Exotic Options

BUSS386. Futures and Options

Professor Ji-Woong Chung

Lecture Outline

- Exotic options
  - Descriptions and uses
  - Pricing with Monte Carlo simulations and binomial trees

- Chapters 26 and 27

Exotic Options

Exotic Options

- **Exotic options** are nonstandard contracts whose payoffs differ from the simple call/put structure. They are often created by slightly modifying ("tweaking") standard options.

- They are designed to solve **specific business or risk-management problems** that cannot be addressed well using ordinary European or American options.
  - Examples: protection only if a price barrier is hit, payoff based on an average price, payoff depending on multiple assets, etc.

- Exotic options are typically engineered and sold by investment banks or professional money managers.
  - Banks hedge these products dynamically and charge fees for structuring and risk management.
  - Clients (corporations, institutions, hedge funds) use them to customize risk exposure.

Exotic Options

- The goal is *not* to memorize formulas for each exotic option.
  - Instead, focus on intuition: how the payoff works and why the exotic structure is needed.

- **Key questions to ask about any exotic option:**
  1. **What problem does this exotic option solve?** (Hedging, cost reduction, targeted payoff, path dependence, barrier protection, etc.)
  2. **Can it be approximated or replicated using a portfolio of standard options?**
     - If yes, this often provides intuition and may help adapt/extend the BSM framework.
  3. **Is it cheap or expensive relative to a standard option that provides a similar payoff?** (Important for traders evaluating relative value.)
  4. **How do we price it when no closed-form formula exists?**
     - Simulation (Monte Carlo)
     - Binomial/trinomial trees
     - PDE methods
     - Replicating portfolios

Binary Options

Path-Independent Exotic Options

- **Binary (digital) options** have payoffs that depend only on whether the underlying ends above or below a strike price—not on how far it moves. They are commonly used to express **pure directional bets** or to build payoffs that jump at a threshold.
- **Cash-or-Nothing Options:** Pay a fixed cash amount if the option finishes in the money.
  - **Call:** pays \$1 if $S_T > K$, otherwise 0

    $$\text{Value} = e^{-r(T-t)} N(d_2)$$

  - **Put:** pays \$1 if $S_T < K$

    $$\text{Value} = e^{-r(T-t)} N(-d_2)$$

- **Asset-or-Nothing Options:** Pay the underlying asset itself if the option finishes in the money.
  - **Call:** pays S_T if $S_T > K$

    $$\text{Value} = S\, e^{-q(T-t)} N(d_1)$$

  - **Put:** pays S_T if $S_T < K$

    $$\text{Value} = S\, e^{-q(T-t)} N(-d_1)$$

Binary Options

Path-Independent Exotic Options

- **Key Relationship to Standard Options:** Consider the portfolio:
  - long one asset-or-nothing call with strike K, and
  - short K cash-or-nothing calls with strike K.
- **What is its payoff at maturity?**
  - If $S_T > K$: payoff $= S_T - K \times 1 = S_T - K$
  - If $S_T \leq K$: payoff $= 0 - K \times 0 = 0$
- **This is exactly the payoff of a standard European call option.**

$$\max(S_T - K, 0)$$

- Binary options are important because they serve as *building blocks*. Many exotic options—and even the standard Black–Scholes call—can be expressed as portfolios of digital (binary) instruments.

Compound Options

- **Compound options** are options whose underlying asset is *another option*. They allow investors to delay committing to a potentially expensive option position.
- Types of compound options:
  - **Call on Call (CoC):** right to buy a call option in the future.
  - **Put on Call (PoC):** right to sell a call option in the future.
  - **Call on Put (CoP):** right to buy a put option in the future.
  - **Put on Put (PoP):** right to sell a put option in the future.
- Such options are useful when:
  - the underlying option is expensive,
  - but the buyer is uncertain whether they will need it, or
  - they want to hedge a future risk that depends on an event that may or may not occur.

Compound Options

Path-Independent Exotic Options

- **Pricing:** Closed-form solutions exist for certain types (e.g., call on call), but in practice, **binomial or trinomial trees (backward induction)** are often used.
- **Illustrative Example: Using a Call on a Cap (Real Corporate Application)**
  - A company is bidding on a large infrastructure project. If they win, they must borrow **$200 million for 2 years**. Their risk: interest rates may rise between **today** and the **contract award date**.
  - Buying a 2-year **interest rate cap** today would hedge this risk—but if they *lose* the bid, the cap is useless and very expensive.
  - Solution: buy a **call option on the cap** (a call-on-call):
    - If the firm wins → they exercise the call, pay the premium, and secure the rate cap.
    - If they lose → the compound option expires worthless; loss is limited to the small premium.
  - This strategy reduces upfront costs and hedges only if needed.
- Compound options provide flexible, lower-cost hedges when the need for an option depends on a future event.

Binomial Tree Diagram for the Compound Call Example

- $S = K = 100, \quad u = 1.2, \quad d = 0.8, \quad \Delta t = 0.5, \quad r = 5\%$
- Risk–neutral probability: $p = \frac{e^{r\Delta t} - d}{u - d} \approx \frac{1.0253 - 0.8}{0.4} \approx 0.5633, \quad 1 - p \approx 0.4367$
- Discount factor per step: $e^{-r\Delta t} \approx 0.9753$

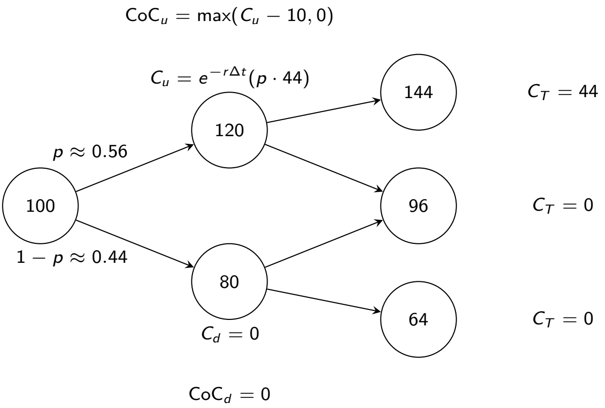

$$\text{CoC}_0 = e^{-r\Delta t}\big(p\,\text{CoC}_u + (1-p)\,\text{CoC}_d\big) \Rightarrow \ \text{CoC}_0 \approx 7.8$$

Gap Options

Path-Independent Exotic Options

- A **gap call** has two strike levels:
  - **Trigger price** K_2 — determines *whether* the option pays,
  - **Payout strike** K_1 — determines *how much* the option pays.
- **Payoff at maturity:**

$$\text{Gap Call Payoff} = \begin{cases} S_T - K_1, & \text{if } S_T > K_2, \\ 0, & \text{if } S_T \leq K_2. \end{cases}$$

- Even if the option finishes deeply in the money (e.g., S_T huge), it only activates when the trigger K_2 is crossed. When triggered, the payoff behaves like a standard call with strike K_1.
- **When are gap options used?** Gap options appear naturally in:
  - Employee stock compensation and performance targets
  - Structured products where "activation" and "payout" are separate

Gap Options

Path-Independent Exotic Options

- **Valuation formula for a gap call:**

$$C_{\text{gap}} = Se^{-qT}N(d_1) \; - \; K_1 e^{-rT}N(d_2),$$

where

$$d_1 = \frac{\ln(S/K_2) + (r - q + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T}.$$

- The gap option uses K_2 in the **activation test** (d_1, d_2), but the payoff uses K_1. This is why the formula looks like a "hybrid" of a digital and a vanilla call.

Illustration

- Pays $S - K_1$ when $S > K_2$. $K_1 = 90$ and $K_2 = 100$
- Does this option cost more or less without the gap with $K = K_1$?

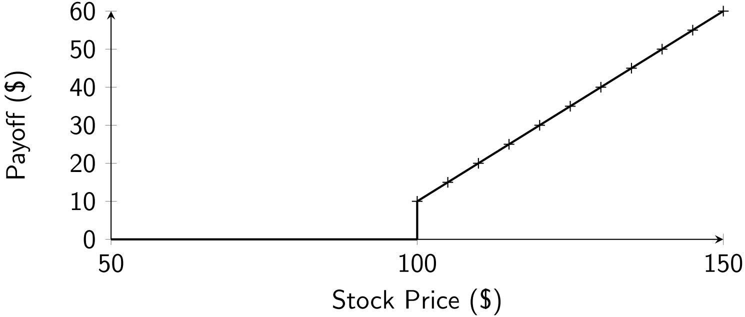

Barrier Options

Path-Dependent Exotic Options with Path-Independent Valuation

- **Barrier options** are options whose payoff depends on whether the underlying asset crosses a specified **barrier level** during the option's life.
  - The option may be *activated* or *extinguished* depending on the path.
  - Therefore: **payoff is path dependent** (not just based on S_T).
- **Knock-Out Options** — cease to exist if the barrier is breached.
  - **Down-and-out**: knocked out if price falls below the barrier.
  - **Up-and-out**: knocked out if price rises above the barrier.
- **Knock-In Options** — come into existence only if the barrier is touched.
  - **Down-and-in**: activated if price falls below the barrier.
  - **Up-and-in**: activated if price rises above the barrier.
- **Rebate Options** — pay a fixed cash amount if the barrier is hit.
  - **Down rebate**: pays if price drops below the barrier.
  - **Up rebate**: pays if price rises above the barrier.

Barrier Options

Path-Dependent Exotic Options with Path-Independent Valuation

- **Important pricing relationship: knock-in + knock-out = vanilla option.**
  - Example: Down-and-in call + Down-and-out call = Standard European call.
  - This allows closed-form barrier valuations using modified BSM formulas.
- **Which is worth more?**
  - A barrier option is *always worth less* than the otherwise identical vanilla option.
  - Why? The barrier introduces conditions that reduce the chance of payoff.

Down-and-In Barrier Option

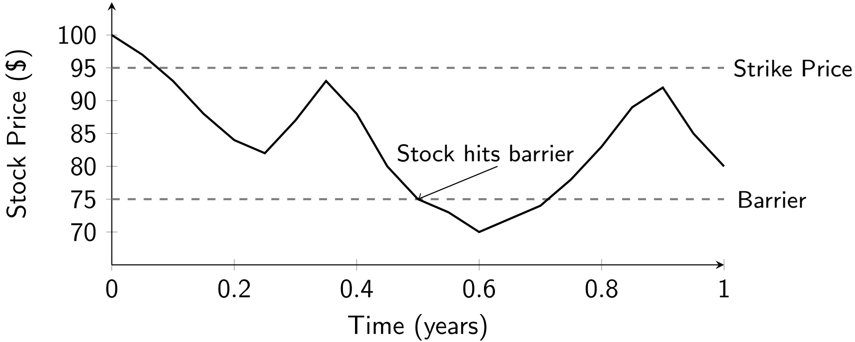

Barrier Options

Pricing and Parity Relationships

- **Barrier Parity Relations** For any barrier level, the knock-in and knock-out versions sum to the corresponding vanilla option:

$$c = c_{UI} + c_{UO}, \qquad c = c_{DI} + c_{DO},$$

$$p = p_{UI} + p_{UO}, \qquad p = p_{DI} + p_{DO}.$$

  These identities allow barrier option valuation by computing one component and backing out the other.

- **Pricing Methods**
  - Binomial/trinomial trees (must track whether barrier is hit).
  - Monte Carlo simulation (works well for knock-out; requires special tricks for knock-in).

- **Example: Currency Put Options (Standard vs. Barrier)** Parameters:

$$x_0 = 0.9, \quad \sigma = 0.10, \quad r_{\$} = 0.06, \quad r_e = 0.03, \quad T = 0.5$$

Table below shows put prices under different strikes and barrier levels.

Strike (K)	**Standard**	**Down-and-In Barrier**		**Up-and-Out Barrier**		
	(\$)	0.8000	0.8500	0.9500	1.0000	1.0500
$K = 0.8$	0.0007	0.0007	0.0007	0.0007	0.0007	0.0007
$K = 0.9$	0.0188	0.0066	0.0167	0.0174	0.0188	0.0188
$K = 1.0$	0.0870	0.0134	0.0501	0.0633	0.0847	0.0869

Lookback Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **(Strike) Floating Lookback Call**
  - Allows the holder to *buy at the minimum price* observed during the option's life. Payoff: $S_T - S_{\min}$.
  - Provides full protection against missing the lowest entry price.
- **(Strike) Floating Lookback Put**
  - Allows the holder to *sell at the maximum price* observed during the option's life. Payoff: $S_{\max} - S_T$.
  - Provides full protection against missing the best exit price.
- **Fixed Lookback Call**
  - Strike is fixed, but payoff depends on the highest price during the life of the option. Payoff: $\max(S_{\max} - K, 0)$.
  - Useful for capturing upside volatility without needing market timing.
- **Fixed Lookback Put**
  - Strike is fixed, but payoff depends on the lowest observed price. Payoff: $\max(K - S_{\min}, 0)$.
  - Provides strong protection against downside volatility.

Lookback Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **Comments**
  - Lookback options are usually **expensive** because the holder receives perfect hindsight on the price path.
  - Closed-form pricing formulas exist under:
    - continuous monitoring of the underlying price,
    - lognormal (GBM) price dynamics.
  - Related to *shout options*: the holder can "freeze" the intrinsic value once during the option's life.
    - A shout option lets the holder lock in the intrinsic value at one chosen time. Final payoff is the maximum of the locked-in value and the European payoff at maturity.

Non-Standard American Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **Bermudan Options** (intermediate between American and European)
  - The holder may exercise only on *specific, pre-specified dates* prior to maturity (not continuously).
  - Strike price or other contractual features may change over time.
  - Often valued using the same techniques as American options (e.g., binomial or trinomial trees), but with exercise allowed only at designated nodes.
  - Useful when issuers want flexibility but not the full early-exercise freedom of American options.
- **Typical Applications**
  - Many callable bonds (callable only on coupon dates).
  - Corporate warrants with staged exercisability or reset features.
  - Employee stock options where reset/ratchet mechanisms change the strike when options become deep out-of-the-money.

Non-Standard American Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **Illustrative Example**
  - A 7-year warrant may allow exercise only on specified dates during years 3–7.
  - The strike could adjust over time, e.g.:
    - Years 3–4: $K = \$30$
    - Years 5–6: $K = \$32$
    - Year 7: $K = \$33$
  - Reflects the firm's desire to control dilution timing and manage incentive alignment.

Asian Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **Definition**
  - Payoff is based on the *average* price of the underlying over a specified period.
  - Averages can be arithmetic or geometric; monitoring can be continuous or discrete.
  - The dependence on the full price path (not just S_T) makes the option path dependent.
- **When Asian Options Are Useful**
  - When economic exposure is naturally tied to an average price (e.g., exchange rates for importers/exporters, energy or commodity procurement, electricity load pricing).
  - When there is concern that the underlying price at a single moment may be distorted or manipulated.
  - When the underlying trades in thin or volatile markets where point-in-time prices are unreliable.
  - Some convertible bonds embed Asian features: conversion triggers or ratios are often based on the *average stock price over a window* (e.g., 20-day average near maturity).

Asian Options

Path-Dependent Exotic Options with Path-Dependent Valuation

- **Valuation Insight**
  - Asian call options are typically **less valuable** than comparable European calls.
  - Reason: averaging *reduces the effective volatility* of the payoff relative to using the terminal price S_T.
  - Lower volatility $\rightarrow$ lower option value.

Basic Types of Asian Options

- **Averaging Method: Arithmetic vs. Geometric**
  - Suppose the underlying price is sampled every h time units from $t = 0$ to T.
  - Arithmetic average: $A(T) = \frac{1}{N}\sum_{i=1}^{N} S_{ih}$
  - Geometric average: $G(T) = (S_h \times S_{2h} \times \cdots \times S_{Nh})^{1/N}$
  - Geometric average is always *less than or equal* to the arithmetic average $\Rightarrow$ geometric Asian options are cheaper.
- **Average Price vs. Average Strike**
  - Asians can average the *underlying price* or the *strike*.
  - Below: payoff formulas for European Asian options.

	Arithmetic	Geometric
Average price call	$\max[0, A(T) - K]$	$\max[0, G(T) - K]$
Average price put	$\max[0, K - A(T)]$	$\max[0, K - G(T)]$
Average strike call	$\max[0, S_T - A(T)]$	$\max[0, S_T - G(T)]$
Average strike put	$\max[0, A(T) - S_T]$	$\max[0, G(T) - S_T]$

Example

Hedging Currency Exposure with Asian Options

- XYZ receives monthly revenue of €100m but incurs costs in USD. Let x_i denote the spot dollar price of one euro in month i. In one year, the dollar value of the 12 revenue payments (with discounting at rate r) is:

$$\text{€}\,100\text{m} \times \sum_{i=1}^{12} x_i \, e^{r(12-i)/12}$$

- **Ignoring interest rates**, the total euro exposure relevant for hedging simplifies to:

$$\sum_{i=1}^{12} x_i = 12 \times \left(\frac{1}{12} \sum_{i=1}^{12} x_i \right)$$

  which is proportional to the **arithmetic average exchange rate** over the year.
- A natural hedge is an **arithmetic average price put option** that places a floor K on the average dollar-per-euro exchange rate:

$$\max\left(0,\ K - \frac{1}{12} \sum_{i=1}^{12} x_i \right)$$

- This protects XYZ against a year-long weakening of the euro, not just a single-day drop.

Example

Hedging Currency Exposure

- **Alternative hedging strategies**
  - A single long-dated European put (1-year maturity).
  - A basket of 12 monthly European puts (each matched to a cash flow).
  - Geometric- or arithmetic-average Asian puts.
  - A currency swap that exchanges fixed EUR cash flows for USD at a fixed rate.
- **Currency option valuation**: Use the Black–Scholes–Merton formula with a constant dividend yield, recognizing that a foreign currency earns the foreign risk-free rate.
- **Example parameters:** Spot $=$ \$0.90/EUR, strike $K = 0.90$, USD rate $r = 6\%$, EUR rate $r_{\text{euro}} = 3\%$, volatility $\sigma = 10\%$.

a. 12 European puts expiring in 1 year	0.2753
b. Basket of 12 monthly options	0.2178
c. 12 geometric-average puts	0.1796
d. 12 arithmetic-average puts	0.1764
e. Currency swap	?

a $>$ b because the long-dated option has more time value (positive Θ). b $>$ c,d because averaging reduces effective volatility. Swap value is typically zero at initiation—no option upside.

Pricing Asian Options

- **Closed-form solution for the geometric-average case**
  - When the underlying follows geometric Brownian motion, the *geometric average* of prices is itself lognormal.
  - Therefore, Asian options based on the geometric average can be priced using a Black-style formula, with an analytically derived mean and variance of the average.[1]
  - Provides a useful benchmark and lower bound on arithmetic-average Asian option prices.

- **Pricing the arithmetic-average case**
  - **Binomial or trinomial trees**: Possible but computationally intensive because the average introduces full path dependence (the number of states grows rapidly with time steps).
  - **Monte Carlo simulation**: The most common approach; handles path dependence naturally. Variance-reduction techniques (control variates using geometric Asians, antithetic sampling, etc.) are widely used.
  - **Approximation methods**: Moment matching, lognormal approximations, or PDE methods can also be applied.

[1] Under GBM, the geometric average has an exact lognormal distribution. The arithmetic average of lognormal variables, however, is *not* lognormal and has no closed form. See Hull, Chapter 26.13.

Exchange Options

Multivariate Options

- **Definition**
  - An exchange option pays off only if one asset outperforms another.

$$\max(0,\, S_T - N_T)$$

  - Useful when the goal is *relative* performance hedging (e.g., commodity spreads, equity relative value, FX crosses).

- **Value of a European Exchange Call (Margrabe Formula)**

$$Se^{-q_S T} N(d_1) \; - \; Ne^{-q_N T} N(d_2)$$

where

$$d_1 = \frac{\ln\left(\frac{S}{N}\right) + (q_N - q_S + \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}, \qquad d_2 = d_1 - \sigma\sqrt{T},$$

and the **effective volatility** of the spread is

$$\sigma = \sqrt{\sigma_S^2 + \sigma_N^2 - 2\rho\sigma_S\sigma_N}.$$

Exchange Options

Multivariate Options

- **Key intuition**
  - This is the Black–Scholes formula adapted to the ratio S/N.
  - Correlation ρ reduces uncertainty in the relative performance: strong positive correlation $\rightarrow$ lower effective volatility $\rightarrow$ cheaper exchange option.
- **Application: Implied Correlation**
  - Given market prices of exchange options, one can *back out* the implied correlation ρ between assets.
  - Widely used in multi-asset derivatives, equity basket products, commodities, and FX.

Quantos

Multivariate Options

- **Definition**
  - A *quanto* is a derivative that lets an investor hold an asset denominated in one currency, but receive payoffs in another currency **without currency risk**.
  - The payoff is converted at a *fixed* (or effectively hedged) exchange rate, regardless of where FX rates move.
- **Example:** Nikkei put warrants traded on the AMEX
  - Payoff is in USD, but is linked to the yen price of the Nikkei relative to a yen-denominated strike.
  - Investors gain exposure to Japan's equity market without USD/JPY exchange-rate uncertainty.

Example: Payoff of Quanto Call vs. Regular Foreign Call

Numerical Illustration

- Consider a call option on a foreign stock:
  - Strike (in foreign currency): $K = 100$
  - Payoff of foreign call (in foreign currency): $\max(S_T - K, 0)$
  - Spot FX today (domestic per foreign): $X_0 = 1.0$
- Two FX scenarios at maturity:
  - **Weak foreign currency**: $X_T = 0.8$
  - **Strong foreign currency**: $X_T = 1.2$
- Payoffs in **domestic currency**:

S_T	$\max(S_T - K, 0)$ (foreign)	Regular call (weak FX) $X_T = 0.8$	Regular call (strong FX) $X_T = 1.2$	Quanto call (fixed $X_0 = 1.0$)
80	0	0	0	0
100	0	0	0	0
120	20	16	24	20
140	40	32	48	40

Regular foreign call: payoff $= X_T \cdot \max(S_T - K, 0)$ depends on FX.

Quanto call: payoff $= X_0 \cdot \max(S_T - K, 0)$ is **independent** of FX at maturity.

Quantos

Multivariate Options

- **Why quantos are attractive**
  - Prevent exposure to both the foreign asset price and the FX rate. A US investor in Japanese stocks normally faces both Nikkei risk and USD/JPY risk.
  - A quanto payoff embeds a **currency forward** that locks the exchange rate at which the foreign payoff is converted.
  - However, the effective notional of the FX exposure changes with the level of the underlying asset.
  - Because of this **variable notional**, a simple FX forward cannot perfectly hedge the position.
  - Hence the name: *"quanto" = quantity-adjusting option*.

Pricing Exotic Options

Pricing Exotic Options

- **Multiple approaches to pricing exotic derivatives**
  - **Modified Black–Scholes–Merton formulas** Closed-form (or semi closed-form) solutions exist for certain structures (e.g., geometric Asian options, barrier options, exchange options). These typically require assumptions such as lognormality and continuous monitoring (See, e.g., Chapter 26 in Hull).
  - **Binomial or trinomial trees** Useful for early-exercise features (e.g., Bermudan, American-style exotics) or when state variables evolve discretely. Complexity increases with path dependence.
  - **Monte Carlo simulation** The most flexible approach, especially when the payoff depends on the full price path, multiple underlyings, or complicated triggers. Works well in high dimensions but requires variance-reduction techniques for accuracy.

Valuation Framework for Exotic Options

- **1. Path-Independent Exotics**
  - Payoff depends only on S_T (not on the path).
  - Typically valued with standard techniques such as BSM or binomial trees.
  - Some have closed-form solutions (e.g., binary/digital options, compound options, option packages).
- **2. Path-Dependent Exotics with Path-Independent Valuation**
  - Payoffs depend on the path, but valuation does not require tracking every path detail.
  - Examples: American options, barrier (knock-in/knock-out), lookback options. These can often be priced using modified BSM formulas, PDE methods, or recombining trees.

Valuation Framework for Exotic Options

- **3. Path-Dependent Exotics with Path-Dependent Valuation**
  - Payoff depends on the entire history of S_t, and valuation requires simulating or tracking that history.
  - Examples: Asian options (arithmetic), mortgage-backed securities (MBS), variance swaps.
  - Monte Carlo simulation or advanced tree methods (non-recombining lattices) are often required.

- **4. Multivariate Exotics**
  - Payoff depends on *multiple* underlying assets, indices, or rates.
  - Examples: Exchange options, quanto options, basket options, rainbow options.
  - Valuation usually requires Monte Carlo simulation or multivariate analytical formulas (when available).

Binomial (Risk-Neutral) Trees

- Recall the one–step binomial model under the risk-neutral measure.
- Assume: $S_0 = 100$, $K = 100$, $T = 1$, $r = 2\%$, $\sigma = 30\%$.
- Up factor under GBM:$u = e^{\sigma\sqrt{T}} = 1.34986$, $\qquad d = \frac{1}{u} = 0.74082$.
- Risk-neutral probability:$p = \frac{e^{rT} - d}{u - d} = 0.4597$.
- The value of any derivative with payoff $f(S_1)$ is:

$$V_0 = e^{-rT}\left[p\, f(S_{1,u}) + (1-p) f(S_{1,d})\right].$$

- For a call option:

$S_{1,d} = 74.082$

$f_{1,d} = 0$

$S_0 = 100$

$p = 0.4597$

$c_0 = 15.731$

$S_{1,u} = 134.986$

$f_{1,u} = 34.986$

Monte Carlo Simulation on Risk-Neutral Trees

- Instead of computing the expectation analytically, we can **simulate** up/down movements.

- In Excel, `RAND()` generates $U \sim \text{Uniform}(0, 1)$.

1. Generate N values of `RAND()`.
   - If `RAND()` $< p$, move to the up state.
   - If `RAND()` $> p$, move to the down state.

2. For each simulation j, obtain the simulated price:

$$S_1^{(j)} \in \{S_{1,u}, S_{1,d}\}.$$

3. Compute the payoff in each simulation:

$$V(S_1^{(j)}) = \max(S_1^{(j)} - K, 0).$$

4. The Monte Carlo estimator of the option value is:

$$\hat{V}_0 = \frac{1}{N} \sum_{j=1}^{N} e^{-rT} V(S_1^{(j)}).$$

Monte Carlo Simulation on Risk-Neutral Trees

- For example, using $p = 0.4587$, the following 10 simulations are obtained.

RAND()	Move	Price at T	Payoff	Discounted
0.457335	up	134.986	34.986	34.293
0.393937	up	134.986	34.986	34.293
0.090053	up	134.986	34.986	34.293
0.878148	down	74.082	0	0
0.658659	down	74.082	0	0
0.759579	down	74.082	0	0
0.798027	down	74.082	0	0
0.061689	up	134.986	34.986	34.293
0.969222	down	74.082	0	0
0.392675	up	134.986	34.986	34.293
			Average	17.147
			Std. error	5.715

- With only $N = 10$, the estimate $\hat{V}_0 = 17.147$ differs significantly from the true value $V_0 = 15.731$.

- Increasing N improves accuracy.

Monte Carlo Simulation on Risk-Neutral Trees

- **How large should N be?**
  - The precision of Monte Carlo estimates improves at rate $1/\sqrt{N}$.
  - Standard error:
    $$s.e. = \frac{\hat{\sigma}}{\sqrt{N}},$$
    where $\hat{\sigma}$ is the sample standard deviation of discounted payoffs.
- In the earlier example, $\hat{\sigma}/\sqrt{10} = 5.715$.
  - A 95% confidence interval is:
    $$\left[\hat{V}_0 - 2s.e.,\ \hat{V}_0 + 2s.e.\right] = [5.717,\ 28.577].$$
  - The interval is extremely wide because $N = 10$ is too small.
- Increasing to $N = 1000$:
  $$\hat{V}_0 = 15.725, \qquad s.e. = 0.52.$$
- Confidence interval becomes:
  $$[14.685,\ 16.765],$$
  which is much tighter.

Multi-step Trees

- As the nubmer of steps increases, the estimate becomes more precise.

Binomial Tree Model

	Stock Assumption			**Option Assumption**			**Tree**			Risk Neutral Prob	
	mu	0.1		T	1		n	10		q*	0.486836
	sigma	0.3		K	100		h	0.1			
	r	0.02		Call or Put	1	(=1 for call	u	1.099514		Price Binomial	3.787
	div yield	0					d	0.909493		Delta Binomial	0.166
	S0	100					p	0.52919			

time ==>	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
i==>	0	1	2	3	4	5	6	7	8	9	10
0	100.000	109.951	120.893	132.924	146.151	160.696	176.687	194.270	213.603	234.859	258.231
1		90.949	100.000	109.951	120.893	132.924	146.151	160.696	176.687	194.270	213.603
2			82.718	90.949	100.000	109.951	120.893	132.924	146.151	160.696	176.687
3				75.231	82.718	90.949	100.000	109.951	120.893	132.924	146.151
4					68.422	75.231	82.718	90.949	100.000	109.951	120.893
5						62.229	68.422	75.231	82.718	90.949	100.000
6							56.597	62.229	68.422	75.231	82.718
7								51.475	56.597	62.229	68.422
8									46.816	51.475	56.597
9										42.579	46.816
10											38.725

Option Pricing Tree

time ==>	0.000	0.100	0.200	0.300	0.400	0.500	0.600	0.700	0.800	0.900	1.000
i==>	0.0	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0
0	12.530	18.247	25.881	35.666	47.656	61.691	77.484	94.868	114.002	135.059	158.231
1		7.156	11.075	16.699	24.431	34.528	46.948	61.294	77.086	94.470	113.603
2			3.465	5.782	9.429	14.947	22.879	33.522	46.551	60.895	76.687
3				1.280	2.345	4.231	7.481	12.871	21.292	33.123	46.151
4					0.275	0.566	1.164	2.396	4.932	10.151	20.893
5						0.000	0.000	0.000	0.000	0.000	0.000
6							0.000	0.000	0.000	0.000	0.000
7								0.000	0.000	0.000	0.000
8									0.000	0.000	0.000
9										0.000	0.000
10											0.000

Multi-step Trees: Monte Carlo Simulation

- Generate stock prices: When $RAND() > p$, go down. Otherwise, go up
- 1,000 simulations of stock prices. Get S_T, compute $f_T e^{-rT}$, take the average.

Option prices by simulation												
	Simulated Call Price									Price Binomial	12.530	
	13.196									Price Black Scholes	12.822	
	0.685											
				Simulation of Risk Neutral Stock Prices								
time ==>	Discounted	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
i==>	Call Payoff	0	1	2	3	4	5	6	7	8	9	10
1	20.479	100	90.949	82.718	75.231	82.718	90.949	100.000	90.949	100.000	109.951	120.893
2	0.000	100	90.949	82.718	75.231	68.422	75.231	68.422	62.229	56.597	62.229	68.422
3	0.000	100	90.949	82.718	90.949	82.718	90.949	82.718	75.231	68.422	62.229	56.597
4	20.479	100	109.951	120.893	109.951	120.893	132.924	120.893	109.951	120.893	109.951	120.893
5	0.000	100	109.951	100.000	90.949	82.718	75.231	68.422	75.231	82.718	75.231	68.422
6	75.169	100	109.951	100.000	90.949	100.000	109.951	120.893	132.924	146.151	160.696	176.687
7	45.238	100	90.949	100.000	109.951	100.000	109.951	100.000	109.951	120.893	132.924	146.151
8	0.000	100	109.951	120.893	109.951	100.000	90.949	82.718	75.231	68.422	75.231	82.718
9	0.000	100	90.949	82.718	75.231	68.422	62.229	56.597	62.229	68.422	75.231	68.422
10	45.238	100	109.951	100.000	90.949	100.000	90.949	100.000	109.951	120.893	132.924	146.151

Why Monte Carlo Simulations?

- **Why do we need Monte Carlo simulation when we already have the binomial tree?**
  - Trees work well when the payoff depends only on the terminal price S_T.
  - Many derivatives have **path-dependent payoffs**: the payoff depends on the *entire sequence* of prices over time, not just the final one.
  - For such securities, the number of distinct paths grows exponentially, making the tree non-recombining and computationally infeasible as n becomes large.

- **Example: Asian call option**

$$\max\left(\frac{1}{T}\sum_{t=1}^{T} S_t - K,\, 0\right)$$

  - The average $\frac{1}{T}\sum S_t$ differs even for paths that end at the same S_T.
  - Thus, knowing the terminal node alone is not enough to price the option.

- Monte Carlo simulation handles path dependence naturally by simulating the entire price path, making it one of the most powerful tools for valuing complex exotics.

Why Monte Carlo Simulations?

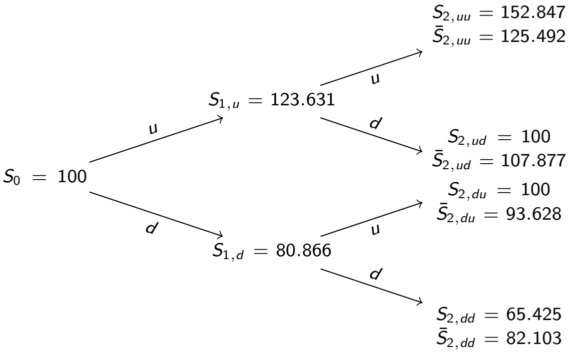

- Even though $S_{2,ud} = S_{2,du} = 100$, the **path averages differ**:$\bar{S}_{2,ud} \neq \bar{S}_{2,du}$. This makes the tree **non-recombining**.
- With n steps, a recombining tree has only $n + 1$ terminal nodes. A *non-recombining* tree has 2^n nodes—computationally explosive.
- Path-dependent options (Asian, lookback, etc.) are therefore much harder to price by tree methods.
- **Monte Carlo simulation** naturally avoids this explosion by simulating sample paths rather than enumerating the entire tree.

Monte Carlo Simulations Without Trees

- **Monte Carlo simulation is not limited to binomial trees.**
  - Trees were useful to introduce risk-neutral pricing.
  - But once no-arbitrage and risk-neutrality are established, we are free to simulate *any* valid risk-neutral process.
- **Key requirement:** Risk-neutral pricing must hold — meaning the asset price process must admit a dynamic replication argument (no-arbitrage).
- On a tree, these no-arbitrage conditions hold by construction.
- Once we accept risk-neutral pricing, we can simulate price paths from any assumed process:
  - Lognormal GBM (as in the Black–Scholes–Merton model),
  - Time-varying volatility models (Heston, GARCH),
  - Stochastic interest rates,
  - Jump–diffusion models.

Monte Carlo Simulations Under Lognormality

- Under the GBM assumption, stock prices evolve as:

$$dS_t = rS_t dt + \sigma S_t dW_t.$$

- Discretized for simulation:

$$S_{t+\Delta t} = S_t \exp\left((r - \tfrac{1}{2}\sigma^2)\Delta t + \sigma \epsilon_t \sqrt{\Delta t}\right),$$

  where $\epsilon_t \sim N(0,1)$.

- Alternatively:

$$\ln\left(\frac{S_{t+\Delta t}}{S_t}\right) \sim N\left((r - \tfrac{1}{2}\sigma^2)\Delta t,\ \sigma^2 \Delta t\right).$$

- This provides a simple algorithm for generating paths:
  1. Draw $\epsilon_t \sim N(0,1)$.
  2. Update $S_{t+\Delta t}$ using the formula above.
  3. Repeat over many time steps to generate a full price path.

Monte Carlo Simulations with Multiple Factors

- Consider an option that pays the maximum of the returns on two stocks (e.g., Google S_t and Apple N_t):

$$\max\left(\frac{S_T}{S_0},\ \frac{N_T}{N_0}\right).$$

- Under risk-neutral pricing, the simulated processes may be:

$$S_{t+\Delta t} = S_t \exp\left((r - \tfrac{1}{2}\sigma_S^2)\Delta t + \sigma_S \epsilon_{1,t}\sqrt{\Delta t}\right),$$
$$N_{t+\Delta t} = N_t \exp\left((r - \tfrac{1}{2}\sigma_N^2)\Delta t + \sigma_N \epsilon_{2,t}\sqrt{\Delta t}\right).$$

- The returns of Google and Apple are likely correlated. Generate correlated shocks via:

$$\epsilon_{2,t} = \rho\,\epsilon_{1,t} + \sqrt{1-\rho^2}\,v_t,$$

where $v_t \sim N(0,1)$ is independent.

Monte Carlo Simulations with Multiple Factors

- For each simulated path i, compute the discounted payoff:

$$V^{(i)} = e^{-rT} \max\left(\frac{S_T^{(i)}}{S_0},\ \frac{N_T^{(i)}}{N_0} \right).$$

- Estimate the option value by averaging across n simulations:

$$\hat{V}_0 = \frac{1}{n} \sum_{i=1}^{n} V^{(i)}.$$

- Example (with $\sigma_S = \sigma_N = 0.3$, $r = 0.02$, $\rho = 0.7$):

$$\hat{V}_0 = 1.134.$$

- Consider instead the “relative return” option:

$$\max\left(\frac{S_T}{S_0} - \frac{N_T}{N_0}, 0 \right).$$

- Simulation yields $\hat{V}_0 = 0.10$. This makes economic sense: if the stocks are positively correlated, their relative performance is less volatile.

Monte Carlo Simulation under the Heston Model

- **Heston dynamics under Q (risk-neutral):**

$$\begin{aligned} dS_t &= rS_t\,dt + \sqrt{v_t}\,S_t\,dW_t^{(1)}, \\ dv_t &= \kappa(\bar{v} - v_t)\,dt + \xi\sqrt{v_t}\,dW_t^{(2)}, \\ dW_t^{(1)}dW_t^{(2)} &= \rho\,dt. \end{aligned}$$

- **Discretization setup:**
  - Choose T, number of steps N, time step $\Delta t = T/N$.
  - Choose parameters $(S_0, v_0, r, \kappa, \bar{v}, \xi, \rho)$.
  - Choose number of paths M.

Monte Carlo Simulation under the Heston Model

- **Step-by-step simulation for each path $i = 1, \dots, M$:**
  1. Initialize $S_0^{(i)} = S_0$, $v_0^{(i)} = v_0$.
  2. For $n = 0, \dots, N-1$:
     - Draw $Z_1, Z_2 \sim N(0,1)$ independently.
     - Set correlated shocks:

$$\epsilon_s = Z_1, \quad \epsilon_v = \rho Z_1 + \sqrt{1-\rho^2}\, Z_2.$$

     - Update variance (Euler, full truncation):

$$v_{t_{n+1}}^{(i)} = \max\left(v_{t_n}^{(i)} + \kappa(\bar{v} - v_{t_n}^{(i)})\Delta t + \xi\sqrt{\max(v_{t_n}^{(i)}, 0)}\,\sqrt{\Delta t}\,\epsilon_v,\ 0\right).$$

     - Update price (log-Euler):

$$S_{t_{n+1}}^{(i)} = S_{t_n}^{(i)} \exp\left((r - \tfrac{1}{2}v_{t_n}^{(i)})\Delta t + \sqrt{v_{t_n}^{(i)}}\sqrt{\Delta t}\,\epsilon_s\right).$$

  3. At T compute payoff $f(S_T^{(i)})$ (e.g., $\max(S_T^{(i)} - K, 0)$).

- **Estimate option price:**

$$\hat{V}_0 = e^{-rT}\,\frac{1}{M}\sum_{i=1}^{M} f(S_T^{(i)}).$$

Summary

- **Main tools for pricing derivatives:**
  - Modified Black–Scholes–Merton formulas,
  - Binomial/trinomial trees,
  - Monte Carlo simulation.
- **Binomial trees:**
  - Essential for American-style options where optimal early exercise must be evaluated.
  - Useful when backward induction and ordered outcomes are important.
- **Monte Carlo simulation:**
  - Widely used in practice for complex or path-dependent derivatives.
  - Steps:
    1. Simulate many sample paths under the risk-neutral measure.
    2. Compute discounted payoffs per path.
    3. Estimate value as the average of payoffs.

Summary

- **Monte Carlo is especially useful for:**
  - Path-dependent payoffs:
    - Asian options, barrier options, lookback options,
    - Options on maxima/minima, rainbow and basket options.
  - Multi-asset derivatives where correlations matter.
- **Also useful for pricing under general dynamics:**
  - Stochastic volatility,
  - Stochastic interest rates,
  - Jump–diffusion processes.
- Increasing computational power has made Monte Carlo one of the most flexible and powerful pricing tools.