n-1}^2: **previous variance realization** (shock/innovation)
- V_L: **long-run average variance rate** (mean-reversion target)

Persistence from $\beta\sigma_{n-1}^2$, fresh information from αu_{n-1}^2, and gravity toward V_L from γV_L. The recursion in σ_{n-j}^2 is what lets a small number of parameters capture long persistence, unlike pure ARCH which would need many lags.

⋆ **GARCH(1,1): The Standard Parametrization.** Set $\omega = \gamma V_L$ to absorb the long-run term:

$$\boxed{\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2} \qquad \alpha \geq 0,\ \beta \geq 0,\ \alpha + \beta < 1$$

$$\boxed{V_L = \frac{\omega}{1 - \alpha - \beta}}$$

Three parameters (ω, α, β), three components, one lag of each. This is the most popular vol model in finance, simple enough to estimate, rich enough to capture persistence and mean reversion. The constraint $\alpha + \beta < 1$ is what guarantees a finite long-run variance, if it fails, vol can drift to infinity.

⋆ **GARCH(1,1) as a Generalization of EWMA.** Recursively substitute σ_{n-j}^2:

$$\sigma_n^2 = \omega + \alpha u_{n-1}^2 + \beta(\omega + \alpha u_{n-2}^2 + \beta\sigma_{n-2}^2) = \omega + \beta\omega + \alpha u_{n-1}^2 + \alpha\beta u_{n-2}^2 + \alpha\beta^2 u_{n-3}^2 + \cdots$$

Setting $\omega = 0$, $\beta = \lambda$, $\alpha = 1 - \lambda$ recovers EWMA exactly.

EWMA is a constrained GARCH(1,1) with no long-run anchor. *GARCH adds the ω term that gives variance somewhere to revert to, and lets $\alpha + \beta < 1$ so old shocks decay rather than persist forever. Everything EWMA does, GARCH does too, plus mean reversion.*

⋆ **Worked Example: GARCH(1,1). Estimated model output:**

$$\sigma_n^2 = 0.000002 + 0.13\, u_{n-1}^2 + 0.86\, \sigma_{n-1}^2$$

- $\alpha = 0.13$: shock weight, how much yesterday's return moves today's variance
- $\beta = 0.86$: persistence, how much yesterday's variance carries forward
- $\omega = 0.000002$: drift toward long-run level
- $\gamma = 1 - \alpha - \beta = 0.01$: long-run anchor weight
- $V_L = \omega/\gamma = 0.0002$: long-run daily variance

Long-run daily volatility: $\sqrt{0.0002} = 1.4\%$. Hence Long-run annual volatility: $\sqrt{252} \times 1.4\% = 22.4\%$.

⋆ **Estimating GARCH by Maximum Likelihood (MLE).** Assume returns u_i are normally distributed with mean zero and variance $v_i = \sigma_i^2$. The likelihood of observing the full return series is the product of normal densities:

$$\boxed{\mathcal{L}(\omega, \alpha, \beta) = \prod_{i=1}^{m} \frac{1}{\sqrt{2\pi v_i}} \exp\left(\frac{-u_i^2}{2v_i}\right)}$$

Take logs and drop constants. Maximizing the log-likelihood is equivalent to:

$$\boxed{\max_{\omega,\alpha,\beta} \sum_{i=1}^{m} \left[-\ln(v_i) - \frac{u_i^2}{v_i}\right]} \quad \text{where} \quad v_i = \omega + \alpha u_{i-1}^2 + \beta\sigma_{i-1}^2$$

Each term penalizes two things: $-\ln(v_i)$ punishes overestimating variance (saying it's huge when nothing happens), and $-u_i^2/v_i$ punishes underestimating (saying it's tiny when a big return occurs). Balance gives the best fit. The variance v_i at each step comes from the GARCH recursion itself, which is why parameters (ω, α, β) enter the optimization.

⋆ **Practical GARCH(p,q) Estimation Workflow.** General case extends the recursion to p lags of squared returns and q lags of conditional variance. Same workflow as GARCH(1,1), but with more parameters and a longer initialization.

Step 1: Prepare data. Compute daily log returns from price history: $u_t = \ln(S_t/S_{t-1})$ Drop the mean ($\bar{u} \approx 0$ for daily) so u_t enters directly as u_t^2 for daily Var estimation.

Step 2: Initialize parameters and seed variances. Pick trial values for the full parameter vector: $\theta = (\omega,\, \alpha_1, \ldots, \alpha_p,\, \beta_1, \ldots, \beta_q)$ subject to $\alpha_i, \beta_j \geq 0$ and $\sum_i \alpha_i + \sum_j \beta_j < 1$.
The recursion needs $\max(p,q)$ lagged values to start. Seed the first $\max(p,q)$ variances and squared returns from sample:

$$\boxed{\sigma_t^2 = u_t^2 = \frac{1}{m_0} \sum_{i=1}^{m_0} u_i^2 \quad \text{for} \quad t = 1, \ldots, \max(p,q)}$$

using first $m_0 \approx 30$–100 returns. Influence decays geometrically.

Step 3: Roll GARCH(p,q) recursion forward through history. For each day $t = \max(p,q) + 1, \ldots, m$, compute conditional variance from p past squared returns and q past conditional variances:

$$\boxed{\sigma_t^2 = \omega + \sum_{i=1}^{p} \alpha_i\, u_{t-i}^2 + \sum_{j=1}^{q} \beta_j\, \sigma_{t-j}^2}$$

Produces the full historical **conditional variance series** $\{\sigma_t^2\}$, generated entirely from the trial parameter vector.

Step 4: Compute conditional log-likelihood. Under conditional normality $u_t \mid \mathcal{F}_{t-1} \sim N(0, \sigma_t^2)$:

$$\boxed{\ell(\theta) = \sum_{t=\max(p,q)+1}^{m} \left[-\ln(\sigma_t^2) - \frac{u_t^2}{\sigma_t^2} \right]}$$

Sum starts at $\max(p,q) + 1$ to skip seeded values that are not model-generated.

Step 5: Optimize parameters. Use a numerical solver to find:

$$\boxed{\hat{\theta} = \arg\max_{\theta}\, \ell(\theta)} \quad \text{subject to} \quad \alpha_i, \beta_j \geq 0,\; \sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{q} \beta_j < 1$$

The solver internally repeats Steps 3–4 for many candidate parameter sets until convergence.

Step 6: Recover model-implied quantities.

- **Long-run daily variance:** $V_L = \dfrac{\hat{\omega}}{1 - \sum_i \hat{\alpha}_i - \sum_j \hat{\beta}_j}$
- **Persistence:** $\sum_i \hat{\alpha}_i + \sum_j \hat{\beta}_j$, closer to 1 means longer memory
- **Mean reversion speed:** $a = \ln\Big(1/(\sum_i \hat{\alpha}_i + \sum_j \hat{\beta}_j)\Big)$
- **Long-run annual volatility:** $\sigma_L^{\text{ann}} = \sqrt{252 \cdot V_L}$

Step 7: Forecast. t-day-ahead variance forecast (mean reversion form, valid for any p, q):

$$\boxed{E[\sigma_{n+t}^2] = V_L + \left(\sum_i \hat{\alpha}_i + \sum_j \hat{\beta}_j \right)^t \left(\hat{\sigma}_n^2 - V_L \right)}$$

For multi-step forecasts beyond $\max(p,q)$ days, use $E[u_{n+k}^2] = \sigma_{n+k}^2$ inside the recursion.

Practical tips and gotchas (general p, q).

- **Lag selection:** use AIC/BIC to compare specifications. GARCH(1,1) almost always wins on parsimony grounds, higher orders rarely improve out-of-sample fit. Box-Jenkins style: start with (1,1), test residuals, increase only if needed.
- **Variance targeting:** fix V_L to historical sample variance, back out $\omega = V_L(1 - \sum \alpha_i - \sum \beta_j)$. Reduces estimation by one parameter regardless of p and q.
- **QMLE in practice:** returns have fat tails, normality is wrong but parameter estimates are still consistent. Correct standard errors via sandwich/HAC formula.
- **Burn-in:** discard first $\sim$ 50–100 fitted variances to remove seed contamination. Length needed grows with $\max(p, q)$.
- **Optimization stability:** more parameters means flatter likelihood and more local maxima. Use multiple starting points, watch for parameters drifting to constraint boundaries (integrated GARCH, $\sum \alpha_i + \sum \beta_j \to 1$).
- **Identification:** GARCH(p,q) with large p, q can have weakly identified parameters. If two specifications give similar likelihood, prefer the smaller one.

⋆ **Variance Targeting: Constrained Estimation.** Instead of estimating ω freely, fix the long-run variance V_L to the **historical sample variance** and back out ω:

$$\boxed{\omega = V_L(1 - \alpha - \beta)} \quad \text{Then estimate only } \alpha \text{ and } \beta \text{ via MLE.}$$

Variance targeting reduces estimation noise by removing one parameter. Long-run vol should not change dramatically from day to day, so anchoring it to a stable historical average produces more stable parameter estimates over time. Standard practice in production GARCH systems.

⋆ **Forecasting Future Volatility**

Rewrite GARCH(1,1) in deviation-from-mean form: $\sigma_n^2 - V_L = \alpha(u_{n-1}^2 - V_L) + \beta(\sigma_{n-1}^2 - V_L)$

Forecast t periods ahead, taking expectations and using: $E[\sigma_{n+t}^2 - V_L] = (\alpha + \beta)\, E[\sigma_{n+t-1}^2 - V_L]$

After repeated substitution, the **t-day forecast** is:

$$\boxed{E[\sigma_{n+t}^2] = V_L + (\alpha + \beta)^t(\sigma_n^2 - V_L)}$$

The deviation from V_L decays geometrically at rate $\alpha + \beta$.

Continuous-time form. Let $V(t) = E[\sigma_{n+t}^2]$ and define $\boxed{a = \ln\Big(1/(\alpha + \beta)\Big)}$:

$$\boxed{V(t) = V_L + e^{-at}\Big[V(0) - V_L\Big]} \quad \text{This is exponential mean reversion toward } V_L \text{ at speed } a.$$

⋆ **Average Variance Over Horizon T (Term Structure Formula).** This is the GARCH(1,1) implied **average daily variance** over the next T days, derived by integrating the continuous-time mean-reversion formula $V(t) = V_L + e^{-at}[V(0) - V_L]$ over $[0, T]$ and dividing by T:

$$\boxed{\sigma_T^2 = \underbrace{\frac{1}{T}\int_0^T V(t)\,dt}_{\substack{\text{average daily variance}\\ \text{over horizon } T}} = \underbrace{V_L}_{\substack{\text{long-run}\\ \text{anchor}}} + \underbrace{\frac{1 - e^{-aT}}{aT}}_{\substack{\text{damping factor}\\ \in(0,1)}} \underbrace{\Big[V(0) - V_L\Big]}_{\substack{\text{current deviation}\\ \text{from long-run}}}}$$

- σ_T^2: **expected average daily variance** over the next T days
- $V(0) = \sigma_n^2$: today's conditional variance (model output)

- $V_L = \omega/(1-\alpha-\beta)$: long-run variance the model reverts to
- $a = \ln\big(1/(\alpha+\beta)\big)$: **mean reversion speed**
- $(1-e^{-aT})/(aT)$: damping factor measuring how much current deviation persists on average over the horizon

The damping factor is the key term: it shrinks the gap between current vol $V(0)$ and long-run vol V_L as horizon grows. Short horizons (T small): factor $\to 1$, $\sigma_T^2 \to V(0)$, today's vol dominates the average. Long horizons (T large): factor $\to 0$, $\sigma_T^2 \to V_L$, long-run vol dominates. This produces the ***volatility term structure****: a curve showing how expected vol differs across maturities.*

Annualized volatility via $\boxed{\sigma_T^{\text{ann}} = \sqrt{252 \cdot \sigma_T^2} = \sqrt{252\left(V_L + \frac{1-e^{-aT}}{aT}\Big[V(0)-V_L\Big]\right)}}$

this is what options pricing and term-structure models actually use. For total expected RV over T days (not average), multiply by T: $E[RV_{0:T}] = T \cdot \sigma_T^2$.

⋆ **Volatility Sensitivity (Term Structure Effect).**

Differentiate σ_T with respect to a shock to instantaneous volatility σ_0: $\boxed{d\sigma_T = \frac{1-e^{-aT}}{aT} \cdot \frac{\sigma_0}{\sigma_T}\, d\sigma_0}$

The damping factor $(1-e^{-aT})/(aT)$ measures how much a short-term vol shock affects volatility at horizon T.

Example: 1% instantaneous vol shock.

Horizon (days)	10	30	50	100	500
Volatility increase (%)	0.90	0.74	0.61	0.41	0.10

A 1% shock to today's vol moves 10-day vol by 0.90% but 500-day vol by only 0.10%, mean reversion damps the shock as horizon lengthens. ***GARCH suggests lower vol shocks for longer maturities****, which is why long-dated implied vols are typically less reactive to short-term events than short-dated implied vols.*

⋆ **GARCH Diffusion Limit: Connection to Stochastic Volatility.**

Start from GARCH(1,1) with $\omega = (1-\alpha-\beta)V_L$: $\sigma_n^2 = (1-\alpha-\beta)V_L + \alpha u_{n-1}^2 + \beta\sigma_{n-1}^2$

Subtract σ_{n-1}^2 from both sides: $\boxed{\sigma_n^2 - \sigma_{n-1}^2 = \underbrace{(1-\alpha-\beta)(V_L - \sigma_{n-1}^2)}_{\text{deterministic mean reversion}} + \underbrace{\alpha(u_{n-1}^2 - \sigma_{n-1}^2)}_{\text{random shock}}}$

Shock term has zero mean and known variance:

Under $u_{n-1} \sim N(0, \sigma_{n-1}^2)$, using $E[u^2] = \sigma^2$ and $E[u^4] = 3\sigma^4$:

$$E[u_{n-1}^2 - \sigma_{n-1}^2] = 0, \qquad \text{Var}(u_{n-1}^2) = E[u^4] - E[u^2]^2 = 3\sigma^4 - \sigma^4 = 2\sigma_{n-1}^4$$

So the shock $(u_{n-1}^2 - \sigma_{n-1}^2)$ has standard deviation $\sqrt{2}\,\sigma_{n-1}^2$.

Continuous-time diffusion limit. **Let $V = \sigma_{n-1}^2$, $\Delta V = \sigma_n^2 - \sigma_{n-1}^2$, $\kappa = 1-\alpha-\beta$. Writing the discrete shock as $Z \sim N(0, \sqrt{2}\,\sigma_{n-1}^2)$ and taking the continuous-time limit:**

$$\boxed{dV = \kappa(V_L - V)\,dt + \alpha\sqrt{2V}\,dz}$$

This is a **CIR/Heston-type stochastic variance process** with:

- $\kappa = 1-\alpha-\beta$: **mean reversion speed** toward V_L

- V_L: **long-run variance** (same as GARCH long-run level)
- $\alpha\sqrt{2V}$: **vol-of-vol** scales with $\sqrt{V}$, variance becomes more volatile when already elevated
- Square-root diffusion keeps $V \geq 0$ always (Feller condition)

GARCH(1,1) is the discrete-time approximation of a continuous-time stochastic volatility model. Same dynamics, just sampled differently. This is why GARCH and Heston produce similar vol forecasts empirically, and why both fit equity vol data well, they are mathematically the same mean-reverting square-root process.

Key distinction: in GARCH the innovation $u_{n-1}^2 - \sigma_{n-1}^2$ is **deterministic given past returns**, once you know u_{n-1} the variance is fixed. In Heston the innovation dz is a **genuinely independent Brownian shock**. So GARCH has **filtered** variance (no latent randomness), while Heston has **truly stochastic** variance with its own random driver.

$\star$ **EGARCH: Exponential GARCH (Asymmetric Volatility).** Standard GARCH treats positive and negative returns symmetrically because u_{n-1}^2 kills the sign, a $+2\%$ and a -2% return contribute identically to next day's variance. Empirically wrong: **down moves spike vol more than equivalent up moves** (the **leverage effect**). **EGARCH** (Nelson 1991) fixes this by adding an explicit sign term and modeling $\ln(\sigma_t^2)$ instead of σ_t^2.

Return and innovation structure:

$$y_t = \mu + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t, \qquad z_t \sim N(0,1)$$

where ε_t is the raw shock and $z_t = \varepsilon_t/\sigma_t$ is the **standardized shock**, unit variance. EGARCH dynamics are written in terms of z_t, not ε_t.

EGARCH(1,1) variance equation: $\boxed{\ln(\sigma_t^2) = \kappa + \beta \ln(\sigma_{t-1}^2) + \alpha \underbrace{\left[\frac{|\varepsilon_{t-1}|}{\sigma_{t-1}} - \sqrt{\frac{2}{\pi}}\right]}_{\text{magnitude effect}} + \xi \underbrace{\frac{\varepsilon_{t-1}}{\sigma_{t-1}}}_{\text{sign effect (leverage)}}}$

- κ: intercept in log-variance space
- $\beta \ln(\sigma_{t-1}^2)$: **persistence**, today's log-variance inherits from yesterday's. $|\beta| < 1$ required for stationarity.
- $\alpha\left[|z_{t-1}| - \sqrt{2/\pi}\right]$: **magnitude effect**. $|z_{t-1}|$ is the absolute standardized shock; $\sqrt{2/\pi} = E[|z|]$ for standard normal, so the bracket has **mean zero**. Large shocks (in either direction) push this term positive, small shocks push it negative. Captures volatility clustering regardless of sign.
- $\xi\, z_{t-1} = \xi\, \varepsilon_{t-1}/\sigma_{t-1}$: **sign effect / leverage**. This is the whole point of EGARCH. If $\xi < 0$, a negative shock adds $\xi z < 0 \cdot (-) > 0$ to $\ln \sigma_t^2$ (raising vol), while a positive shock of equal size **reduces** it. Asymmetric response recovered.

Why $\ln(\sigma_t^2)$ and not σ_t^2? Three reasons: (1) ***automatic positivity****: $\sigma_t^2 = \exp(\ln \sigma_t^2) > 0$ always, no parameter constraints needed; (2)* ***no sign restrictions on α, β****: standard GARCH needs $\alpha, \beta \geq 0$ to keep $\sigma_t^2 > 0$, EGARCH frees this because exp handles positivity; (3)* ***richer dynamics****: log allows multiplicative rather than additive effects, which matches how vol actually behaves in data.*

The sign effect in detail. The leverage term ξz_{t-1} produces asymmetric vol responses:

$$\text{Impact on } \ln \sigma_t^2 \;=\; \alpha\left[|z_{t-1}| - \sqrt{2/\pi}\right] + \xi z_{t-1}$$

Breaking out by sign of z_{t-1}:

- **Positive shock** ($z_{t-1} > 0$): contributes $(\alpha + \xi) z_{t-1}$ (magnitude of shock on vol)
- **Negative shock** ($z_{t-1} < 0$): contributes $(\alpha - \xi)|z_{t-1}|$

With $\xi < 0$: negative shocks get weight $\alpha - \xi > \alpha + \xi$ of positive shocks, so **down moves amplify vol more than up moves**. This directly matches equity data: SPX drops cause bigger VIX spikes than SPX rallies of equal magnitude.

EGARCH is the go-to model when asymmetric vol response matters, equity indices, single stocks, anything with a visible leverage effect. If the asymmetry is weak (commodities, some FX pairs), standard GARCH usually suffices. Test ξ for significance: $\hat{\xi}$ statistically < 0 confirms leverage.

⋆ **GARCH-M: GARCH-in-Mean.** Adds **volatility as a regressor in the mean equation**, embedding the **risk-return tradeoff** directly into the model: investors demand higher expected return when vol is higher. Consistent with asset pricing theory (CAPM, ICAPM, Merton's ICAPM).

GARCH-M model:

$$\boxed{y_t = \mu + \lambda\sigma_t + \varepsilon_t, \qquad \varepsilon_t = \sigma_t z_t} \qquad \sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta\sigma_{t-1}^2$$

- μ: baseline expected return (unconditional mean)
- $\lambda\sigma_t$: **risk premium**, extra expected return proportional to today's vol
- $\varepsilon_t = \sigma_t z_t$: random shock scaled by current vol
- σ_t^2: standard GARCH(1,1) recursion

Interpretation of λ (the price of risk). λ measures **extra return per unit of volatility**. In asset pricing theory:

- **CAPM / ICAPM prediction:** $\lambda > 0$, risk-averse investors demand compensation for vol
- **Typical equity data:** $\hat{\lambda}$ small and often statistically insignificant. The weak evidence is called the "risk-return puzzle".
- **Alternative variants:** $\lambda\sigma_t^2$ (variance in mean) or $\lambda \ln \sigma_t^2$ (log-variance in mean), theory doesn't pin down which scaling is correct

GARCH-M couples the mean and variance equations: when vol rises, expected return rises. Useful when modeling assets where returns and vol comove, risk assets around events, emerging markets, credit-sensitive securities. Estimation is harder because σ_t appears in both equations, MLE must be done jointly. In practice $\hat{\lambda}$ is small and imprecise for broad equity indices, so GARCH-M is more common in academic work than production.

7.3 Value at Risk (VaR)and Expected Shortfall (ES)

⋆ **Coherent Risk Measure.** A **coherent risk measure** is one that behaves the way risk *should* behave under four common-sense rules: bigger losses give higher risk, diversification never hurts, scaling a position scales the risk, and adding cash reduces risk one-for-one. **ES satisfies all four; VaR fails subadditivity**, which is the main theoretical reason regulators have shifted toward ES.

A **risk measure** $\rho(X)$ assigns a number to a loss random variable X. A **coherent risk measure** satisfies four properties (Artzner et al. 1999):

- **Monotonicity:** worse outcomes imply higher risk. **If** $X \leq Y$ **then** $\rho(X) \leq \rho(Y)$**.**
- **Subadditivity:** diversification cannot increase risk. $\rho(X+Y) \leq \rho(X) + \rho(Y)$
- **Positive Homogeneity:** scaling a position scales the risk proportionally. **For all** $\lambda \geq 0$**,** $\rho(\lambda X) = \lambda \rho(X)$**.**
- **Translation Invariance:** adding cash reduces risk one-for-one. **For any deterministic cash** c**,** $\rho(X+c) = \rho(X) - c$**.**

These four properties are what a "well-behaved" risk measure should satisfy. Subadditivity is the most important in practice: it guarantees combining two positions never makes the total risk larger than the sum of individual risks. ***ES is coherent. VaR is not*** *(it can violate subadditivity in fat-tailed cases), which is the main theoretical reason regulators have shifted toward ES.*

⋆ **Three Risk Definitions:**

- **Value at Risk (VaR):** loss level not exceeded with probability α.
  Answers: ***"How much can I lose at a certain probability?"***
- **Expected Shortfall (ES):** expected loss given that the loss exceeds VaR.
  Answers: ***"If things go bad, how bad on average?"*** Also called **Conditional Tail Expectation (CTE)**. Captures tail severity, not just frequency.
- **Economic Capital (EC):** capital needed to remain solvent after an economic shock at confidence α.
  Answers: ***"How much capital do I need to survive a crisis?"*** Internal alternative to regulatory capital based on full economic view.

VaR tells you the threshold of bad outcomes; ES tells you how bad it gets beyond that threshold. VaR ignores everything in the tail past its quantile, ES averages over it. For fat-tailed distributions ES is much more informative than VaR.

⋆ **VaR and ES Formulas (Normal Loss Distribution).** Assume losses $\sim N(\mu, \sigma^2)$ over horizon T days. Then:

$$\boxed{\mathrm{VaR}_\alpha = \mu T + \sigma\sqrt{T}\,\Phi^{-1}(\alpha)} \qquad \boxed{\mathrm{ES}_\alpha = \mu T + \sigma\sqrt{T}\,\frac{\phi(\Phi^{-1}(\alpha))}{1-\alpha}}$$

- α: confidence level (e.g. 0.99)
- μ: mean loss, typically assumed zero over short horizon (1 day)
- σ: loss standard deviation per day
- $\sqrt{T}$: time scaling under i.i.d. assumption
- $\Phi^{-1}(\alpha)$: standard normal quantile, e.g. $\Phi^{-1}(0.99) = 2.326$
- $\phi(\cdot)$: standard normal pdf

VaR: take the expected loss over horizon T, plus the loss volatility scaled to T, multiplied by the standard normal quantile at confidence α, this is the loss level that gets exceeded only $(1-\alpha)$ of the time.

ES: same structure, but the quantile is replaced by a tail multiplier $\phi(\Phi^{-1}(\alpha))/(1-\alpha)$, which is the conditional mean of a standard normal given it exceeds the quantile, this gives the average loss **inside** the tail, not just the threshold.

ES is always larger than VaR at the same α because it averages over the tail beyond VaR. The ratio $\phi(\Phi^{-1}(\alpha))/(1-\alpha)$ is the "tail multiplier": at $\alpha = 0.99$ it equals 2.665, vs $\Phi^{-1}(0.99) = 2.326$, so ES is about 14% larger than VaR for normal losses.

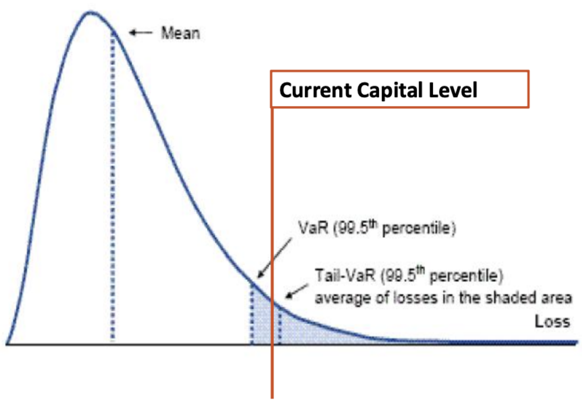

*The plot shows the **loss distribution**: x-axis = loss size, y-axis = probability density. **VaR** is the loss level at the 99.5th percentile (only 0.5% of outcomes are worse), **Tail-VaR / ES** is the average of all losses in the shaded tail beyond VaR. The **Current Capital Level** is the actual cash buffer the firm holds, if it sits to the right of VaR (as drawn) the firm is adequately capitalized to survive a 99.5% adverse event, if it sits to the left the firm is undercapitalized and would fail under that stress.*

Worked Example: Single-Asset VaR. Position \$10MM, daily vol 2% (32% annualized), daily volatility exposure \$10MM $\times$ 2% = \$200,000, confidence 99% ($\Phi^{-1}(0.99) = 2.326$).

$$\text{Daily VaR}_{99\%} = \$200{,}000 \times \sqrt{1} \times 2.326 = \$465{,}200; \quad \text{10-day VaR}_{99\%} = \$200{,}000 \times \sqrt{10} \times 2.326 = \$1{,}470{,}000$$

Company XYZ. Position \$5MM, daily vol 1%, exposure \$50,000.

$$\text{Daily VaR}_{99\%} = \$116{,}300, \qquad \text{10-day VaR}_{99\%} = \$367{,}800$$

*The structure is always the same: **VaR = (position × vol) × $\sqrt{T} \times \Phi^{-1}(\alpha)$.** Time scales as $\sqrt{T}$, confidence enters through the normal quantile, position and vol enter linearly.*

⋆ **Portfolio VaR with Correlation.** For a two-asset portfolio with bivariate normal returns and correlation ρ:

$$\boxed{\sigma_{\text{ABC+XYZ}} = \sqrt{\sigma_{\text{ABC}}^2 + \sigma_{\text{XYZ}}^2 + 2\rho\,\sigma_{\text{ABC}}\sigma_{\text{XYZ}}}}$$

Example numbers. $\sigma_{\text{ABC}} = \$200{,}000$, $\sigma_{\text{XYZ}} = \$50{,}000$, $\rho = 0.3$:

$$\sigma_{\text{ABC+XYZ}} = \sqrt{200{,}000^2 + 50{,}000^2 + 2(0.3)(200{,}000)(50{,}000)} = \$220{,}200$$

10-day Portfolio VaR at 99%: $\$220{,}200 \times \sqrt{10} \times 2.326 = \$1{,}620{,}100$.

This is the Markowitz portfolio variance formula applied to dollar volatilities. Correlation matters: $\rho = 1$ means risks add ($\sigma_p = \sigma_1 + \sigma_2$), $\rho = 0$ means quadratic addition ($\sigma_p = \sqrt{\sigma_1^2 + \sigma_2^2}$), $\rho < 0$ means partial offset.

⋆ **Diversification Benefit. Definition:** sum of marginal (standalone) VaRs minus joint portfolio VaR:

$$\boxed{\text{Diversification Benefit} = \underbrace{\sum_i \text{VaR}_i}_{\text{undiversified EC}} - \underbrace{\text{VaR}_{\text{portfolio}}}_{\text{diversified EC}}}$$

If you held each asset separately, you would need to reserve standalone VaR for each ($\sum_i$ VaR$_i$); but held together, losses don't all hit on the same day, so the joint portfolio VaR is smaller, the gap (sum minus joint) is the capital you "save" from imperfect correlation.

Numerical example. Marginals: \$1,471,300+\$367,800 = \$1,839,100. Joint portfolio VaR: \$1,620,100. Diversification benefit: \$1,839,100 − \$1,620,100 = \$219,000 (about 12% of undiversified).

⋆ **VaR Calculation Approaches.** Four standard methodologies in practice:

1. **Historical Simulation:** apply observed historical market changes directly to current portfolio. Non-parametric, no distributional assumptions, captures real fat tails. Standard for trading book VaR.
2. **Scenario Approach:** apply specific stress paths (Dodd-Frank/CCAR bank stress test, designer scenarios). Forward-looking, regulator-driven.
3. **Approximations:** linear (delta) and quadratic (delta-gamma) approximations for portfolio P&L given small market moves. Fast, but breaks down for nonlinear products and large moves.
4. **Model Building (Full Valuation):** model the underlying economics, re-evaluate every position at each Monte Carlo scenario. Most accurate but computationally expensive.

⋆ **1. Historical Simulation: Steps.**

1. Build historical database of all market variables and their daily changes
2. Apply each historical daily change to current values to create a set of outcome scenarios (decide between absolute and percentage changes per asset class)
3. Re-evaluate the current portfolio under each simulated outcome
4. Rank the resulting losses, compute VaR at α and ES as average of tail

The key insight: historical simulation ***reuses past market moves as proxies for tomorrow's possible moves****. You don't simulate from a model, you replay history. Captures non-normality and correlation structure for free since both are baked into real data, but assumes the past distribution is representative of the future.*

Generate scenarios . For each historical day $i = 1, 2, \ldots, m$, simulate the next-day value of every market variable v by applying that day's observed percentage change to today's level:

$$\boxed{v_{n+1}^{(i)} = v_n \times \frac{v_i}{v_{i-1}}}$$

- v_n: today's (current) value of the variable
- v_i/v_{i-1}: ratio of the variable on historical day i to day $i-1$, i.e. that day's gross return
- $v_{n+1}^{(i)}$: simulated value of the variable on day $n+1$ under scenario i

This is the heart of ***historical simulation****: rather than assuming a distribution, you treat each past day's percentage move as one possible realization of tomorrow's move. With m days of history you generate m alternative scenarios for tomorrow, revalue the full portfolio under each, sort the resulting P&Ls, and read off VaR (the α-quantile loss) and ES (average of losses beyond VaR). For rates and similar variables, absolute changes are sometimes used instead of percentage changes:* $v_{n+1}^{(i)} = v_n + (v_i - v_{i-1})$.

$$11{,}022 \times \frac{11{,}173}{11{,}219} = 10{,}977$$

Scenario	DJIA	FTSE 100	CAC 40	Nikkei 225	Portfolio Value ($000s)	Loss ($000s)
1	10,977	9,569	6,205	115	10,021	-21
2	10,926	9,676	6,293	113	10,019	-19
3	11,070	9,454	6,088	112	9,954	46
…	…….	…….	…….	…….	…….	……..
499	10,831	9,383	6,051	113	9,857	143
500	11,223	9,763	6,371	111	10,134	-134

$$\frac{10{,}977}{11{,}223} * 4{,}000 + \frac{9{,}569}{9{,}763} * 3000 + \frac{6{,}205}{6.371} * 1{,}000 + \frac{115}{111} * 2{,}000 = 10{,}021$$

Scenario Number	Loss ($000s)
494	478
339	345
349	282
329	277
487	253
227	218
131	205

99% one-day VaR

Reading off risk measures from ranked losses. After computing portfolio value under all 500 historical scenarios and sorting the resulting losses from worst to best, the risk measures are read directly off the ranked tail:

- **99% one-day VaR** = 5th worst loss out of 500, since $5/500 = 1\%$ of scenarios are worse. Result: **\$253,000**.
- **99% one-day ES** = average of the 5 worst losses $(478 + 345 + 282 + 277 + 253)/5$. Result: **\$327,000**.

*VaR is the **threshold**: the loss level that gets exceeded only 1% of the time. ES is the **average severity past that threshold**, the expected loss conditional on being in the worst 1% of outcomes. ES is always larger than VaR because it integrates over the tail rather than reading off a single quantile, which is why regulators have shifted toward ES as a more informative tail-risk measure.*

⋆ 2. Scenario Approach

Apply specific historical or hypothetical stress patterns to current market variables. Common scenarios: COVID (2020), 2008-2009 Financial Crisis, 2001-2002 dot-com bust, 1987 crash, 1970s stagflation, Great Depression

Adapt to current conditions: interest rates may be much lower or higher than during historical stresses; equity markets may be at elevated or depressed levels relative to the historical baseline; modify scenarios to reflect company-specific risks or management views.

*Scenario analysis is more flexible than historical simulation because you can **design** the shock pattern instead of replaying observed history. This matters when current conditions are unlike anything in the historical database (negative rates, post-COVID recovery), or when management wants to stress-test a specific risk thesis.*

⋆ CCAR Comprehensive Capital Analysis and Review.

Step 1: Fed publishes scenario (input, prescribed). The Fed releases a **synthetic** 9-quarter macro trajectory each year: **not historical replay**, custom-designed to be more severe than any single past quarter. Three scenarios provided: Baseline, Adverse, Severely Adverse. Each specifies ~ 28 variables per quarter:

- **Macro:** GDP growth, unemployment, CPI
- **Rates:** 3M UST, 10Y UST, BBB corp yield, mortgage rate
- **Markets:** DJTR, VIX, HPI, CRE

The bank receives this trajectory as fixed input, no choice in the path.

Step 2: Bank uses internal models calibrated on real historical data. Banks do not replay history, they use **regression-style models** trained on decades of past data that map macro variables to losses. The historical data lives in the model coefficients, not in the scenario itself. Generic form:

$$\text{Loss}_t = f(\text{UR}_t,\ \text{GDP}_t,\ \Delta\text{HPI}_t,\ \ldots)$$

Example loss models by line of business:

- **Mortgage charge-offs:** $\text{CO}_t = \beta_0 + \beta_1 \text{UR}_t + \beta_2 \Delta\text{HPI}_t$
- **Credit card losses:** $\text{Loss}_t = f(\text{UR}_t, \text{GDP}_t, \text{personal income}_t)$
- **C&I loan losses:** $\text{Loss}_t = f(\text{GDP}_t, \text{BBB spread}_t, \text{unemployment}_t)$
- **Trading book:** full revaluation under Fed equity/rate/FX shocks (not statistical model)
- **Pre-Provision Net Revenue (PPNR):** $\text{PPNR}_t = \text{NII}_t + \text{Fees}_t - \text{OpEx}_t$, each component a function of rates and GDP

Step 3: Quarter-by-quarter capital projection (output). For each quarter t, plug Fed macro values into models, get projected losses and revenue, update capital:

$$\text{Capital}_t = \text{Capital}_{t-1} + \text{PPNR}_t - \text{Losses}_t - \text{Taxes}_t - \text{Dividends}_t$$

Then capital ratio: $$\text{CET1 Ratio}_t = \frac{\text{CET1}_t}{\text{Risk-Weighted Assets}_t}$$

Must stay above 4.5% + Stress Capital Buffer every quarter, not just at the end.

7.3.1 Model Building Approach (VaR)

⋆ **Model Building Approach.** The most general VaR methodology: instead of replaying history (Historical Simulation) or applying a single stress (Scenario Approach), specify a joint statistical model for all market variables and simulate scenarios to build a full loss distribution.

- Specify **future joint distribution** of all market variables
- Generate large number of economic scenarios
- Calculate value of all assets/liabilities in each scenario
- Aggregate portfolio value under every scenario
- Calculate VaR or ES from the resulting loss distribution

Alternative to historical simulation:

- Introduces **subjectivity** of defining economic scenarios (every assumption is a modeling choice)
- Must specify **marginal distribution** of each factor along with the **joint distribution**
- **Dependency modeling (copulas)** important, especially for tail risk

Compared to historical simulation, the model-building approach is more flexible (can capture stress regimes not in history) but introduces subjectivity. Joint distribution matters more than marginals for tail risk: two assets with the same marginal distributions but different copulas can give very different VaR.

⋆ **Specifications:**

Economic factors to specify: all factors affecting asset/liability valuations: **Equity returns, Interest rates, Credit spreads, Implied volatility**

Three sub-approaches by valuation complexity:(Modeling Complexity of Positions

- **Linear Approximation:** delta only
- **Quadratic Approximation:** delta + gamma
- **Full Valuation:** re-price every position from scratch in each scenario

Linear and quadratic are cheap but lossy approximations of the true valuation. Full valuation is exact but computationally expensive, especially for large derivative books. Choice depends on portfolio composition and computational budget.

⋆ **Linear Model.** Assume portfolio value change is **linearly related** to economic factor shocks:

$$\boxed{\Delta P = \sum_{i=1}^{N} \Delta V_i = \sum_{i=1}^{N} w_i \sum_{j=1}^{K} \delta_{i,j} \Delta f_j}$$

- V_i: value of asset position i
- w_i: units invested in asset i
- $\delta_{i,j} = \partial V_i / \partial f_j$: **first-order sensitivity (delta)** of asset i to factor j
- Δf_j: shock to economic factor j
- K: number of factors, N: number of assets

*Linear model works for **linear contracts**: stocks, forwards, futures, swaps. **Fails for nonlinear contracts** like options because it ignores curvature, a 1% stock move and a −1% stock move produce equal but opposite P&L under the linear model, but for an option they don't.*

Why Linear Fails: Impact of Option Gamma. Options have nonlinear payoffs, so equal-magnitude up and down moves do not produce symmetric P&L:

- **Positive gamma** (long options): payoff convex, gains larger than losses of equal magnitude ⇒ **right-skewed P&L distribution**
- **Negative gamma** (short options): payoff concave, losses larger than gains ⇒ **left-skewed P&L distribution**

The linear model forces a symmetric normal distribution onto an inherently skewed payoff, missing exactly the tail risk that matters most for VaR.

⋆ **Asset Price vs Option Price Mapping.** Even when the underlying **asset price** is normally distributed, the **option price** is not, because the asset-to-option payoff curve is nonlinear (kinked at the strike for a call). A symmetric normal distribution on the asset price gets mapped through the curved payoff function into a **skewed, non-normal distribution** on the option value.

This is why VaR for option books cannot be computed with linear factor sensitivities alone: the mapping itself transforms the distribution shape. You need at least quadratic (gamma) terms or full revaluation to capture the actual P&L distribution.

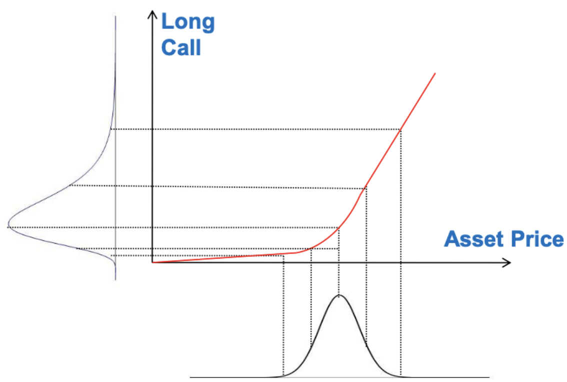

The bottom bell shows a **normal distribution of asset prices**, symmetric and thin-tailed. The red curve is the **long call payoff**: flat below strike (worthless), then bending upward and becoming linear deep ITM. The dotted lines map each asset price through the payoff to its corresponding option value, equal-spaced asset prices map to **unequal-spaced** option values because the curve bends. The result on the left is the **option value distribution**: heavily skewed with a spike near zero (calls expire worthless when asset stays below strike) and a long right tail (rare large gains when asset rallies). **Symmetric input becomes skewed output** purely because the payoff function is nonlinear, this is exactly why linear VaR models fail for option books and why quadratic or full valuation is needed.

⋆ **Quadratic Model (Delta-Gamma).** Add second-order Taylor terms to capture curvature:

$$\boxed{\Delta P = \sum_{i=1}^{N} w_i \sum_{j=1}^{K} \delta_{i,j} \Delta f_j + \frac{1}{2} \sum_{i=1}^{N} w_i \sum_{j=1}^{K} \sum_{k=1}^{K} \gamma_{i,j,k}\, \rho_{ijk}\, \Delta f_j \Delta f_k}$$

- $\delta_{i,j} = \partial V_i / \partial f_j$: first derivative (delta)
- $\gamma_{i,j,k} = \partial^2 V_i / (\partial f_j \partial f_k)$: **second derivative (gamma / cross-gamma)**
- ρ_{ijk}: correlation between factor shocks j and k

Quadratic captures the ***curvature*** *of nonlinear payoffs, like a delta/gamma hedge in options or a duration/convexity approximation in bonds. Much better than linear for option books, still fast since you only compute Greeks once and reuse them per scenario, but breaks down for large moves where third-order and higher terms matter.*

⋆ **Instrument Approximations: Required Greeks by Asset Class.** Different instruments have different sensitivity exposures, the table shows which Greeks are needed per (instrument, factor) combination:

Factor		**Stock**	**Risky Bond**	**Risk-Free Bond**	**Put Option**
2*Equity	Delta	X	0	0	X
	Gamma	0	0	0	X
2*Interest Rate	Delta	0	X	X	X
	Gamma	0	X	X	X
2*Credit Spread	Delta	0	X	0	0
	Gamma	0	X	0	0
2*Implied Vol	Delta	0	0	0	X
	Gamma	0	0	0	X

Stocks need only equity delta (linear). Bonds need rate delta and gamma (plus credit for risky bonds). Put options have nonzero exposure across nearly every factor including implied vol (vega), this is why options are the hardest instruments to approximate.

⋆ **Full (Exact) vs Approximate Models.** The **exact** portfolio change is the difference in valuations across all factors:

$$\boxed{\Delta V = V(S_t, r_t, CS_t, IV_t) - V(S_0, r_0, CS_0, IV_0)} \quad \text{Exact}$$

Taylor-expanding to second order produces the linear and quadratic approximations:

$$\Delta V \approx \underbrace{\frac{\partial V}{\partial t}dt}_{\substack{\theta \\ \approx 0 \text{ daily}}} + \underbrace{\frac{\partial V}{\partial S}dS + \frac{\partial V}{\partial r}dr + \frac{\partial V}{\partial CS}dCS + \frac{\partial V}{\partial IV}dIV}_{\text{linear terms (deltas)}}$$

$$+ \underbrace{\frac{1}{2}\frac{\partial^2 V}{\partial S^2}(dS)^2 + \frac{1}{2}\frac{\partial^2 V}{\partial r^2}(dr)^2 + \frac{1}{2}\frac{\partial^2 V}{\partial CS^2}(dCS)^2 + \frac{1}{2}\frac{\partial^2 V}{\partial IV^2}(dIV)^2}_{\text{quadratic terms (gammas)}}$$

$$+ \textit{ cross-product terms} + \textit{higher-order terms (only the exact valuation includes all of them).}$$

*The Greeks (δ, γ) are derivatives evaluated **once** at the current state and reused across scenarios. Full valuation **re-prices each position from scratch** in each scenario, capturing all higher-order effects exactly but at much higher computational cost. The theta term $\partial V/\partial t \cdot dt$ is approximately zero for daily horizons.*

⋆ **Monte Carlo Approach: Workflow.**

- Calculate initial portfolio value (Economic Value at $t = 0$)
- Generate set of economic variables at horizon T (1-day, 1-year, ...)
- Value the portfolio using the simulated economic variables
- Repeat N times to build distribution of portfolio loss ΔP
- Calculate risk metrics: VaR (quantile), ES (tail average)
- Scale to desired horizon

⋆ **Economic Variable Projections (Multi-Factor Diffusion).** This is the **Step 2 of Monte Carlo**: how to actually generate one scenario for all four economic factors simultaneously, with realistic correlation structure between them. Each factor follows a discretized Brownian motion (drift + scaled random shock), and all shocks are drawn from a **joint multivariate normal** so they move together the way markets actually do.

Joint diffusion equation:

$$\boxed{\begin{bmatrix} S_t \\ r_t \\ CS_t \\ IV_t \end{bmatrix} = \underbrace{\begin{bmatrix} S_0 \\ r_0 \\ CS_0 \\ IV_0 \end{bmatrix}}_{\text{current values}} + \underbrace{\begin{bmatrix} \mu_S \\ \mu_r \\ \mu_{CS} \\ \mu_{IV} \end{bmatrix} dt}_{\text{drift over } dt} + \underbrace{\begin{bmatrix} \sigma_S Z_{S,t} \\ \sigma_r Z_{r,t} \\ \sigma_{CS} Z_{CS,t} \\ \sigma_{IV} Z_{IV,t} \end{bmatrix}}_{\text{correlated random shocks}} \sqrt{dt}}$$

The shocks $Z = (Z_S, Z_r, Z_{CS}, Z_{IV})'$ are drawn from a **multivariate normal** $Z \sim N(0, \Sigma)$ where Σ encodes the correlations between factor shocks:

$$\Sigma = \begin{bmatrix} 1 & * & * & * \\ \rho_{S,r} & 1 & * & * \\ \rho_{S,CS} & \rho_{r,CS} & 1 & * \\ \rho_{S,IV} & \rho_{r,IV} & \rho_{CS,IV} & 1 \end{bmatrix} \quad \text{e.g.} \quad \begin{bmatrix} 1 & * & * & * \\ 0.4 & 1 & * & * \\ 0.6 & -0.4 & 1 & * \\ 0.8 & -0.1 & 0.7 & 1 \end{bmatrix}$$

Typical daily vols: $\sigma_S = 0.17$, $\sigma_r = 0.01$, $\sigma_{CS} = 0.01$, $\sigma_{IV} = 0.03$.

Where it fits in the workflow:

- One application of this equation generates **one scenario** of factor values at horizon T
- Repeat $N = 10{,}000$ times to get a full set of joint scenarios
- For each scenario, plug (S_t, r_t, CS_t, IV_t) into the pricing model (or linear/quadratic approximation) to compute portfolio value
- Sort the resulting losses, read off VaR and ES

The whole point of the matrix form is to model factors ***jointly****, not independently. Equity vol moves up when stocks fall ($\rho_{S,IV} = 0.8$ here, but note this is correlation of shocks, hence positive between equity down move and IV up move depends on sign convention). Credit spreads widen when stocks fall ($\rho_{S,CS} = 0.6$), capturing the risk-on/risk-off pattern. Without correlation the model would assume all four factors move independently, which dramatically underestimates joint tail risk in a real crisis.*

In practice, generating correlated draws is done via ***Cholesky decomposition****: factor $\Sigma = LL'$, draw independent $Z' \sim N(0, I)$, then correlated draws are $Z = LZ'$. For daily horizons drift μ is set to zero since it is dwarfed by $\sigma\sqrt{dt}$, longer horizons must include drift.*

Drift μ can be set to zero for daily steps since it is dwarfed by $\sigma\sqrt{dt}$. Longer horizons need to account for drift. To generate correlated normal draws in practice: take Cholesky decomposition $\Sigma = LL'$, generate independent $Z' \sim N(0, I)$, then correlated draws are $Z = LZ'$.

⋆ **Modeling Approaches Comparison.**

- **Model Building Approach:** requires assumptions for market variables. **Joint distribution is key for VaR**. Enhance with non-normal distributions and copulas to aggregate risks accurately. Use linear/quadratic approximations when full valuation is computationally infeasible.
- **Historical Simulation:** lets historical data determine the distributions. Easy to backtest: how often was the actual 1-day loss greater than the 99%/1-day VaR?
- **Stress Testing (CCAR):** directly determines the economic environment via prescribed scenarios. Defined over a given horizon: instantaneous, 1-year, multi-year.

In practice banks combine all three: historical simulation for daily trading book VaR, Monte Carlo full valuation for complex derivatives, scenario analysis and CCAR for regulatory stress testing. No single approach is universally best, each has different strengths, costs, and validation properties.

7.4 Credit VaR

⋆ **Credit VaR.** Extends VaR to losses driven by **credit defaults** rather than market price moves. The challenge: defaults are rare, low-probability binary events with strong tail dependence, so naive VaR methods don't work well. **Complications:**

- **Tail estimation:** large numbers of scenarios needed to estimate tails accurately because defaults are rare events
- **Default correlation:** hard to estimate, defaults are rare and clustering is hard to observe directly

What credit risk covers (and what it doesn't):

- **Permanent loss** due to default within the horizon, the main thing
- **Future losses beyond horizon** sometimes evaluated, induced by VaR-stressed default rates
- **Market value changes from spread movements** are covered under **market risk**, not credit risk, even though they look similar (these can subside post-event, default losses cannot)

Estimate T-year Credit VaR: The standard analytic formula for portfolio credit losses is the **Vasicek ASRF model** (Asymptotic Single Risk Factor), which gives the loss level not exceeded with probability X over a T-year horizon.

The key conceptual difference from market VaR: market VaR asks "how bad can prices move?"; Credit VaR asks "how many defaults could happen jointly, and what would the total loss be?". The math is fundamentally different because default is binary (default or not) while price moves are continuous.

⋆ **Gaussian Copula Model: Asset Value Decomposition.** The standard model for default correlation. Each firm i's asset value is decomposed into a **systemic factor** F and an **idiosyncratic factor** ε_i:

$$\boxed{A_i = \sqrt{\rho}\,F + \sqrt{1-\rho}\,\varepsilon_i}$$

where F and ε_i are independent standard normals, $F \sim N(0,1)$, $\varepsilon_i \sim N(0,1)$.

- F: **common (systemic) factor**, hits all firms simultaneously (e.g. the macro/credit cycle)
- ε_i: **idiosyncratic** firm-specific shock, uncorrelated across firms
- ρ: copula correlation, controls how much weight goes on the common vs idiosyncratic component

Pairwise correlation between firms i and j.

$$\rho[A_i, A_j] = E\Big[(\sqrt{\rho}F + \sqrt{1-\rho}\varepsilon_i)(\sqrt{\rho}F + \sqrt{1-\rho}\varepsilon_j)\Big] = E[\rho F^2] + E[\sqrt{\rho(1-\rho)}F(\varepsilon_i+\varepsilon_j)] + E[(1-\rho)\varepsilon_i\varepsilon_j]$$

Since $E[F^2] = 1$, $E[F\varepsilon] = 0$, and $E[\varepsilon_i\varepsilon_j] = 0$ for $i \neq j$: $\boxed{\rho[A_i, A_j] = \rho}$

This is the elegant trick of the Gaussian copula: a single correlation parameter ρ pins down the dependence between every pair of firms. High $\rho \Rightarrow$ firms move together, default clusters in bad systemic states. Low $\rho \Rightarrow$ firms default independently. Empirically $\rho \approx 0.10$ to 0.20 for retail, higher for corporates.

⋆ **Gaussian Copula: Default Probability Conditional on Systemic Factor.** A firm **defaults** when its asset value falls below a threshold L. Given the realized systemic factor F, the conditional probability of default is:

$$PD = \Pr(A_i < L \mid F) = \Pr\Big(\sqrt{\rho}F + \sqrt{1-\rho}\varepsilon_i < L\Big)$$

Solving for ε_i:

$$= \Pr\left(\varepsilon_i < \frac{L - \sqrt{\rho}F}{\sqrt{1-\rho}}\right) = \Phi\left(\frac{L - \sqrt{\rho}F}{\sqrt{1-\rho}}\right)$$

Setting the threshold. The default threshold L is chosen so that the **unconditional probability of default** equals the firm's PD. Since $A_i \sim N(0,1)$ unconditionally:

$$\boxed{L = \Phi^{-1}(PD)}$$

*This is the inverse map: if a firm has 1% unconditional default probability, the threshold is $\Phi^{-1}(0.01) = -2.326$. Asset values below this trigger default. The conditional formula then says: given a bad systemic state ($F < 0$), the conditional default probability rises above PD; given a good state ($F > 0$), it falls below PD. This captures **default clustering**: many firms default together when F is bad.*

⋆ **Vasicek Portfolio Loss Formula (ASRF).** Apply the Gaussian copula to a **large portfolio** of loans, all with the same $PD = Q(T)$, same pairwise correlation ρ, and same loss given default. By the law of large numbers (idiosyncratic shocks ε_i average out), the portfolio loss rate is fully determined by the common factor F.

Loss formula at confidence $X\%$: we are $X\%$ certain that default rate is less than:

$$\boxed{V(X,T) = L_{\text{exposure}} \cdot \Phi\left(\frac{\Phi^{-1}[Q(T)] + \sqrt{\rho}\,\Phi^{-1}(X)}{\sqrt{1-\rho}}\right) \cdot (1 - RR)}$$

- L_{exposure}: total loan portfolio exposure
- $Q(T)$: probability of default for each loan over horizon T
- ρ: copula correlation parameter (pairwise asset correlation)
- X: confidence level (e.g. 0.999 for 99.9% Credit VaR)
- $\Phi^{-1}(X)$: standard normal quantile, captures the systemic stress
- RR: recovery rate (fraction recovered when a loan defaults)

Credit VaR: the loss at confidence X minus the expected loss:

$$\boxed{\text{Credit VaR} = V(X,T) - EL}$$

The formula's elegance: a single closed-form expression for the loss percentile of a large credit portfolio. The systemic factor stress $\Phi^{-1}(X)$ replaces the actual F at the desired confidence level. As $\rho \to 0$ the correlation term vanishes and the formula reduces to $V \to L \cdot Q(T) \cdot (1 - RR)$ (no tail risk, pure expected loss). As $\rho \to 1$ all firms default together, the tail explodes. This is why Credit VaR is extremely sensitive to the correlation assumption.

7.5 Economic Value (EV) and Economic Capital (EC)

⋆ **Economic Value (EV).** A firm's balance sheet measured from a market-based, economic perspective. Defined as the gap between the market value of assets and the market value of liabilities:

$$\boxed{EV = MVA - MVL \quad \text{(current market conditions)}}$$

- **Assets** held at **traded market value**; non-traded assets and liabilities valued via models (e.g., **risk-neutral** pricing for insurance contracts)
- Avoids **book-value accounting**, where assets sit at amortized cost until impaired and hide real economic gains/losses
- **More volatile** than book value because it moves with market conditions every day
- **EV can be negative**, unlike a stock price; negative EV means liabilities exceed assets at fair market levels

The point of EV is to ask: "if we had to liquidate today at fair market prices, what would we be left with?" Book accounting can mask trouble for years (an underwater bond held to maturity still shows par value on the books); EV refuses that fiction. SVB is the canonical case: book value showed healthy assets, EV showed massive unrealized losses on long-dated UST holdings once rates rose. The downside of EV is volatility, a firm whose EV swings with daily markets may look unstable even when its long-run economics are fine.

⋆ **Economic Capital (EC).** The same balance sheet identity, but evaluated under a **stressed market environment** chosen to align with the firm's risk appetite:

$$\boxed{EC = MVA - MVL \quad \text{(after a projected severe economic shock)}}$$

- Goal: keep $EV > 0$ after the stress, i.e., remain solvent through a severe adverse event
- Common severity: **99.5% adverse shock** (the Solvency II / insurance industry standard)
- $EV < 0$ after shock implies the firm is **economically bankrupt** and may need regulatory intervention even if its accounting books still look healthy

EC is just EV recomputed under a bad scenario. It answers: "how much economic cushion would we have left after a 1-in-200-year event hits the markets?" This is the firm's internal alternative to *regulatory capital, anchored in market reality rather than rule-of-thumb risk weights. The required capital is whatever amount keeps EC above zero under the chosen stress.*

⋆ **Economic Factor Risks (Risk Bucketing).** Total economic risk is split into separate **factor buckets**, each shocked and modeled independently: **Interest Rates**, **Equity**, **Credit Spreads**, **Implied Volatility**, and **Foreign Exchange**.

Two aggregation methods for combining factor risks into total EC:

- **Full valuation:** run a single joint stress where all factors are shocked together, then re-price every position. Captures all nonlinear cross-factor effects but expensive.
- **Correlation-matrix aggregation:** shock each factor separately to get marginal risks, then combine using a correlation matrix. Cheap and closed-form but linear in cross-factor effects.

Sub-bucketing within a factor (example: interest rates).

- Split rate risk into **PCA factors** (level, slope, curvature) instead of using a single parallel shift
- Split into **up-rate vs down-rate** shocks to capture asymmetric exposures (e.g., callable bonds, MBS prepayment)

Bucketing lets each risk type be modeled with its own appropriate shock structure (rates use PCAs, equity uses log-normal moves, credit uses spread widening) and then combined at the end. Without this separation you would need one giant joint distribution covering everything at once, which is much harder to calibrate and harder to interpret. The trade-off is that correlation assumptions between buckets become a critical, hard-to-estimate input that drives the final number.

⋆ **Correlation-Matrix Aggregation Formula.** For each economic factor k at confidence α, apply a marginal shock and revalue the portfolio. The set of value changes feeds the EC formula.

- V_0^{Port}: portfolio economic value at $t = 0$ under **current** conditions
- V_0^k: portfolio economic value at $t = 0$ under a **shock to factor k only**
- $M_k = V_0^k - V_0^{\text{Port}}$: **change in value** from the factor-k shock
- $M = (M_1, M_2, \dots, M_K)$: row vector stacking all factor value changes
- Σ_{EC}: **correlation matrix** between economic factors

Economic Capital: $\boxed{EC = \sqrt{M\,\Sigma_{EC}\,M^T}}$

Pipeline:

1. Compute baseline V_0^{Port} under current conditions
2. For each factor $k = 1, \dots, K$: apply a marginal shock to factor k only, revalue the portfolio to get V_0^k
3. Compute the value-change vector entries: $M_k = V_0^k - V_0^{\text{Port}}$
4. Stack into the vector $M = (M_1, \dots, M_K)$
5. Plug into the quadratic form: $EC = \sqrt{M\,\Sigma_{EC}\,M^T}$

This is the same Markowitz portfolio-variance formula generalized to risk factors instead of asset returns. Each M_k is the marginal dollar risk from factor k; Σ_{EC} controls how those marginals combine. With $\Sigma_{EC} = I$ (uncorrelated factors), the formula reduces to $\sqrt{\sum_k M_k^2}$, pure Pythagorean addition. With Σ_{EC} all ones (perfectly correlated factors), it collapses to $|\sum_k M_k|$, simple linear addition. Real Σ_{EC} sits between, giving a ***diversification benefit*** *relative to the worst-case linear sum.*

Efficiency point: only $K + 1$ revaluations are needed (one baseline plus one per factor), not a full Monte Carlo with thousands of joint scenarios. The trade-off is that all interaction effects between factors enter ***only*** *through the correlation matrix, missing any nonlinear cross-factor effects.*

⋆ **Interest Rate Aggregation via PCA (Special Case).** Decompose interest rate risk into the first **3 PCAs (level, slope, curvature)** and use these as the rate sub-buckets, replacing the standard single parallel shift.

PCs are orthogonal by construction, so their pairwise correlations are zero. The correlation matrix among rate PCs is the identity:

$$\boxed{EC_{\text{Rates}} = \sqrt{\Delta M \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \Delta M^T} = \sqrt{M_1^2 + M_2^2 + M_3^2}}$$

where M_1, M_2, M_3 are the portfolio value changes under PC1, PC2, and PC3 shocks respectively.

The orthogonality of PCs does the heavy lifting here: because level/slope/curvature shocks are statistically independent by construction, their dollar impacts simply add in quadrature. No correlation guessing needed for the inside of the rate bucket. This is exactly why PCA-based stress testing is preferred over ad hoc parallel-and-twist shifts: you get a clean orthogonal decomposition for free, and the aggregation formula collapses to a Euclidean norm.

⋆ **Combined Factor Aggregation (Block-Structured Form).** Combine the rate PCA block (orthogonal inside) with cross-correlations to other factors (equity, credit, IV, FX). The full correlation matrix is block-structured:

$$\boxed{VaR = \sqrt{\Delta M \begin{pmatrix} 1 & 0 & 0 & & & \\ 0 & 1 & 0 & & & \\ 0 & 0 & 1 & & & \\ \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & 1 & & \\ \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & 1 & \\ \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & \rho_{i,j} & 1 \end{pmatrix} \Delta M^T}}$$

- **Top-left 3×3 block:** identity, the 3 rate PCs (orthogonal, no estimated correlations needed inside)
- **Bottom rows:** other factors (equity, credit, IV, FX), correlated both with each other and with the rate PCs
- **Off-diagonal $\rho_{i,j}$:** cross-correlations between rate factors and non-rate factors, estimated from historical data or expert judgment

The matrix splits cleanly into "stuff that's orthogonal by math" (the PCA block) and "stuff that needs correlation assumptions" (everything else). Using PCA inside the rate bucket means you only estimate correlations across factor categories, not within rates, which is where most of the modeling judgment is hard to defend. This is also why PCA shocks tend to be regulators' preferred rate stress: the orthogonality is mathematical, not assumed.

Compared to Monte Carlo for capital aggregation:

- **No simulation needed:** only $K + 1$ portfolio revaluations (one baseline plus one per shocked factor)
- **Closed-form aggregation:** the result follows from a single matrix multiplication
- **Transparent attribution:** each factor's marginal contribution is explicit, so capital can be allocated back to its source business line

This is the production workhorse for capital aggregation at insurers and banks: fast, auditable, and easy to update when correlations or shock magnitudes change. The price you pay is missing nonlinear cross-factor effects, the things only a full joint Monte Carlo would catch (e.g., how an extreme rates move feeds into changes in equity volatility). For most balance sheets the linear approximation is good enough; portfolios heavy in cross-factor optionality may need a Monte Carlo overlay on top.

7.6 Asset Liability Management (ALM)

⋆ **ALM Core Idea.** Match the **interest rate sensitivities** of one portfolio (assets) against another (liabilities) so that joint movements under a rate shock cancel. The goal: control the joint sensitivity of assets and liabilities under a chosen valuation metric (PV, GAAP income, or capital) and a chosen rate shock structure.

- **Pension funds:** match projected benefit claims against assets held in reserve
- **Banks:** match book of loans (often with embedded options like prepayment) against deposit base
- **Insurance companies:** match policyholder liabilities against the supporting asset portfolio

Three sensitivity metrics depending on what you care about:

- **Mark-to-Market** (economic): match changes in market values of assets and liabilities
- **GAAP income** (financial): match changes in reported earnings under a rate move
- **Capital** (regulatory): minimize changes to the firm's capital position

Two implementation methods:

- **Construct matching portfolios** with equivalent sensitivities (duration/convexity matching)
- **Trade interest-sensitive derivatives** (UST futures/options, SOFR swaps) to layer hedges on top

ALM is the answer to the question: "what set of assets do we need to hold so that whatever happens to rates, our balance sheet stays roughly intact?" Different institutions pick different metrics depending on what they are forced to report on. A bank may care most about quarterly GAAP income; a pension fund cares about long-run economic surplus; a regulated insurer cares about solvency capital. The metric chosen drives which hedging method is most appropriate.

⋆ **Cash-Flow Matching (Strictest Hedge).** Hedge a liability with known fixed cash flows by holding an asset portfolio (e.g., fixed-coupon bonds) whose cash flows arrive at the same dates with the same magnitudes:

$$\text{Hedge CF}(t) = \text{Liability CF}(t) \quad \forall t$$

- If matched exactly: **perfect replicating portfolio**, zero residual risk under any rate scenario
- Easiest for **deterministic cash flows** (fixed or floating with no optionality)
- Visually: liability bars and hedge bars cancel at every maturity, leaving a near-zero **Net** bar

Cash-flow matching is the gold standard because it dominates every other hedge: if every dollar coming in matches every dollar going out at the exact same time, no rate path can hurt you. The catch is realism. Most real liabilities have prepayment options, surrender options, or behavioral elements (deposits, lapses) that make the future cash flows themselves stochastic, so an exact static match is impossible.

⋆ **Cash-Flow Matching Optimization.** When exact matching is impossible (most real cases), solve a least-mismatch optimization. Choose hedge weights α_i on assets H^i to minimize the present-value-weighted absolute mismatch between hedge cash flows and liability cash flows.

- L_t: liability cash flow at time t (over horizon T)
- H^i: hedge asset i producing cash flow H^i_t at time t (across K assets)
- α_i: position weight (units) in hedge asset i (decision variable)
- $P(0,t)$: discount factor from t back to today (PV weight)
- E: number of cash-flow scenarios (for stochastic CFs)

Objective:

$$\boxed{\min_{\alpha_i} \sum_{e=1}^{E} \sum_{t=1}^{T} P(0,t) \left| \sum_{i=1}^{K} \alpha_i H^i_{t,e} - L_{t,e} \right|}$$

Optionally impose non-negativity ($\alpha_i \geq 0$, no shorting) or position-size limits.

This is a **static hedge**, weights are chosen once and held. No rebalancing. For dynamic strategies, see duration/convexity matching below.

The formula reads: at every time t in every scenario e, compute net cash mismatch (hedge inflows minus liability outflows), take its absolute value (penalize over- and under-funding equally), and discount it back to today. Sum across all times and scenarios to get total weighted mismatch, then choose hedge weights to minimize this number. Using the absolute value rather than squared error ensures the objective stays linear (LP-solvable) and treats a \$1 surplus the same as a \$1 shortfall, matching the economic reality that both create funding/reinvestment problems.

⋆ **Why Cash-Flow Matching Often Fails in Practice.** Cash-flow matching eliminates rate risk **only if** cash flows and discounting are deterministic.

- It is **stronger than moment matching** because it matches the entire cash-flow path, not just summary statistics
- But it is **not very practical** when assets or liabilities contain optionality (callable bonds, MBS prepay, insurance lapses, demand deposits)

The pragmatic alternative: match interest-rate sensitivities (not full cash flows).

- **First and second order: duration/convexity matching**
- **Partial sensitivities: key rate durations and convexities**

The hierarchy of approaches matters: cash-flow matching is the cleanest in theory but rarely feasible; duration matching trades exactness for tractability; key rate durations refine duration matching by recognizing that not all parts of the curve move together. Each step down the ladder gives up some precision in exchange for being implementable on real, optionality-laden balance sheets.

⋆ **Bond Interest Rate Sensitivity (Taylor Expansion).** For a bond with price $B(r,t)$ (fixed or floating coupon), expand around current rate and time:

$$\boxed{\Delta B(r,t) = \frac{\partial B}{\partial t}\Delta t + \frac{\partial B}{\partial r}\Delta r + \frac{1}{2}\frac{\partial^2 B}{\partial r^2}(\Delta r)^2 + \cdots}$$

Two distinct sensitivity types:

- **Passage of time** (Δt): captures pull-to-par and theta. **Unhedgeable**, time always moves forward at the same rate for everyone.
- **Changes in interest rates** (Δr): captures the price impact of rate moves. **Hedgeable locally** via duration (first-order) and convexity (second-order).

The Taylor expansion is the foundation of every fixed-income hedging metric. The $\partial B/\partial r$ term defines duration; the $\partial^2 B/\partial r^2$ term defines convexity. Time decay is its own thing, real but unavoidable, you cannot hedge against the calendar. The third-order and higher terms are usually ignored for small rate moves; they re-enter discussions only when you stress test with large shifts (50bp+).

⋆ **Portfolio Duration and Convexity.** At the portfolio level, duration and convexity are simple **value-weighted averages** of the component instruments:

$$\boxed{D_{\text{Port}} = \sum_{i=1}^{N} w_i D_i} \qquad \boxed{C_{\text{Port}} = \sum_{i=1}^{N} w_i C_i}$$

- N: number of instruments in the portfolio
- w_i: market value share of instrument i ($\sum w_i = 1$)
- D_i, C_i: duration and convexity of instrument i

Immunization (the ALM payoff): match $D_{\text{Assets}} = D_{\text{Liabilities}}$ and $C_{\text{Assets}} = C_{\text{Liabilities}}$ so that the surplus $A - L$ is **locally insensitive** to small parallel yield shifts. Examples: pension funds matching pension liabilities, insurance companies matching policy liabilities.

Linearity of duration is what makes ALM tractable: you do not need to model the portfolio as a single instrument, you just take the value-weighted average of component sensitivities. The word "locally" is critical, immunization protects against ***small*** *parallel shifts only. Large moves trigger the convexity term, and non-parallel moves (steepening, flattening) leak through entirely. This is the gap that key rate durations are designed to plug.*

⋆ **Key Rate Duration (KRD).** Standard duration assumes a **parallel shift** of the entire yield curve, which is unrealistic. Key Rate Duration measures the portfolio's sensitivity to a shock applied to one segment of the curve at a time, holding all other rates fixed.

Choice of key rate buckets:

- **Individual tenors:** 1Y, 2Y, 5Y, 10Y, 30Y, etc.
- **Ranges of tenors:** short, medium, long; or pre-specified bands like 0-1Y, 1-3Y, 3-5Y, 5-10Y, 10Y+
- **Critical:** chosen nodes/ranges must form a **basis spanning the entire term structure**, no gaps

$$\boxed{D = -\frac{1}{PV}\frac{\partial PV}{\partial y}} \qquad \boxed{KRD_j = -\frac{1}{PV}\frac{\partial PV}{\partial y_j}}$$

Key additivity property: if the chosen key rates partition a parallel shift, $\Delta y(t) = \sum_{j=1}^{K} \Delta y_j(t)$ for all t, then on a **first-order basis**:

$$\boxed{D = \sum_{j=1}^{K} KRD_j}$$

KRD is duration localized to one segment of the curve. The intuition: a 30-year bond has zero sensitivity to a shock in the 2-year rate, but standard duration would treat all rate movements as if they affected every cash flow proportionally. KRDs let you see exactly which part of the curve drives your portfolio risk, a 30-year liability with 30Y KRD = 15 and zero elsewhere is fundamentally different from a portfolio with KRD spread evenly across all tenors, even if both have the same total duration.

⋆ **KRD Conceptual Example: How the Calculation Works.** Given a bond with cash flows over several years and a base discount curve:

- Compute **Base Value (A)**: PV under the original spot curve

- Compute **Parallel Shift Value (B)**: PV after shifting every spot rate up 1%, gives modified duration $D = -(B - A)/A/1\%$
- Compute **KR1 Value (C)**: PV after shifting only the 1Y rate up 1%, all other rates unchanged
- Compute **KR2 through KRn**: PV after shifting only the 2Y, 3Y, ... rate up 1% one at a time
- Each $KRD_j = -(\text{Value}_j - A)/A/1\%$ measures fractional PV change per unit shift at that key rate

What the output tells you:

- Each KRD_j is the bond's exposure to that specific bucket of the curve
- Sum of KRD_j approximately equals modified duration (exact for full revaluation under a parallel partition shift, since the second-order convexity term is what creates the small gap for finite shocks)
- % of duration attributable to each bucket = KRD_j/D, useful for risk attribution

The procedure is mechanical: shift one rate at a time, revalue, take the percentage change. What makes it powerful is the ***decomposition****: a bond's total interest rate risk gets split into a vector of bucket risks, and you can hedge each bucket independently. A long bond gets most of its KRD weight at the long end; a short bond at the short end; a barbell shows weight at both ends with little in the middle. This level of granularity is impossible with a single duration number.*

⋆ **KRD Identities for Finite Shocks.** For **infinitesimal** shocks, KRD is exactly additive: $D = \sum_j KRD_j$. For **finite** yield curve shocks, second-order terms enter and exact additivity breaks.

Second-order Taylor expansion of PV under multi-bucket shock:

$$\boxed{\Delta PV \approx \sum_{j=1}^{K} \frac{\partial PV}{\partial y_j}\Delta y_j + \frac{1}{2}\sum_{i=1}^{K}\sum_{j=1}^{K} \frac{\partial^2 PV}{\partial y_i \partial y_j}\Delta y_i \Delta y_j}$$

Substituting the KRD definition:

$$\boxed{\Delta PV \approx PV \sum_{j=1}^{K} KRD_j \Delta y_j + \frac{1}{2}\sum_{i,j} \frac{\partial^2 PV}{\partial y_i \partial y_j}\Delta y_i \Delta y_j}$$

When KRDs sum to modified duration exactly:

- **Nonlinearity** comes from the convexity of exponential discounting: PV is $\sum CF_t \cdot e^{-y_t t}$, which is curved in y
- **Second-order terms vanish only under infinitesimal shocks**; for finite shocks they show up as a small residual
- **Full revaluation under a partition** (as opposed to pure linear approximation) sums exactly for any given shock size, because revaluation captures the convexity term naturally

The takeaway: when you shock individual rates and revalue from scratch, the KRDs you get back will sum exactly to modified duration even under a 1% shock, because every nonlinear bit gets baked into the revaluation. But if you build KRDs from analytical first derivatives only, the sum drifts off duration as shock size grows. This is why most practitioners use full revaluation methods, the convexity adjustment is invisible but correct.

⋆ **SVB ALM Mishap (2023): A Cautionary Case Study.** Silicon Valley Bank's collapse is the textbook example of what happens when ALM is done badly. Three failures stacked on top of each other.

Failure 1: Duration mismatch.

- Funded by **short-term deposits** (concentrated in Silicon Valley tech firms; balances mostly above the $250K FDIC insurance cap, so uninsured and flight-prone)
- Bought **long-term assets** (UST bonds purchased during the low-rate era to pick up yield)
- Created **large negative convexity**: as rates rose, asset values fell sharply while the deposit base became increasingly skittish

Failure 2: Liability modeling misestimated.

- Assumed deposits were **stickier** (longer effective duration) than reality
- No mechanism actually prevented rapid withdrawal
- Once depositors lost confidence, the model assumptions collapsed and a **classic bank run** began

Failure 3: Management actions made it worse.

- Optimized for **net income** rather than capital preservation
- Originally held pay-fixed swaps (a hedge against rising rates), but these caused **negative income during the low-rate period** because the assets being hedged were classified as **Held to Maturity** and thus showed no mark-to-market gain to offset the swap drag
- **Removed the swaps just before/at the start of the rising rate cycle**, leaving the bank fully exposed when rates spiked and depositors fled

SVB violated almost every ALM principle in this section. They had massive duration mismatch (the asset side moved very differently from the liability side under a rate shock), they used optimistic behavioral assumptions about liability stickiness that did not hold under stress, and they removed the one hedge that would have saved them because it was hurting accounting income in the short run. The lesson is that ALM is fundamentally about ***economic*** *sensitivity, hedging based on accounting categorizations (HTM versus AFS) rather than economic exposure leaves the firm vulnerable to the very stress the hedge was supposed to protect against.*

7.7 Commodity Derivatives

⋆ **Commodity Markets and Stylized Facts.** Main markets: **Energy** (Crude Oil, Natural Gas, largest), **Precious Metals**, **Industrial Metals**, **Agriculture**. Traded on CME, ICE, LSE, SHFE, DCE, ZCE.

Two facts that drive the model:

- Commodity prices **mean revert** (analogous to interest rates), unlike equities
- Model must **calibrate to current spot and futures curve** (no arbitrage at $t = 0$)

Mean reversion is the structural difference: commodity prices reflect supply/demand fundamentals that pull prices back toward production cost. Equities have no such anchor. This is why commodities use Hull-White-style models, not GBM.

⋆ **Commodity Price Process.** Mean-reverting process on $\ln S$ so S stays positive:

$$\boxed{d\ln(S) = [\theta(t) - a\ln(S)]\,dt + \sigma\,dW_t}$$

- a: **mean reversion speed** (how strongly prices get pulled back)
- σ: **log-price volatility** (constant)

- $\theta(t)$: **time-dependent drift function**, the only free function, calibrated numerically to fit the futures curve

Same form as Hull-White on rates. We build a symmetric zero-mean tree in $\ln S$ *space, then shift each time slice by an offset* α_t *so the tree's expectation matches the market futures price. The set of offsets* $\{\alpha_t\}$ *is the numerical realization of* $\theta(t)$*.*

⋆ **Tree Framework: Variables and Geometry.** The tree lives in $\ln S$ space, with **node levels** indexed by integer j relative to a center.

Tree parameters:

$$\boxed{\Delta R = \sigma\sqrt{3\Delta t}} \qquad \boxed{j_{\max} = \text{ceil}\left(\frac{0.184}{a\,\Delta t}\right)}$$

- ΔR: spacing in $\ln S$ between adjacent levels (chosen to match Brownian variance over Δt)
- j: **node level** ($j = 0$ is center, $j = +1$ one above, etc.). $\ln S$ value at level j is $j\,\Delta R$
- $j_{\max}$: **truncation level** at which branching geometry must change
- Above $+j_{\max}$: branches flip to "mid/down/down-down" (no more pure up)
- Below $-j_{\max}$: branches flip to "up-up/up/mid" (no more pure down)

The truncation $j_{\max}$ *is what enforces mean reversion in the tree. Without it, branches drift outward forever. The constant* $0.184 = 1 - \sqrt{2/3}$ *is exactly the level where the standard probabilities would otherwise turn negative; the geometry has to flip there.*

⋆ **Risk-Neutral Probabilities by Moment Matching.** At each node, the probabilities (p_u, p_m, p_d) are chosen so the discrete tree step reproduces the first two moments of the continuous SDE.

Three constraints, three unknowns:

1. Sum to one: $p_u + p_m + p_d = 1$
2. Match local mean: $p_u\,\Delta R - p_d\,\Delta R = -aj\,\Delta R\,\Delta t$
3. Match variance: $p_u\,\Delta R^2 + p_d\,\Delta R^2 - (\text{mean})^2 = \sigma^2\,\Delta t$

Let $M = -aj\,\Delta t$ (local drift in units of ΔR). Solving gives:

Interior node ($|j| < j_{\max}$):

$$\boxed{p_u = \tfrac{1}{6} + \tfrac{M^2+M}{2}, \quad p_m = \tfrac{2}{3} - M^2, \quad p_d = \tfrac{1}{6} + \tfrac{M^2-M}{2}}$$

Top edge ($j = +j_{\max}$, branches mid/down/down-down):

$$\boxed{p_u = \tfrac{7}{6} + \tfrac{M^2+3M}{2}, \quad p_m = -\tfrac{1}{3} - M^2 - 2M, \quad p_d = \tfrac{1}{6} + \tfrac{M^2+M}{2}}$$

Bottom edge ($j = -j_{\max}$): mirror of top edge (swap $p_u \leftrightarrow p_d$).

Probabilities depend only on level j*, not on time or on the offsets. Same level* $\Rightarrow$ *same probabilities across all time slices. They are computed once at tree construction and reused everywhere. Edge formulas exist because the branching geometry physically changes at* $\pm j_{\max}$ *to enforce mean reversion.*

⋆ **Calibration: Determining Offsets α_t.** Apply a time-dependent shift α_t to each time slice so the tree's risk-neutral expectation matches the market's futures price. Actual price at any node:

$$\boxed{S_{\text{node}} = \exp\Big(\alpha_t + j\,\Delta R\Big)}$$

Calibration target (no-arbitrage at $t = 0$):

$$\boxed{F(0,t) = \widehat{E}[S(t)] = \sum_{\text{nodes at } t} P_{\text{reach}}(\text{node}) \cdot \exp(\alpha_t + j\,\Delta R)}$$

where $P_{\text{reach}}(\text{node})$ is the total probability of reaching that node from the root (sum of all path probabilities leading to it).

Procedure (sequential, one slice at a time):

1. Compute reaching probability for each node at time t using the transition probabilities
2. Solve $\sum_{\text{nodes}} P_{\text{reach}} \cdot e^{\alpha_t + j\Delta R} = F(0,t)$ for α_t (1D root-finding in one variable)
3. Move to next time slice and repeat

Time 0 (single root node, $P_{\text{reach}} = 1$):

$$\alpha_0 = \ln(S_0)$$

Each α_t is the level shift that aligns the tree's expectation with the observed futures price at that horizon. The sequence $\{\alpha_0, \alpha_1, \ldots\}$ is the numerical $\theta(t)$. After calibration, the tree is arbitrage-free with respect to the entire futures curve.

⋆ **Consistency Check.** After calibration, at every non-edge node the local expectation must reproduce the next-period futures price:

$$\boxed{p_u\, S_{\text{up}} + p_m\, S_{\text{mid}} + p_d\, S_{\text{down}} = F(0, t + \Delta t)}$$

If this fails, the calibration is wrong.

⋆ **Option Valuation by Backward Induction.** Once the tree is calibrated, value any path-independent option by walking backward from maturity to root.

Procedure:

1. **Terminal payoff** at maturity nodes: e.g., $V = \max(K - S, 0)$ for a put, $\max(S - K, 0)$ for a call
2. **Continuation value** at each earlier node:
$$V_{\text{cont}} = e^{-r\,\Delta t}\Big(p_u V_{\text{up}} + p_m V_{\text{mid}} + p_d V_{\text{down}}\Big)$$
3. **Exercise check** (American only): $V_{\text{exercise}} = \max(\text{intrinsic}, 0)$
4. **Take the max**: $V_{\text{node}} = \max(V_{\text{cont}}, V_{\text{exercise}})$
5. Repeat back to root; root value is the option price

Discount at the risk-free rate r because probabilities are already risk-neutral (calibrated to futures, which are risk-neutral by construction). For European options, skip the exercise check. The same tree handles puts, calls, knockouts, anything path-independent.

⋆ **Two-Factor Extensions.** Single-factor model misses curve shape and vol dynamics. Two extensions:

Gibson-Schwartz (stochastic convenience yield δ_t):

$$\boxed{\begin{aligned} dS &= [r - \delta_t] S\, dt + \sigma_S\, S\, dW_t^1 \\ d\delta_t &= \kappa[b - \delta_t]\, dt + \sigma_\delta\, dW_t^2 \end{aligned}}$$

Convenience yield δ_t is the implicit benefit of holding physical commodity (avoiding stockouts), itself mean reverting. Drives futures curve shape (contango vs backwardation).

Eydeland-Geman (stochastic volatility):

$$\boxed{\begin{aligned} dS_t &= \kappa(\mu(t) - S_t)\, dt + \sqrt{V_t}\, dW_t^S \\ dV_t &= \alpha(\bar{V} - V_t)\, dt + \sigma_V \sqrt{V_t}\, dW_t^V \\ \text{corr}&(dW_t^S, dW_t^V) = \rho \end{aligned}}$$

Heston-style stochastic variance, captures vol clustering and smile in commodity options.

Single-factor models calibrate to the futures curve but misprice options sensitive to curve shape evolution (calendar spreads) or vol dynamics (cliquets). Two-factor models add a second source of randomness at the cost of harder calibration and a 2D tree or PDE.

7.8 Real Options

⋆ **Real Options.** Method to value **non-financial options embedded in real assets**, e.g., a company's option to expand, abandon, defer, or contract a physical project. Apply derivatives valuation techniques to corporate decisions.

Methodology:

- Define stochastic processes for the key underlying variables (commodity price, sales, costs)
- Use **risk-neutral valuation** (calibrate to traded markets where possible)
- Apply **risk premium adjustments** for non-traded quantities (sales volume, demand)

Real options is the alternative to standard NPV. NPV uses a single "best estimate" cash flow path and one discount rate; real options recognizes that management has flexibility to react as uncertainty resolves. A project that looks unprofitable on NPV may be valuable once you account for the ability to abandon if things go badly or expand if they go well.

NPV Fails for Optionality. The standard NPV approach is to discount expected cash flows at the asset's real-world required return. For options this breaks down because the **required return on an option is not the same as on its underlying**.

Setup: stock at S_0 goes to S_u or S_d in one period. Strike K. Risk-free rate r, real-world required return μ.

Risk-neutral probability (used with r as discount rate): $\boxed{S_u\, p + S_d(1-p) = S_0\, e^{r\,\Delta t} \quad \Rightarrow \quad p}$

Real-world probability (used with stock's μ): $\boxed{S_u\, p^* + S_d(1-p^*) = S_0\, e^{\mu\,\Delta t} \quad \Rightarrow \quad p^*}$

Implied real-world discount rate for the option:

- Price the option correctly using risk-neutral p and r
- Then back-solve for the rate that, paired with real-world p^*, reproduces that price
- Result: **call discount rate is much higher than μ; put discount rate can be negative**

*Calls amplify upside, so they are riskier than the stock and require a higher return. Puts pay off when the stock falls, so they hedge market risk and require a **negative** expected return in a CAPM sense. There is no single "project discount rate" that works across all options. This is why real options must be priced under the risk-neutral measure, not by guessing an appropriate hurdle rate.*

⋆ **Types of Real Options.**

- **Abandonment:** terminate project early if it becomes unprofitable (analogous to American put)
- **Expansion:** scale up production when conditions are favorable (analogous to American call)
- **Contraction:** scale down production when conditions worsen
- **Option to defer:** delay committing capital until uncertainty resolves
- **Option to extend life:** prolong a project past its original horizon

Each maps to a familiar derivatives payoff, just with project value (not stock price) as the underlying. The key is that management is not locked into the original plan; flexibility itself has value.

⋆ **Valuation Framework (General Procedure).** For a project whose cash flows depend on a commodity, equity, or other modeled underlying:

Pipeline:

1. **Build the underlying tree** (e.g., commodity Hull-White trinomial), calibrated to current spot and futures
2. **At each node, compute the period cash flow**: $\mathrm{CF}_{\text{node}} = \text{units} \cdot (S_{\text{node}} - \text{var cost}) - \text{fixed cost}$
3. **Roll backward**: at each node, value = current $\mathrm{CF}_{\text{node}} + e^{-r\Delta t}\, E[\text{value at children}]$
4. **At nodes with optionality**: compare continuation value vs exercise payoff, take the max
5. **Project value at root** = tree root value − initial investment

Risk-neutral pricing on the calibrated tree handles all the discounting correctly. The futures curve is risk-neutral by construction, so cash flows derived from prices in the tree are already risk-adjusted. Discount each step back at r, not at any project-specific hurdle rate.

⋆ **Base Project (No Optionality).** Discount expected cash flows at the risk-free rate, using futures prices as the risk-neutral expected commodity prices:

$$\boxed{NPV_{\text{base}} = -I_0 + \sum_{t=1}^{T} P(0,t)\left[\text{units} \cdot (F(0,t) - \text{var cost}) - \text{fixed cost}\right]}$$

- I_0: initial investment outlay
- $F(0,t)$: t-year futures price (already risk-neutral expected spot)
- $P(0,t) = e^{-rt}$: risk-free discount factor

This is just NPV with the right discount rate and the right expected price. If the result is negative, the project loses money on a static commit-and-run basis. The same number falls out of full backward induction on the calibrated tree (sanity check), confirming the tree reproduces the futures curve.

⋆ **Option to Abandon (American Put on Project).** Allow management to walk away at any node where continuation value is negative.

Backward induction with abandonment:

1. Terminal nodes: value = 0 (project ends naturally)
2. At each earlier node, compute **continuation value**:

$$V_{\text{cont}} = \mathrm{CF}_{\text{node}} + e^{-r\Delta t}\Big(p_u V_{\text{up}} + p_m V_{\text{mid}} + p_d V_{\text{down}}\Big)$$

3. **Abandon value** = 0 (or salvage value S if any)
4. Take the max: $V_{\text{node}} = \max(V_{\text{cont}}, 0)$

Project value with option: $\boxed{NPV_{\text{abandon}} = NPV_{\text{base}} + \text{Abandonment Option Value}}$ The option value is the difference between the tree-based valuation **with** the abandon-or-continue choice and the no-optionality NPV.

Abandonment is most valuable when the underlying is in deep loss territory. Lower-branch nodes (where commodity prices are far below break-even) typically trigger abandonment; upper branches keep going. The option turns negative-value scenarios into zero, capping downside while leaving upside intact, exactly the put payoff structure.

⋆ **Option to Expand (American Call on Incremental Capacity).** Pay extra capital to scale up production when conditions are favorable.

Setup at each node:

- Production scales by factor ϕ (e.g., +20%)
- Fixed costs scale with the same factor
- Variable cost per unit unchanged
- Expansion requires additional investment I_{exp}

Expansion payoff at a node: $\boxed{V_{\text{exp}} = \phi\,(\text{remaining project value at node}) - I_{\text{exp}}}$

Backward induction with expansion:

1. Compute continuation value as before
2. Compute exercise (expansion) payoff
3. Take the max: $V_{\text{node}} = \max(V_{\text{cont}}, V_{\text{exp}})$

Expansion is most valuable on upper branches where commodity prices are high; the incremental capacity captures additional upside that would otherwise be left on the table. Like a call, the option ignores poor outcomes (you just don't expand) and pays off in good ones.

⋆ **Non-Traded Underlyings: MPR Adjustment.** Some real options have underlyings with **no liquid traded market** (sales volume, demand, customer growth). Risk-neutral valuation requires adjusting the real-world drift using the **Market Price of Risk (MPR)**.

Underlying dynamics (real-world):

$$\frac{dX}{X} = \mu_{RW}\,dt + \sigma_X\,dW_t^X$$

Market dynamics (real-world):

$$\frac{dM}{M} = \mu_m\,dt + \sigma_m\,dW_t^m, \qquad E[dW_t^X\,dW_t^m] = \rho\,dt$$

Risk-neutral drift adjustment:

$$\boxed{\mu_{RN} = \mu_{RW} - \sigma_X\,\rho\,\lambda_m}$$

where λ_m **is the MPR (Sharpe ratio of the market) and the product $\sigma_X\,\rho\,\lambda_m$ is the asset-specific risk premium.**

Asset valuation:

- **Traded asset** (stock, commodity): risk-neutral drift is r (or futures growth rate), discount cash flows at r
- **Non-traded asset**: risk-neutral drift is $\mu_{RW} - \sigma_X\rho\lambda_m$, discount cash flows at r after this adjustment

The market is incomplete when X is non-traded: no portfolio of liquid assets perfectly hedges X, so risk-neutral pricing is not unique. The MPR adjustment is the CAPM-style correction: subtract the risk premium that the market would assign to X based on its correlation with the broader market. Once the drift is adjusted, the rest of the valuation proceeds normally at r.

⋆ **MPR Estimation (CAPM-Style).** Treat λ_m as the market Sharpe ratio: $\boxed{\lambda_m = \frac{\mu_m - r}{\sigma_m}}$

Then the implied price of risk for non-traded X: $\boxed{\lambda_X = \rho\,\lambda_m = \rho\,\frac{\mu_m - r}{\sigma_m}}$

- ρ: correlation between % changes in X and the market
- σ_m: market return volatility
- μ_m: real-world expected market return
- r: risk-free rate

Same logic as CAPM beta, applied to a non-traded driver. If X is uncorrelated with the market ($\rho = 0$), then $\lambda_X = 0$ and no risk-adjustment is needed (real-world drift equals risk-neutral drift). If X is highly correlated with the market, the risk premium is large and the risk-neutral drift falls noticeably below the real-world drift, lowering project value. The whole adjustment is a single-factor CAPM applied to whatever "the underlying really is".

7.9 Joint Equity and Interest Rate Model

⋆ **Joint Risk-Neutral Dynamics.** Equity follows log-normal dynamics with a **stochastic** risk-free rate; rates follow Hull-White; the two Brownian motions are correlated.

Equity (log-normal under stochastic rates): $\boxed{\frac{dS}{S} = (r(t) - q)\,dt + \sigma_S\,dZ}$

Interest rate (Hull-White): $\boxed{dr(t) = [\theta(t) - a\,r(t)]\,dt + \sigma_r\,dW}$

Correlation between the two shocks: $\boxed{E[dZ \cdot dW] = \rho\,dt}$

Hull-White $\theta(t)$ that fits the initial term structure:

$$\boxed{\theta(t) = \frac{\partial f_{0,t}^M}{\partial T} + a\,f_{0,t}^M + \frac{\sigma_r^2}{2a}\left(1 - e^{-2at}\right)}$$

where $f_{0,t}^M$ is the initial market forward curve.

Two state variables now drive equity option prices: the equity itself ***and*** *the short rate. Bond prices are no longer trivial discount factors $e^{-r(T-t)}$ but stochastic quantities $P(t,T)$. The $\theta(t)$ function is the level shift that makes the model arbitrage-free with respect to today's yield curve, automatically calibrated as soon as the forward curve is fitted.*

⋆ **Modified Black-Scholes for Stochastic Rates.** Standard Black-Scholes assumes constant r. Generalizing to stochastic rates: replace the discount factor $e^{-r(T-t)}$ with the **zero-coupon bond price** $P(t,T)$, and replace constant volatility with the **total volatility** $\sigma_{T,t}$ that integrates equity, rate, and correlation contributions.

Call and put formulas:

$$\boxed{C(t,T,K) = S(t)\,e^{-q(T-t)}\,N(d_1) - K\,P(t,T)\,N(d_2)}$$

$$\boxed{P(t,T,K) = K\,P(t,T)\,N(-d_2) - S(t)\,e^{-q(T-t)}\,N(-d_1)}$$

Modified d_1, d_2:

$$\boxed{d_1 = \frac{\ln\!\left(\frac{S(t)}{K\,P(t,T)}\right) - q(T-t) + \frac{1}{2}\sigma_{T,t}^2}{\sigma_{T,t}}, \qquad d_2 = d_1 - \sigma_{T,t}}$$

Two structural changes vs standard BS: (1) the strike is discounted using $P(t,T)$ (the actual bond price, observable in the market) instead of $e^{-r(T-t)}$; (2) volatility in d_1, d_2 is the total integrated volatility $\sigma_{T,t}$, not just $\sigma_S\sqrt{T-t}$. The form of the formula stays familiar; the inputs change.

⋆ **Total Variance Decomposition.** The total variance $\sigma^2_{T,t}$ that enters the option formula is the sum of three components: equity variance, rate variance, and their covariance.

Rate variance (from Hull-White integration):

$$\boxed{V(t,T) = \frac{\sigma_r^2}{a^2}\left[(T-t) + \frac{2}{a}e^{-a(T-t)} - \frac{1}{2a}e^{-2a(T-t)} - \frac{3}{2a}\right]}$$

Equity variance (constant vol over horizon):

$$\boxed{\sigma_S^2(T-t)}$$

Total variance (sum plus covariance):

$$\boxed{\sigma^2_{T,t} = V(t,T) + \sigma_S^2(T-t) + \frac{2\rho\,\sigma_r\,\sigma_S}{a}\left[(T-t) - \frac{1}{a}\Big(1 - e^{-a(T-t)}\Big)\right]}$$

Conceptually:

$$\boxed{\text{Total Variance} = \text{Rate Variance} + \text{Equity Variance} + \text{Covariance} = \text{Implied Variance}}$$

The covariance term is positive when $\rho > 0$ (rates and equity move together) and negative when $\rho < 0$ (the more typical regime: stocks fall when rates rise, lifting bond yields). Negative correlation ***reduces*** *total variance, lowering long-dated option prices vs the constant-rate Black-Scholes value. For short maturities the rate contribution is small and the formula collapses back toward standard BS; the difference matters most for long-dated options where rate uncertainty accumulates.*

⋆ **Calibration Strategy. Two practical approaches depending on the product:**

- **Joint calibration** (rates + equity options simultaneously): used for products with strong dependence on both, e.g., long-dated equity-linked structured notes, callable equity products
- **Sequential calibration** (fix rate parameters, calibrate equity vol): used when equity is the primary driver and rate optionality is secondary. Rate parameters (σ_r, a) pinned from swaption market; correlation ρ from history or expert view; only σ_S floats to fit equity option prices

Important: bond prices always fit by construction. The Hull-White $\theta(t)$ is calibrated to the initial forward curve, so $P(0,T)$ matches market for every T regardless of which calibration path is chosen. Calibration freedom only concerns the volatility and correlation inputs, never the term structure itself.

The implied volatility quoted in equity option markets is the ***total*** *volatility, blending equity, rate, and correlation effects. Using it directly in standard Black-Scholes (with constant r) double-counts: BS treats rates as deterministic and absorbs all the variance into pure equity vol. For short-dated options the misattribution is small; for long-dated options it can be material, which is why this joint model exists.*

7.10 EXTRA Ideas:

⋆ **Distribution Chain: Prices, Returns, and Squared Returns.** The three distributions arise sequentially from the same underlying normality assumption:

$$\underbrace{S_t}_{\text{Log-Normal}} \xrightarrow{r_t = \ln(S_t/S_{t-1})} \underbrace{r_t}_{\text{Normal}} \xrightarrow{\text{square}} \underbrace{r_t^2}_{\text{Chi-squared}}$$

Normal → Log-Normal (prices from returns). Assume log returns are normal: $r_t \sim N(\mu, \sigma^2)$. Then:

$$S_t = S_{t-1}e^{r_t} \quad \Rightarrow \quad \ln S_t \sim N(\ln S_{t-1} + \mu,\, \sigma^2)$$

S_t is the exponential of a normal random variable, which by definition is **log-normal**. This is why we model prices as log-normal: log returns being normal is a natural assumption, and exponentiation then guarantees $S_t > 0$ always.

Normal → Chi-Squared (squared returns from returns). Standardize: $z_t = r_t/\sigma \sim N(0, 1)$. Squaring a standard normal gives a **chi-squared with 1 degree of freedom** by definition:

$$z_t^2 \sim \chi^2(1) \quad \Rightarrow \quad r_t^2 = \sigma^2 z_t^2 \sim \sigma^2\chi^2(1)$$

Taking expectations using $E[\chi^2(1)] = 1$:

$$\boxed{E[r_t^2] = \sigma^2 \cdot E[z_t^2] = \sigma^2 \cdot 1 = \sigma^2}$$

Each squared return is an **unbiased** but **noisy** single draw of the true variance, with $\mathrm{Var}(r_t^2) = 2\sigma^4$ (highly dispersed). Averaging M squared returns:

$$\frac{1}{M}\sum_{i=1}^{M} r_i^2 \xrightarrow{LLN} \sigma^2$$

This is the foundation of **Realized Variance**: average enough intraday squared returns within one day and the noise cancels, leaving a precise estimate of that day's true variance.