B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

Introduction

1.1 Aims of the course

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

Why do we want to study financial modelling? Hasn't an over-reliance on mathematics in finance led to significant social costs?

- Simply ignoring finance isn't going to work.
- Understanding the system, and why it operates the way it does, is the first step to effectively improving it
- Utimately, if we want to act in the real world, economics and finance are going to be part of what we need to do.

What are we trying to do in this course?

- Build financial models and understand how they can be used
- Understand where the models will fail, and where we need to take particular care
- Develop technical proficiency which will allow us to work with better models

There are a huge number of books on financial derivatives. Here is a selection, worth consulting for background reading.

- Steven E. Shreve, *Stochastic calculus for finance I: The binomial asset pricing model*, Springer 2004
  (A superb probabilistic account of the binomial model.)
- Steven E. Shreve, *Stochastic calculus for finance II: Continuous-time models*, Springer 2004
  (A superb first text on stochastic calculus for finance with many examples.)
- Alison Etheridge, *A course in financial calculus*, CUP 2002
  (An excellent primer on stochastic calculus for finance.)

- Paul Wilmott, Sam Howison and Jeff Dewynne, *The mathematics of financial derivatives: A student introduction*, CUP 1995
  (A decent first text on the PDE aspects of the subject.)
- Tomas Björk, *Arbitrage theory in continuous time*, 3rd Ed., OUP 2009
  (A good all-round text which covers many topics outside the scope of the course.)
- John C. Hull, *Options, futures and other derivatives*, 8th Ed., Pearson 2011
  (A bestseller that has a more financial as opposed to mathematical bias, and was one of the first textbooks on the subject, becoming a mainstay of many trading rooms.)

- Jean Jacod and Philip Protter, *Probability essentials*, Springer 2003
  (An excellent text on measure-theoretic probability, good for background.)
- Geoffrey R. Grimmett and David R. Stirzaker, *Probability and random processes*, 3rd Ed., OUP 2001
  (An excellent and encyclopedic background probability text.)

B8.3 Mathematical Models of Financial Derivatives

Introduction

1.2 Assumptions

Samuel Cohen

Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

What are the key concerns of financial modelling?

1. Avoid being exploited.
2. Control, manage and understand your risk.
3. Make a profit (usually by managing risk).

As we will see, the first concern is critical in practice.

We will need some economic assumptions when building financial models. The basic assumptions which underlie the models in the course are:

1. There is a riskless investment (a bank account or bonds) which grows at a constant, continuously compounded rate r. If M_t is invested at time t then it grows to $M_T = M_t\, e^{r(T-t)}$ at time $T > t$. A guaranteed amount of B_T paid at time T is worth $B_t = B_T\, e^{-r(T-t)}$ at time $t < T$. Borrowing and lending rates are both assumed equal to r.
2. There are no trading costs; if an asset can be bought for S_t at time t it can be sold for S_t at time t, and any amount can be bought or sold at the same price.

3. Assets are infinitely divisible, so it possible to own 0.432 shares for instance. This is not a major issue as forwards, calls and puts are usually written on $1,000$s or $10,000$s of shares, rather than one share.
4. Short-selling (i.e. holding negative quantities of an asset) is allowed. This is *often* true.
   In many markets, one can *borrow* assets (for a fixed time, for a fee, which we ignore) and sell them, which allows you to own a negative quantity of the asset. This is known as a *covered* short, and is usually seen as a normal part of a well-functioning market.
   In other markets, you can sell something without owning it, provided you deliver within a specified period (usually 2 days). This is known as a *naked* short, and is quite controversial, as it has been linked to negative effects on asset prices.

The key concept which will alow us to build arguments is arbitrage:

An arbitrage is an investment which costs nothing (or less) to set up at time t, $X_t \leq 0$, *but at a later time* $T > t$ *has:*

1. *zero probability of having a negative value,* $\mathbb{P}(X_T < 0) = 0$;
2. *strictly non-zero probability of having a strictly positive value,* $\mathbb{P}(X_T > 0) > 0$.

We assume that *no arbitrage opportunities exist.* (In practice they do, but when institutions exploit them supply and demand causes prices to realign in order to eliminate them.)

- No-arbitrage also is often seen in terms of the ‘law of one price’
- If I have two ways of constructing the same payoff within the market, and one of them is cheaper than the other, then I can construct a portfolio which
  - Buys the cheaper version of the claim, sells the more expensive
  - At expiry has no risk, but has made an initial profit.
- We will often use no-arbitrage in this format.

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD
Mathematical Institute

Financial Products

2.1 Forward contracts

Samuel Cohen
Hilary Term 2021

- Imagine you have a contract in which you will receive USD in one year, but have to pay costs in GBP at that time.
- At current exchange rates this deal is profitable, but you are concerned about it becoming unprofitable if there is a fall in USD relative to GBP.
- How can you manage your risk?
  For simplicity, let's assume (for now) that no interest is paid on USD.

Definition

A forward is an agreement entered into

- by two parties at time t
- in which the holder (who has the long position) promises to pay the agreed forward price $F_t > 0$ for an asset at some given maturity date $T > t$, and
- the writer (who has the short position) promises to deliver the asset at time T for the forward price F_t.
- Both parties are obliged to go through with the transaction regardless of the asset price, $S_T > 0$, at maturity.
- Under normal circumstances, neither party has to pay to enter the agreement at time t.

Consider first an agreement to sell the asset at time T (so selling our USD); this is known as the *short* position. In this case, the forward may be hedged by

- borrowing cash equal to S_t, the price of the asset at time t,
- buying the asset, holding it to maturity then delivering it in return for F_t.

The payoff for doing this is $F_t - e^{r(T-t)}S_t$, and has no risk or initial cost. As there must be no arbitrage, we know that

$$F_t - e^{r(T-t)}S_t \leq 0.$$

For the long position, consider

- short-selling the asset at t,
- putting the money in a risk-free account and then
- using the forward to buy back the asset for F_t and close out the short sale.

The payoff for doing this is $e^{r(T-t)}S_t - F_t$, and by no arbitrage

$$e^{r(T-t)}S_t - F_t \leq 0$$

As we have assumed no transaction costs, the forward prices on each side of the deal are the same. Therefore, if there is no arbitrage, then

$$F_t = S_t \, e^{r(T-t)}.$$

The payoff diagram for the forward for the long position is a plot of the value of the forward to the holder at maturity against the value of the asset at maturity which is $S_T - F_t$.

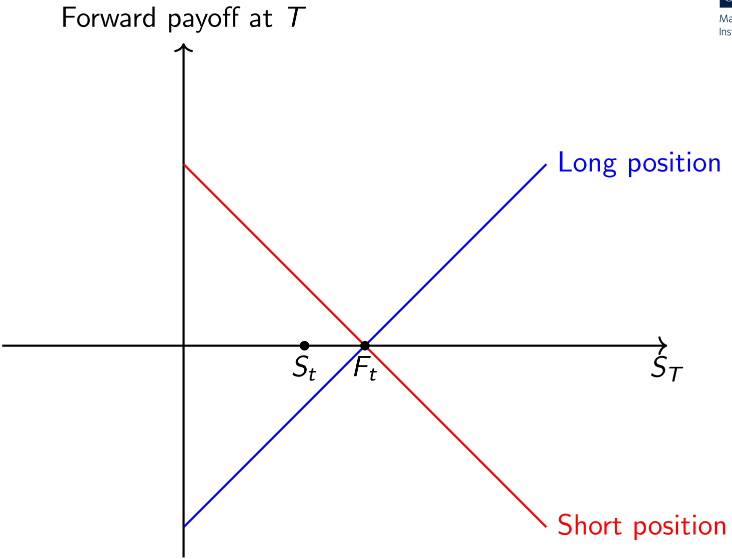
Forward payoff at T
Long position
S_t
F_t
S_T
Short position

Key points:

- The forward value is based on the *current* asset price and the *observed* interest rate. *It does not depend on whether the asset is being fairly priced.*
- We did not need to model the evolution of asset prices in the future.
- If interest rates are positive, and there is no cost/benefit to carrying for the asset (e.g. warehousing costs, foreign interest payments, etc), then the forward price is *above* the current ('spot') price of the asset. (Question: What happens in our argument above if USD pays interest at rate $\hat{r}$?)

- As we approach the expiry date, the forward and spot prices converge.
- Forwards cost nothing to enter, so provide easy exposure to risk.
- Unlike other many assets we will see, you can't purchase the same forward tomorrow that you purchase today (as the forward price changes).

Criticism

So, what are the possible flaws in this analysis (in addition to our earlier assumptions)?

- We have ignored any cost/benefit of holding the asset. This is fair enough for stocks or foreign currency (after accounting for dividends and interest), but is difficult for a lot of commoditites, where warehousing is expensive.
- A related issue is that our no-transaction-cost assumption is generally good if the contract is *cash-settled*. If settlement is in real assets, then you may face large transaction costs on the asset side. In practice there may also be a (small) transaction cost in the forward market, so the forward prices available on each side of the deal can be a little different.
- We have assumed that there is no default risk.

Mathematical Institute

The default risk issue is very significant, and has lead to the development of ‘Futures’ markets. These are very similar to forwards but:

- They have standardized terms (expiry dates, strikes)
- They are traded on an exchange, rather than over the counter
- They are cleared (so your contract is with a clearing house, rather than the person who bought the other side)
- This allows them to be bought and sold freely, as you don’t need to keep track of who holds the other side.
- A margin account is needed – this is an account of cash (or other liquid assets), held at the clearing house, which is used to offset changes in the value of your position.

- Forward contracts have existed in some form since antiquity – suggestions in the Code of Hammurabi (18th century BC) and in Aristotle's Politics (4th century BC).
- Formal markets for forwards developed during Tulipmania in Holland in the 1630s.
- Futures are more recent – the earliest example is the Dojima rice exchange in Osaka, Japan (1697).
- These became common for agricultural products in the late 19th century (e.g. Chicago Board of Trade 1864 – now part of CME Group), but financial futures (on currencies, interest rates, stock market indices, etc...) were only developed in the 1970s. In many of these markets, futures contracts are the main form of trading in the asset.

University of Oxford
Mathematical Institute

B8.3 Mathematical Models of Financial Derivatives

Financial Products

2.2 Call and put options

Samuel Cohen

Hilary Term 2021

- Suppose we are interested in the fixed-price aspect of a forward, but do not like the risk that we will be out-of-pocket if the asset falls.
- This leads us into the world of *options*. Options are common on equity (i.e. shares of companies).
- Let's assume again that a contract is written on a share which pays no dividends and doesn't cost anything to hold.

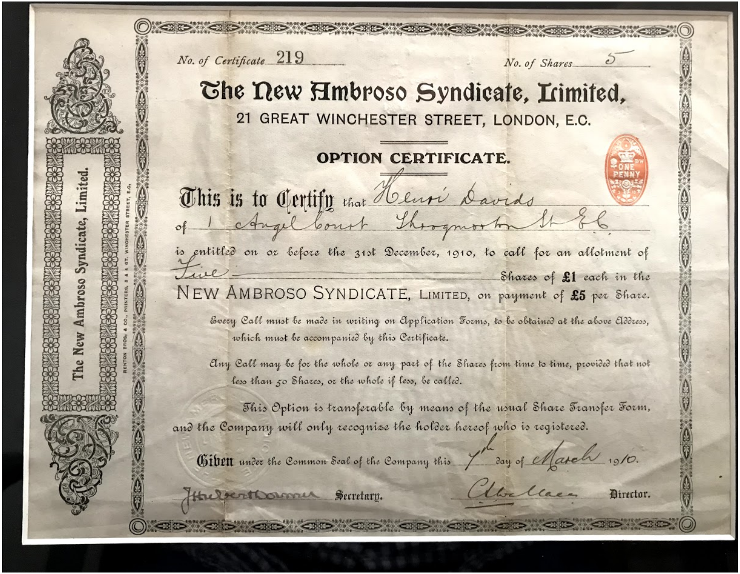

The New Ambroso Syndicate, Limited.

RENTON BROS. & CO., PRINTERS, 3 & 4 GT. WINCHESTER STREET, E.C.

No. of Certificate 219 *No. of Shares* 5

The New Ambroso Syndicate, Limited,

21 GREAT WINCHESTER STREET, LONDON, E.C.

OPTION CERTIFICATE.

ONE PENNY

This is to Certify that Henri Davids

of 1 Angel Court Throgmorton St E.C.

is entitled on or before the 31st December, 1910, to call for an allotment of Five Shares of **£1** each in the NEW AMBROSO SYNDICATE, LIMITED, on payment of **£5** per Share.

Every Call must be made in writing on Application Forms, to be obtained at the above Address, which must be accompanied by this Certificate.

Any Call may be for the whole or any part of the Shares from time to time, provided that not less than 50 Shares, or the whole if less, be called.

This Option is transferable by means of the usual Share Transfer Form, and the Company will only recognize the holder hereof who is registered.

Given under the Common Seal of the Company this 7th day of March 1910.

Secretary. Director.

Definition

A call option is a contract with an expiry date $T > t$ and a strike $K > 0$ in which:

- the holder (who has the long position) has the *right* to buy the underlying share for the strike at the expiry date;
- the writer (who has the short position) is *obliged* to deliver the share for the strike if the holder exercises their right.

The value (of the long position) of the call at expiry is

$$\max(S_T - K, 0) = (S_T - K)^+.$$

Unlike forwards, the holder has to pay a positive amount for the call option (this is a consequence of no arbitrage).

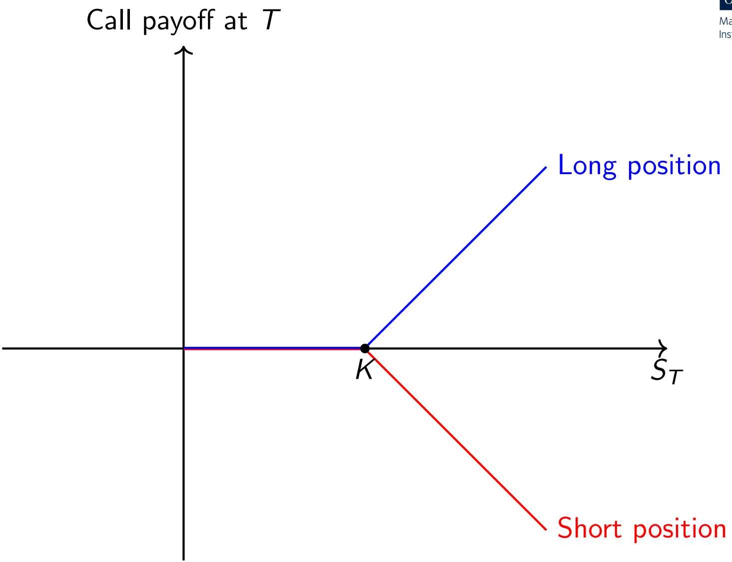
Call payoff at T
Long position
K
S_T
Short position

Some terminology

- The *Strike* or *expiry price* is K, which is agreed when *writing* the contract
- *Spot* is the the same as the price of the underlying asset S_t
- *Premium* is the price of the option. It consists of:
  - The *Intrinsic value*, the value of the option (long position) if it were to be exercised now
  - The *Time value*, the value of the option above intrinsic value.
- An option is *at the money* if the strike and the spot are the same (so no intrinsic value).
  - It is *in the money* if it has positive intrinsic value, *out of the money* if it has zero (or negative) intrinsic value.
  - The *moneyness* of the option is S_t/K.
- *Maturity* is the length of time until the *expiry date*
- For some options you have the right to *exercise* the option at various times.

Definition

A put option is a contract with an expiry date $T > t$ and a strike $K > 0$ in which

- the holder (who has the long position) has the *right* to sell the underlying share for the strike at the expiry date;
- the writer (who has the short position) is *obliged* to buy the share for the strike if the holder exercises their right.

The value (of the long position) of the call at expiry is

$$\max(K - S_T, 0) = (K - S_T)^+$$

The holder has to pay a positive amount for a put option (by no arbitrage).

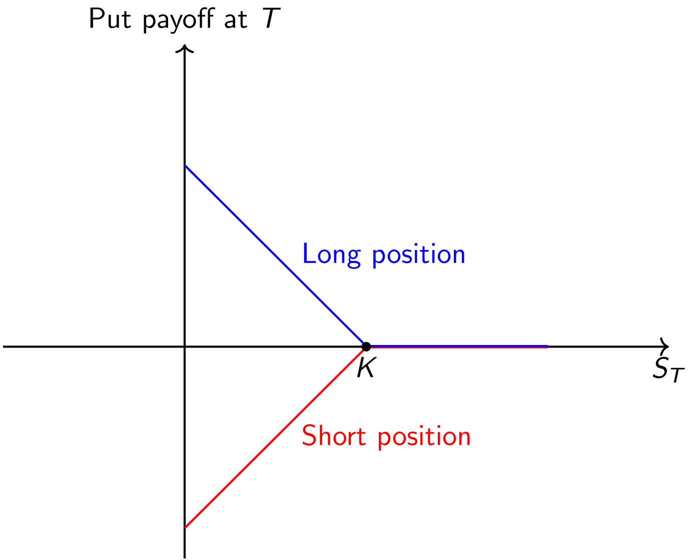
Put payoff at T
Long position
K
S_T
Short position

European vs American options

- An option which may be exercised *only* at its expiry date is called a European option.
- One which may be exercised at any time up to and including its expiry date is called an American option.
- Other styles exist – we will see some of these later in the course.
- In this course, options will be assumed European unless stated otherwise.

Unlike for forwards, we cannot give a price to an option without building a model for the evolution of stock prices. However, there is a relationship betwen put and call options which must hold.

Theorem (Put–Call parity)

Consider a European call and a European put option, written on the same underlying asset with (spot) price S_t. Suppose the risk-free discount factor is from t to T is $e^{-r(T-t)}$ and there is no default risk. Then their respective time-t prices C_t and P_t must satisfy

$$C_t - P_t = S_t - Ke^{-r(T-t)}$$

or their exists an arbitrage opportunity.

Proof

UNIVERSITY OF OXFORD
Mathematical Institute

Consider purchasing a single call and selling a single put, for initial cost $C_t - P_t$. This portfolio has no cashflows before expiry time T and, at expiry, has (unknown) payoff

$$C_T - P_T = (S_T - K)^+ - (K - S_T)^+ = S_T - K.$$

While its value is unknown, this payoff is identical, in every state of the world, to a portfolio consisting of one stock and a debt of K at time T. At time t, the latter portfolio has value

$$S_t - Ke^{-r(T-t)}.$$

If these two portfolios have different initial prices, then by purchasing the cheaper one and shortselling the other, an arbitrage profit can be made. □

B8.3 Mathematical Models of Financial Derivatives

The binomial model

3.1 A single step model

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

University of Oxford Mathematical Institute

The *simplest* model for a random share price is the one-step binomial model.

- The asset price is S_t at time t.
- At time T it can be either

$$\begin{cases} S_T = S^u & \text{with probability } p > 0, \\ S_T = S^d < S^u & \text{with probability } 1 - p > 0. \end{cases}$$

- No arbitrage implies that

$$S^d < S_t \, e^{r(T-t)} < S^u.$$

- We assume a risk-free investment is available, which grows from 1 at time t to $R = e^{r(T-t)}$ at time T.

An option with payoff function $f(S_T)$ at time T is written on this asset.

- For a call, we have $f(s) = (s - K)^+$.
- For a put, we have $f(s) = (K - s)^+$.
- For a forward, we have $f(s) = s - K$.

At expiry our option's value is

$$V_T = \begin{cases} V^u = f(S^u) & \text{with probability } p \\ V^d = f(S^d) & \text{with probability } 1 - p \end{cases}$$

The problem is to find the current value of the option V_t. There are at least three (related) ways to do this.

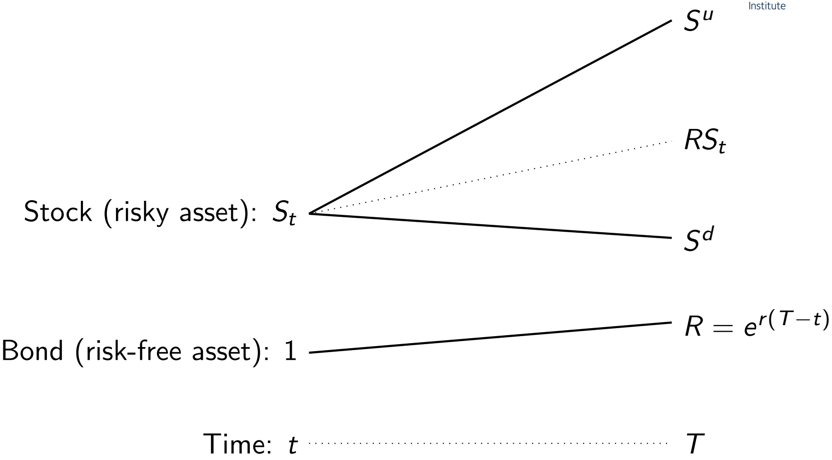

Figure: Underlying asset prices in a one-step binomial model

Approach 1: Delta hedging argument

Key idea:

- We set up a trading strategy in the underlying stock and the option which guarantees a fixed payoff (at time T).
- This strategy doesn't require any cash inputs between times t and T.
- The initial cost of setting up this portfolio must be the same as the value of the payoff (as there's no risk)
- We use this to solve for the initial option price.

- At time t, set up a portfolio long an option and short Δ_t shares. The initial value is

$$\Pi_t = V_t - \Delta_t S_t,$$

and hold this portfolio fixed until time T.

- Choose Δ_t so that the portfolio has the same value regardless of whether the up-state or the down-state occurs,

$$V^d - \Delta_t S^d = V^u - \Delta_t S^u.$$

- This gives

$$\Delta_t = \left(\frac{V^u - V^d}{S^u - S^d} \right).$$

- This portfolio is *risk-free* and so must grow at the *risk-free rate*, or there would be an arbitrage opportunity.
- This implies that

$$(V_t - \Delta_t S_t)\, e^{r(T-t)} = V^u - \Delta_t S^u = V^d - \Delta_t S^d$$

- When we solve for V_t we find that

$$V_t = e^{-r(T-t)} V^u - \left(\frac{V^u - V^d}{S^u - S^d}\right)\left(e^{-r(T-t)} S^u - S_t\right)$$

giving a closed-form value V_t.

Approach 2: Self-financing replication argument

University of Oxford
Mathematical Institute

Key idea:

- We set up a trading strategy in the underlying stock and bonds, which gives the same payoffs as the option.
- The strategy doesn't require any cash inputs between times t and T.
- The initial cost of setting up the strategy must be the same as the option price.

In this version of the pricing argument we see that *the price of the option is simply the cost of setting up a self-financing portfolio that perfectly covers the option writer's liability at expiry* T*.*

- At time t set up a portfolio with ϕ_t shares and ψ_t bonds (bonds grow at the risk-free rate).

$$\Phi_t = \phi_t \, S_t + \psi_t.$$

- Hold this portfolio fixed and choose ϕ_t and ψ_t so that the portfolio has value V^u in the up-state and V^d in the down-state

$$\Phi^u = \phi_t \, S^u + \psi_t \, e^{r(T-t)} = V^u,$$
$$\Phi^d = \phi_t \, S^d + \psi_t \, e^{r(T-t)} = V^d.$$

Solving for ϕ_t and ψ_t gives

$$\phi_t = \left(\frac{V^u - V^d}{S^u - S^d} \right), \quad \psi_t = \left(\frac{S^u V^d - S^d V^u}{S^u - S^d} \right) e^{-r(T-t)}.$$

- As this portfolio perfectly replicates the option payoff (and has no other cash flows), its value at t must equal

$$V_t = \Phi_t = \phi_t S_t + \psi_t,$$

  which simplifies to the same formula for V_t as earlier.
- Note that $\Phi \equiv V$, $\psi \equiv \Pi$ and $\phi \equiv \Delta$; either argument amounts to a simple rearrangement of the symbols in the other.

From these two approaches, we can draw the following conclusions:

- our model for the share price is *complete* in the sense that we can replicate *any* payoff (i.e., solve one equation for Δ_t in the delta-hedging argument or two equations in two unknowns in the replication argument).
- The number of stocks we hold is 'almost' the derivative of the payoff V with respect to the underlying S (except for discretization).
- We have not assumed that the stock is being 'fairly' priced, but have found the *only* price for the option which is consistent with the stock price and does not lead to arbitrage.

Approach 3: Arrow-Debreu securities

We'll now consider a rather different way to obtain the price of an option in a binomial model.

- This involves creating a fictional probability measure (the 'risk-neutral measure') such that the price of all traded assets (in particular the option) is the discounted expected value.
- It's easiest to motivate this by introducing some 'basis securities' (in the same way as we have basis vectors in linear algebra).

Definition

An Arrow–Debreu security is a contract which, at an expiry time T, pays \$1 in one particular state of the world, and nothing otherwise.

- Usefully, this means that, in our Binomial model, we can define two Arrow–Debreu securities (A^u and A^d), and the price of the option at expiry is $V_T = V^u A^u_T + V^d A^d_T$.
- By no arbitrage, as V^u and V^d are fixed constants, and there are no intermediate cashflows, the price of V at time t must be

$$V_t = V^u A^u_t + V^d A^d_t.$$

- We then need only to find the price of A^u and A^d.

- Using our earlier formula (with $V^u = 1, V^d = 0$),

$$A_t^u = e^{-r(T-t)} \frac{e^{-r(T-t)}S^u - S_t}{S^u - S^d} = e^{-r(T-t)}q$$

where $q = \frac{e^{r(T-t)}S_t - S^d}{S^u - S^d} \in (0,1)$.

- Similarly $A_t^d = e^{-r(T-t)}(1-q)$. This can also be seen directly by no-arbitrage, as $A_T^u + A_T^d = 1$.
- Our earlier formula for a general option then simplifies to

$$V_t = V^u A_t^u + V^d A_t^d = e^{-r(T-t)}\big(qV^u + (1-q)V^d\big).$$

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD

Mathematical Institute

The binomial model

3.2 No arbitrage and the Fundamental Theorem of Asset Pricing

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- In the last lecture, we introduced the Arrow–Debreu securities, and saw how to expand the price of a general option in terms of the Arrow–Debreu prices
- By no arbitrage, we see that $A_t^u = e^{-r(T-t)}q$ where $q = \frac{e^{r(T-t)}S_t - S^d}{S^u - S^d} \in (0,1)$.
- By interpreting q as a ‘probability’, we get an interesting mathematical fiction.

We give the time-dependent version...

Definition

In the binomial model, the measure $\mathbb{Q}$ constructed by

$$\mathbb{Q}\Big(S_t \uparrow S^u | \mathcal{F}_t\Big) = \frac{e^{r(T-t)}S_t - S^d}{S^u - S_d},$$

where $\mathcal{F}_t$ is the information observed by time t, is called the *risk-neutral measure*. It has the property that any traded asset has price

$$V_t = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[V_T|\mathcal{F}_t].$$

We have seen that this is a well defined probability, assuming no-arbitrage.

- There is a close connection between this result and martingale theory.
- If we consider the discounted price $e^{-rt}V_t$, we see that

$$e^{-rt}V_t = \mathbb{E}_{\mathbb{Q}}[e^{-rT}V_T|\mathcal{F}_t],$$

  that is, discounted prices are martingales under the risk neutral measure.
- This is despite the fact that these assets are not risk-free!
- We do not expect $\mathbb{Q} = \mathbb{P}$ generally, as this would imply that prices don't reflect risk-aversion.

- The existence of this measure $\mathbb{Q}$ is, in fact, *equivalent* to no-arbitrage.
- It's uniqueness is equivalent to prices being determined uniquely.
- These two statements are known as the 'Fundamental Theorem of Asset Pricing', and hold in much more general models than the binomial model (but the proof becomes delicate in continuous time...)

Theorem (Fundamental Theorem of Asset Pricing)

1. *Assuming no arbitrage or transaction costs, and deterministic interest rates, there exists a probability measure* $\mathbb{Q}$ *such that the price of a payoff* X_T *at time* t *is given by* $e^{-r(t-T)}\mathbb{E}_{\mathbb{Q}}[X_T|\mathcal{F}_t]$. *(Also,* $\mathbb{Q}$ *is equivalent to the real-world probability measure in the sense of measures.)*
2. *The probability measure* $\mathbb{Q}$ *is unique if and only if all payoffs are traded (or can be replicated from traded claims).*

Proof

UNIVERSITY OF OXFORD
Mathematical Institute

See Etheridge (2002), S1.5 and S1.6 for a proof in a general discrete time and price model. Here we give a sketch of the proof over a single step.

- Assume no arbitrage or transaction costs.
- As there are no transaction costs, the prices of all avalable assets must be *linear*, that is, $\Pi(aX + Y) = a\Pi(X) + \Pi(Y)$ for any payoffs X, Y and any constant $a \in \mathbb{R}$.
- If we assume there are finitely many possible outcomes, then a payoff X can be represented by a vector $(x_1, ..., x_N)$ in $\mathbb{R}^N$ for some N (the number of outcomes). This is the Arrow–Debreu representation.

Proof

- Consequently, we can think of Π as a linear operator mapping $\mathbb{R}^N \to \mathbb{R}$. From algebra, we know that such an operator can always be written as a matrix, in particular,

$$\Pi(X) = \sum_i \pi_i x_i.$$

- Considering the case when $x_i \equiv 1$, so the payoff is constant, by no arbitrage we have $\Pi(1) = e^{-rt}$, which implies $\sum_i \pi_i = e^{-rt}$.
- Considering an Arrow–Debreu security, we see $\pi_i = e^{-rt} q_i$, and no arbitrage guarantees $q_i \geq 0$ and $q_i > 0$ if this outcome happens with nonzero probability.
- This guarantees $(q_1, ..., q_N)$ is a probability vector, as desired.

Proof

UNIVERSITY OF OXFORD
Mathematical Institute

- Uniqueness of $\mathbb{Q}$ is clearly the same as uniqueness of the prices of Arrow–Debreu securities.
- If we can replicate the Arrow–Debreu securities using traded assets, then we can replicate all possible claims, and their prices are unique.
- Otherwise, there are Arrow–Debreu securities where the price is only determined up to some bounds (by no-arbitrage), and the resulting $\mathbb{Q}$ is not unique.
- Hence $\mathbb{Q}$ is unique iff all claims can be replicated.

$\square$

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD

Mathematical Institute

The binomial model

3.3 Multiple time-steps

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

- We split the interval $[t, T]$ into n steps of length $\delta t = (T - t)/n$, say

$$t_0 = t, \quad t_{m+1} = t_m + \delta t, \quad t_n = T, \quad \text{for} \quad m = 1, 2, \dots n,$$

and build a binomial, or sometimes a binary, tree starting from S_t.

- It common practice to set

$$S^{\omega u}_{t_{m+1}} = u\, S^{\omega}_{t_m}, \quad S^{\omega d}_{t_{m+1}} = d\, S^{\omega}_{t_m},$$

where $u > 1$ and $0 < d < 1$ are constants and, frequently, $u \times d = 1$.

- Here ω denotes the path to the current node on the tree, for example after two steps $\omega \in \{uu, ud, du, dd\}$.

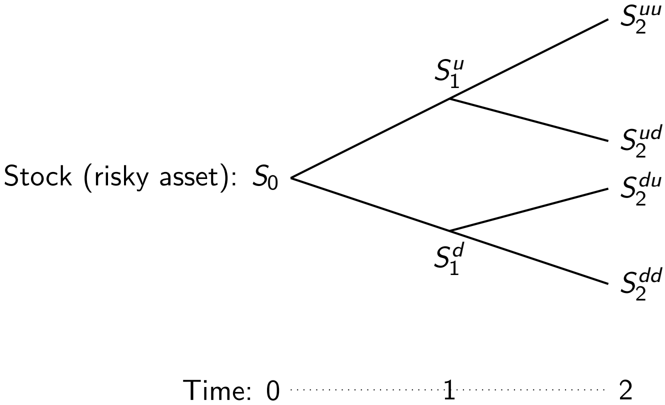

Figure: Underlying asset prices in a two-step binomial model

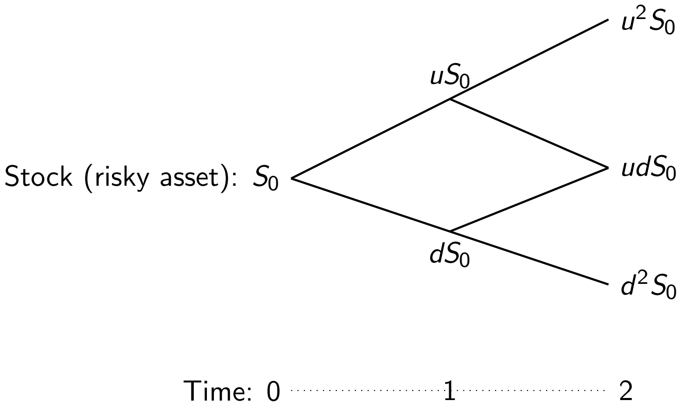

Figure: Underlying asset prices in a recombining two-step binomial model

No-arbitrage in the share price tree requires

$$0 < \left(\frac{S^{\omega}_{t_m} e^{r\delta t} - S^{\omega d}_{t_{m+1}}}{S^{\omega u}_{t_{m+1}} - S^{\omega d}_{t_{m+1}}} \right) = \left(\frac{e^{r\delta t} - d}{u - d} \right) < 1.$$

Over each step the risk-neutral pricing formula gives

$$V^{\omega}_{t_m} = e^{-r\delta t} \left(q\, V^{\omega u}_{t_{m+1}} + (1-q)\, V^{\omega d}_{t_{m+1}} \right), \quad q = \left(\frac{e^{r\delta t} - d}{u - d} \right), \qquad (1)$$

which requires us to work backwards from $t_n = T$, where we know the option prices from its payoff. This is sometimes called *dynamic programming.*

The Δ-hedging parameter at each step becomes

$$\Delta_{t_m}^{\omega} = \left(\frac{V_{t_{m+1}}^{\omega u} - V_{t_{m+1}}^{\omega d}}{S_{t_{m+1}}^{\omega u} - S_{t_{m+1}}^{\omega d}} \right)$$

and the replicating portfolio (at each step) is

$$\phi_{t_m}^{\omega} = \left(\frac{V_{t_{m+1}}^{\omega u} - V_{t_{m+1}}^{\omega d}}{S_{t_{m+1}}^{\omega u} - S_{t_{m+1}}^{\omega d}} \right), \quad \psi_{t_m}^{\omega} = \left(\frac{S_{t_{m+1}}^{\omega u} \, V_{t_{m+1}}^{\omega d} - S_{t_{m+1}}^{\omega d} \, V_{t_{m+1}}^{\omega u}}{S_{t_{m+1}}^{\omega u} - S_{t_{m+1}}^{\omega d}} \right) e^{-r\delta t}.$$

Recall that at time t_m and in state ω, $\phi_{t_m}^{\omega}$ is the number of shares we hold and $\psi_{t_m}^{\omega}$ is the amount of cash hold in order that we perfectly replicate the option's value in the two possible future states.

Let S_t be the value of a share and B_t be the value of a bond (i.e., cash) at time t. If at time t a portfolio has ϕ_t shares and ψ_t in cash then the value of the portfolio is

$$\Phi_t = \phi_t \, S_t + \psi_t \, B_t.$$

Let

$$\delta S_t = S_{t+\delta t} - S_t, \quad \delta B_t = B_{t+\delta t} - B_t, \quad \delta \Phi_t = \Phi_{t+\delta t} - \Phi_t$$

so, in general,

$$\delta \Phi_t = \phi_t \, \delta S_t + \psi_t \, \delta B_t \\ + (S_t + \delta S_t) \, \delta \phi_t + (B_t + \delta B_t) \, \delta \psi_t$$

If it turns out that

$$(S_t + \delta S_t)\,\delta\phi_t + (B_t + \delta B_t)\,\delta\psi_t = 0,$$

then any money to buy $\delta\phi_t$ new shares at $t + \delta$ comes from selling $\delta\psi_t$ bonds (i.e., borrowing the same amount of cash) and vice versa. If this is the case, we call the portfolio *self-financing* over $[t, t + \delta t)$ and we find that

$$\delta\Phi_t = \phi_t \delta S_t + \psi_t \delta B_t, \tag{2}$$

which is usually known as the *self-financing equation*.

- The ‘self-financing’ condition on strategies is implicit in the definition of an arbitrage
- This requirement is enough to ensure we have sensible pricing *through time*
- We can weaken it to the requirement that no assets are added to the portfolio (removing assets doesn’t change the definition substantially)

- The replication strategy given above is self-financing; over any interval $[t_m, t_{m+1})$ both $\phi^{\omega}_{t_m}$ and $\psi^{\omega}_{t_m}$ are fixed, so both $\delta\phi^{\omega}_{t_m} = 0$ and $\delta\psi^{\omega}_{t_m} = 0$.
- By construction, the replicating portfolio set up at t_m in state ω is guaranteed at time t_{m+1} to have the value of $V^{\omega u}_{t_{m+1}}$ in the up-state (ωu) and $V^{\omega d}_{t_{m+1}}$ in the down-state (ωd).
- So, although the number of shares and the amount of cash changes from $(\phi^{\omega}_{t_m}, \psi^{\omega}_{t_m})$ to $(\phi^{\omega\, u/d}_{t_m}, \psi^{\omega\, u/d}_{t_m})$ as we go from t^-_{m+1} to t^+_{m+1}, the *value* of the replicating portfolio does not; as we re-adjust the portfolio at t_{m+1}, we sell however many shares are necessary to buy the required number of bonds and vice versa.
- This establishes that under all possible circumstances in the binomial model, the (ϕ, ψ) strategy both replicates the option's payoff and is self-financing.

B8.3 Mathematical Models of Financial Derivatives

The binomial model

3.4 Computation and American Options

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

Mathematical Institute

Given we've got a way of pricing options, we can now think about how to implement it to value options over multiple steps.
Two main methods are possible:

- Using the tree structure for recursive computation
- Using the risk-neutral measure for computation

We'll cover the second approach first.

Risk-neutral measure computation

- In our basic binomial model, suppose that the risk neutral probability $q = \frac{e^{r\delta t} - d}{u - d}$ is constant.
- We then know that the $\mathbb{Q}$-probability of S_0 moving to $u^m d^{n-m} S_0$ by time T (with $\delta t = (T - t)/n$) is given by the binomial probability

$$\binom{n}{m} q^m (1-q)^{n-m}$$

- We can therefore calculate the price of a European option with payoff $f(S_T)$ as the discounted expected value

$$\begin{aligned} V_0 &= e^{-r(T-t)} \mathbb{E}_{\mathbb{Q}}[f(S_T)] \\ &= e^{-r(T-t)} \sum_m \binom{n}{m} q^m (1-q)^{n-m} f(u^m d^{n-m} S_0). \end{aligned}$$

- For $f(S) = (S-K)^+$, i.e. a call option, we can simplify further.
- Let $M = \left\lceil \frac{-n\log(d)+\log(K/S_0)}{\log(u)-\log(d)} \right\rceil$ (where $\lceil x \rceil$ is the smallest integer $\geq x$).
- Then,

$$\begin{aligned}
V_0 &= e^{-r(T-t)} \sum_m \binom{n}{m} q^m (1-q)^{n-m} (u^m d^{n-m} S_0 - K)^+ \\
&= e^{-r(T-t)} \sum_{m=M}^{n} \binom{n}{m} q^m (1-q)^{n-m} (u^m d^{n-m} S_0 - K) \\
&= \left(\sum_{m=M}^{n} \binom{n}{m} (uqe^{-r\delta t})^m (d(1-q)e^{-r\delta t})^{n-m} \right) S_0 \\
&\quad - \left(\sum_{m=M}^{n} \binom{n}{m} q^m (1-q)^{n-m} \right) e^{-r(T-t)} K
\end{aligned}$$

- Now notice that $\tilde{q} := uqe^{-r\delta t} = 1 - d(1-q)e^{-r\delta t}$.
- This means that, if we have $\Psi(x; n, q)$ as the upper tail of the Binomial(n, q) distribution (i.e. $\Psi(x; n, q) = \mathbb{P}[X \geq x]$ for $X \sim B(n, q)$), we can simplify

$$
\begin{aligned}
V_0 &= \Big(\sum_{m=M}^{n} \binom{n}{m} \tilde{q}^m (1-\tilde{q})^{n-m} \Big) S_0 \\
&\qquad - \Big(\sum_{m=M}^{n} \binom{n}{m} q^m (1-q)^{n-m} \Big) e^{-r(T-t)} K \\
&= \Psi(M; n, \tilde{q}) S_0 - \Psi(M; n, q) e^{-r(T-t)} K.
\end{aligned}
$$

- This is the discrete version of the Black–Scholes formula.
- The price of a put option follows by Put–Call parity.

Tree-based computation

- The earlier computation works for all European options.
- When we have American options, or other path-dependent problems (which we will see later in the course), we need alternative tools.
- The most basic tool is to work with the tree-structure given by our binomial model.
- To move one-step down the tree, we use the formula we've already derived:

$$V_t^{\omega} = e^{-r\delta t}(qV_{t+1}^{\omega u} + (1-q)V_{t+1}^{\omega d})$$

- To demonstrate, let's start with a simple European Put option, and assume $T = n = 2$, $e^{-r\delta t} = 0.95$, $u = 1/d = 1.2$, $S_0 = 1$, $K = 1.2$. Hence $q = 0.5981$ (all results to 4dp)

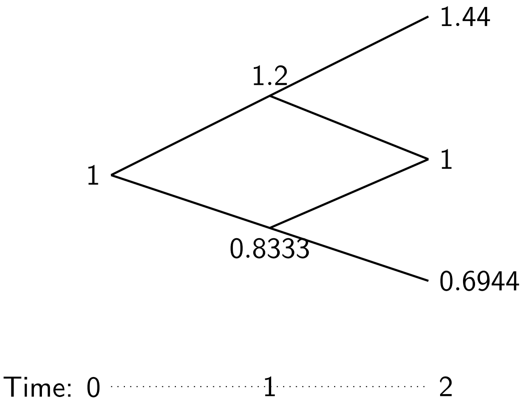

Figure: The underlying stock price

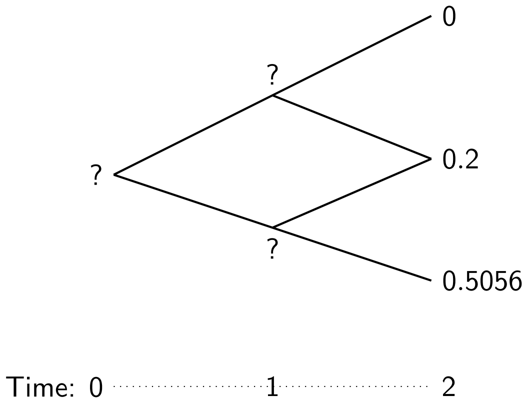

Figure: The Option value (terminal value $(1.2 - S_T)^+$)

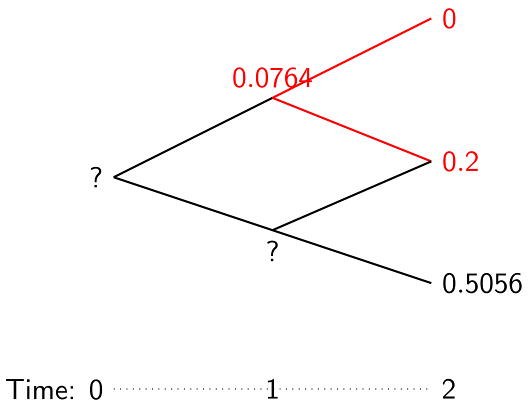

Figure: The Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV^{\omega u}_{t+1} + (1-q)V^{\omega d}_{t+1}) = 0.0764.$$

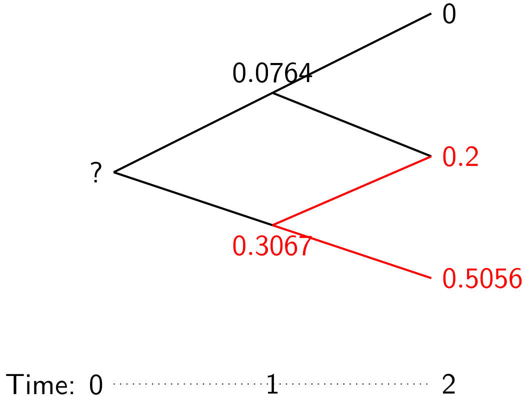

Figure: The Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV^{\omega u}_{t+1} + (1-q)V^{\omega d}_{t+1}) = 0.3067.$$

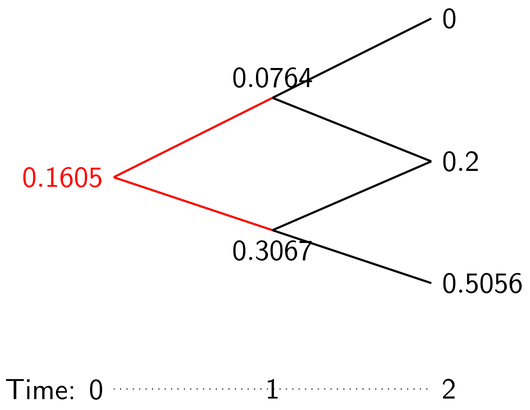

Figure: The Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV_{t+1}^{\omega u} + (1-q)V_{t+1}^{\omega d}) = 0.1605$$

American Options

- American options can be exercised at any time before expiry.
- Consequently, we need to include this decision in the calculation of the price.
- It is the *holder* of the option who has the right to choose whether to exercise, and we assume will do so to maximize the value.

At each node on the tree the option holder has two choices:

- hold the option until the next step, in which case its values is given by the recursion we've already seen; or
- exercise the option at this step and receive the payoff.

A rational investor will choose the one which makes the option most valuable to them and so if $P^{\omega}_{t_m}$ represents the payoff at the current node then

$$V^{\omega}_{t_m} = \max\Big(e^{-r\delta t}\left(q\,V^{\omega u}_{t_{m+1}} + (1-q)\,V^{\omega d}_{t_{m+1}}\right),\ P^{\omega}_{t_m}\Big)$$

Let's consider this for the same values as the option considered above.

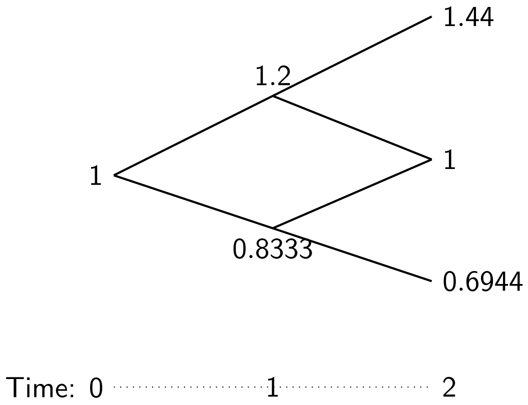

Figure: The underlying stock price

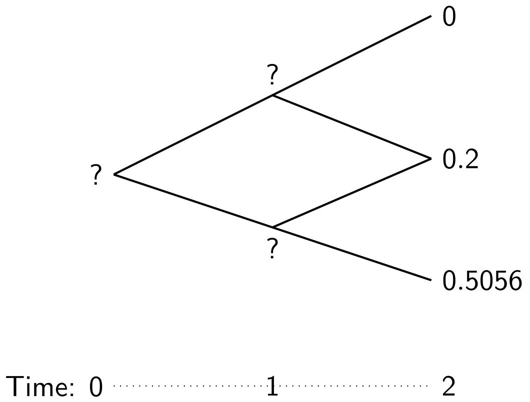

Figure: The American Option value (terminal value $(1.2 - S_T)^+$)

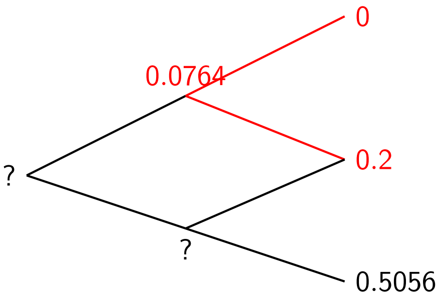

Figure: The American Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV_{t+1}^{\omega u} + (1-q)V_{t+1}^{\omega d}) = 0.0764 \text{ (no exercise)}$$
$$(K - S_t^{\omega})^+ = (1.2 - 1.2)^+ = 0$$

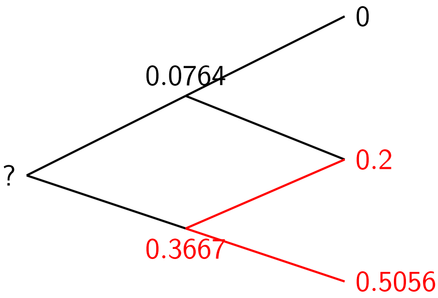

Figure: The American Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV^{\omega u}_{t+1} + (1-q)V^{\omega d}_{t+1}) = 0.3067$$
$$(K - S^{\omega}_t)^+ = (1.2 - 0.8333)^+ = 0.3667 \text{ (exercise!)}$$

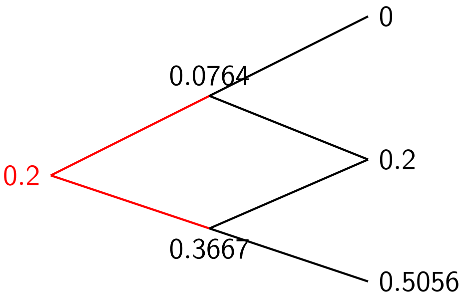

Figure: The American Option value (terminal value $(1.2 - S_T)^+$)

$$e^{-r\delta t}(qV_{t+1}^{\omega u} + (1-q)V_{t+1}^{\omega d}) = 0.1834$$
$$(K - S_t^{\omega})^+ = (1.2 - 1)^+ = 0.2000 \text{ (exercise!)}$$

B8.3 Mathematical Models of Financial Derivatives

Brownian Motion and Martingales

4.1 Definitions

Samuel Cohen
Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

- While the binomial model is easy to work with in some ways, we will gain a lot from moving to a continuous time and space setup, where we can use calculus.
- To do this, we need to build the basic tools for modelling, in particular we need some familiarity with the theory of stochastic processes.
- Arguably the most fundamental stochastic process is *Brownian motion*, so we spend a bit of time getting familiar with it.
- Our aim is to build a working theory, rather than to be completely rigorous with all our calculations (see B8.2 for this!)

Definition

A *stochastic process* is a sequence of random variables indexed by a parameter, for example, $(W_t)_{t\geq 0}$. For each fixed $t \geq 0$, W_t is a random variable.

Unless specified otherwise, in this course we assume that

- time ranges through $[0, T]$, where $T > 0$ is some expiry time (after which we don't need a mathematical model).
- (in-)equalities between random variables and stochastic processes are taken to hold *with probability 1* (also known as *almost surely*).

It will be useful to have some of the basic terminology of more general stochastic processes

- We assume that there is a space of possible outcomes Ω (for example, the space of all possible paths of the stock price).
- To avoid axiom-of-choice nonsense, with this we have a collection $\mathcal{F}$ of subsets of Ω, which we call *events*.
  - Formally, $\mathcal{F}$ is closed under taking complements, countable unions, and contains Ω, so is a σ-algebra.
  - Intuitively, $\mathcal{F}$ is all the events where we will eventually know whether they happen or not.
- A probability measure $\mathbb{P}$ gives probabilities $\mathbb{P}(A)$ for each $A \in \mathcal{F}$.
  - Formally $\mathbb{P}$ should be countably additive over disjoint sets and $\mathbb{P}(\Omega) = 1$.

- To model the flow of information, we say that we have a collection $\{\mathcal{F}_t\}_{t\geq 0}$ of subsets of $\mathcal{F}$ (in particular, sub-σ-algebras), which describe what we know at time t.
- Intuitively, the events $A \in \mathcal{F}_t$ are those for which, at time t, we know whether A will happen or not.
- As we can't forget information, we have the property that $\mathcal{F}_s \subseteq \mathcal{F}_t$ for $s \leq t$.
- With this setup, we can define the *conditional expectation* given $\mathcal{F}_t$, written $\mathbb{E}[\cdot|\mathcal{F}_t]$.
  - $\mathbb{E}[X|\mathcal{F}_t]$ is the random variable Y which minimizes $\mathbb{E}[(X-Y)^2]$ among random variables known at time t (formally, among $\mathcal{F}_t$-measurable random variables).
  - To avoid having to assume $\mathbb{E}[X^2] < \infty$, there is an alternative construction, see B8.1/2.

The conditional expectation has various useful properties:

- It is linear and monotone
  - $\mathbb{E}[aX + Y|\mathcal{F}_t] = a\mathbb{E}[X|\mathcal{F}_t] + \mathbb{E}[Y|\mathcal{F}_t]$ for all X, Y and all $a \in \mathbb{R}$.
  - If $X \geq Y$ in every state of the world then $\mathbb{E}[X|\mathcal{F}_t] \geq \mathbb{E}[Y|\mathcal{F}_t]$.
- It satisfies the *Tower law*: For $s \leq t$

$$\mathbb{E}\left[\mathbb{E}[X|\mathcal{F}_t]\middle|\mathcal{F}_s\right] = \mathbb{E}[X|\mathcal{F}_s].$$

Given the tower law, there is a particularly important class of stochastic processes: the *martingales*.

Definition

A stochastic process X is a *martingale* if

- $\mathbb{E}[|X_t|] \leq \infty$ for all $t \geq 0$
- For any $s \leq t$, we have $X_s = \mathbb{E}[X_t | \mathcal{F}_s]$.

Knowing a process is a martingale seems not to tell us much (for example, we don't know how its variance changes through time), but it does give us enough to prove many surprising properties.

Suppose X is a martingale

- Then $\mathbb{E}[X_t] = X_0$ for all $t > 0$
- X admits limits from the right and left at every time
- If $\sup_t \mathbb{E}[X_t^2] < \infty$ then X converges as $t \to \infty$
- If X is continuous, then its paths have infinite length (more on this later)
- If X is continuous, it will have no nontrivial intervals on which it is increasing or decreasing
- ...

- We will also sometime need random variables taking values in $[0, T]$, this will be called *random times*
- If a random time τ has the property that $\{\tau \leq t\} \in \mathcal{F}_t$ for all t, then τ is called a *stopping time*.
- Stopping times have the property that "when you reach time τ, you will know you've reached time τ".
- An extraordinary property of martingales is that, for any two stopping times $\rho \leq \tau$,

$$M_\rho = \mathbb{E}[M_\tau | \mathcal{F}_\tau].$$

Definition

A process $(W_t)_{t \geq 0}$ is a *Brownian motion* if

- $\forall\ s \geq 0,\ t \geq 0,\ \big(W_{t+s} - W_t\big)$ is normally distributed with zero mean and variance s,

$$\mathbb{E}[\, W_{t+s} - W_t \,] = 0, \quad \mathbb{E}\big[\, (W_{t+s} - W_t)^2 \,\big] = s,$$

- if $0 \leq p \leq q \leq s \leq t$ then $(W_q - W_p)$ and $(W_t - W_s)$ are independent,
- the map $t \mapsto W_t$ is continuous, and
- $W_0 = 0$ (this is really a convention, it saves some writing).

It is not obvious that such a thing exists, but there are a number of ways of constructing it (see Etheridge S3.1 and S3.2, for example, or B8.2 lectures!).

Note that if $(W_t)_{t\geq 0}$ is a Brownian motion then so too are:

- $\hat{W}_t = W_{(t+t_0)} - W_{t_0}$ for any constant $t_0 \geq 0$ (or even any stopping time t_0)!;
- $\tilde{W}_t = c\,W_{(t/c^2)}$ for any constant $c > 0$.

It's easy to check that W is a martingale, as are:

$$W_t^2 - t \qquad \text{and} \qquad \exp(\theta W_t - \theta^2 t/2)$$

B8.3 Mathematical Models of Financial Derivatives

Mathematical Institute

Brownian Motion and Martingales

4.2 Quadratic variation

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- We need to define an integral against a Brownian motion or related process
- Given a process X, and an integrand H, the discrete 'integral' could be defined by

$$\sum_t H_t(X_{t+\delta t} - X_t)$$

If we know that X is differentiable with respect to t, then we can write $X_{t+\delta t} - X_t \approx X'_t \delta t$, and our limiting integral would then be

$$\int H_t X'_t dt.$$

- Sadly, this doesn't work for most interesting stochastic processes...

Theorem

A Brownian motion is not differentiable at any point, with probability one.

Proof (sketch):

- At a point $t_0 \geq 0$, we know that $\hat{W}_t = W_{t+t_0} - W_{t_0}$ is also a Brownian motion.
- If Brownian motion were differentiable at some time, then there's a stopping time t_0 at which it's differentiable.
- In this case, the limit

$$\lim_{t \to 0} \frac{W_{(t+t_0)} - W_{t_0}}{t} = \lim_{t \to 0} \frac{\hat{W}_t}{t}$$

would exist.

It is enough to show that the limit $\lim_{t\to 0} \frac{\hat{W}_t}{t}$ does not exist. Let A_n and B_n be defined by

$$A_n = \left\{ \frac{|\hat{W}_t|}{t} > n : \text{ for some } t \in \left(0, \frac{1}{n^4}\right] \right\},$$

$$B_n = \left\{ \frac{|\hat{W}_t|}{t} > n : \text{ at } t = \frac{1}{n^4} \right\}.$$

Clearly we have $B_n \subseteq A_n$ and so

$$\mathbb{P}(A_n) \geq \mathbb{P}(B_n) = \mathbb{P}\left(\frac{|\hat{W}_{1/n^4}|}{1/n^4} > n \right)$$
$$= \mathbb{P}\left(|n^2 \, \hat{W}_{1/n^4}| > \frac{1}{n} \right) = \mathbb{P}\left(|\tilde{W}_1| > \frac{1}{n} \right).$$

- As $n \to \infty$ we have $\mathbb{P}(|\tilde{W}_1| > 1/n) \to 1$.
- Therefore $\lim_{n\to\infty} \mathbb{P}(A_n) = 1$ which means that in this limit there is (with probability one) always some $0 < t \leq 1/n^4$ with $|\hat{W}_t|/t > n$.
- This shows that (with probability one) the limit which defines the derivative of a Brownian motion can not exist.

□

The fundamental issue is that Brownian motion (and, in fact, every nontrivial continuous martingale) is *too rough* to be differentiable. We will next see how we can measure the 'roughness' of a martingale in a useful way.

For a partition of $[0, t]$, $t_0 = 0 < t_1 < t_2 < \cdots < t_n = t$, let $|\pi| = \max_{0\leq k<n}(t_{k+1} - t_k)$.

Mathematical Institute

Definition

The 'quadratic variation' of a random process X_t is defined (if it exists) by

$$[X]_t = \mathbb{P}\text{-}\lim_{|\pi|\to 0} \sum_{k=0}^{n-1} (X_{t_{k+1}} - X_{t_k})^2.$$

Importantly, the limit here should be taken *in probability*, in the sense that

$$\mathbb{P}\left[\left|[X]_t - \sum_{k=0}^{n-1} (X_{t_{k+1}} - X_{t_k})^2\right| > \epsilon\right] \to 0 \text{ as } \epsilon \to 0.$$

Other common notation for quadratic variation include $[X]_t$, $\langle X\rangle_t$, $[X, X]_t$ and $\langle X, X\rangle_t$. These are identical for continuous processes.

Lemma

If X is continuously differentiable on $[0, t]$ then $[X]_t = 0$.

Proof: As $X_{t_{k+1}} - X_{t_k} = X'(\xi_k)(t_{k+1} - t_k)$ for some $\xi_k \in [t_k, t_{k+1}]$ we have

$$\sum_{k=0}^{n-1}(X_{t_{k+1}} - X_{t_k})^2 = \sum_{k=0}^{n-1} X'(\xi_k)^2(t_{k+1} - t_{t_k})^2$$

$$\leq |\pi| \sum_{k=0}^{n-1} X'(\xi_k)^2(t_{k+1} - t_k)$$

and as $|\pi| \to 0$, by convergence of Riemann integrals,

$$\sum_{k=0}^{n-1} X'(\xi_k)^2(t_{k+1} - t_k) \to \int_0^t X'(u)^2\, du < \infty. \quad \square$$

Lemma

If X is an increasing continuous function, then $[X]_t = 0$.

Proof:

- Consider a partition so that $|X_{t_{k+1}} - X_{t_k}| < \epsilon$ for all k.
- As X is continuous and increasing, such a partition exists and can be taken to have at most $\left\lceil \frac{X_T - X_0}{\epsilon} \right\rceil$ intervals.
- However, this means that

$$\sum_{k=0}^{n-1} |X_{t_{k+1}} - X_{t_k}|^2 \leq \left\lceil \frac{X_T - X_0}{\epsilon} \right\rceil \epsilon^2 \to 0.$$

$\square$

Theorem

The quadratic variation of a Brownian motion is given by $[W]_t = t$.

In fact, basically the same proof works for a general martingale (but we won't prove this), and can even be used to give a (in some ways better) definition of the quadratic variation...

Theorem

If X *is a martingale and you know that* $\{X_t^2 - Y_t\}_{t \geq 0}$ *is a martingale, for* Y *a continuous increasing process, then* $Y = [X]$.

Conversely, if X *is a martingale and* $\mathbb{E}\big[[X]_t\big] < \infty$, *then* $X_t^2 - [X]_t$ *is a martingale and* $\mathbb{E}\big[X_t^2\big] = \mathbb{E}\big[[X]_t\big]$.

Proof of Theorem

UNIVERSITY OF OXFORD
Mathematical Institute

We expand the expression for the quadratic variation and make use of our knowledge of the normal distribution.

Let $\{t_j\}_{j=0}^n$ denote the endpoints of the intervals that make up the partition π of $[0, t]$. Write $S = \sum_j (W_{t_j} - W_{t_{j-1}})^2$
First observe that

$$|S - t|^2 = \left| \sum_{j=1}^{n} \{ |W_{t_j} - W_{t_{j-1}}|^2 - (t_j - t_{j-1}) \} \right|^2.$$

Write $\delta_j = \left|W_{t_j} - W_{t_{j-1}}\right|^2 - (t_j - t_{j-1})$. Then

$$|S - t|^2 = \sum_{j=1}^{n} \left(\delta_j^2 + 2 \sum_{k>j} \delta_j \delta_k \right).$$

Note that since Brownian motion has independent increments,

$$\mathbb{E}\left[\delta_j \delta_k\right] = \mathbb{E}\left[\delta_j\right] \mathbb{E}\left[\delta_k\right] = 0 \qquad \text{if } j \neq k.$$

Also

$$\mathbb{E}[\delta_j^2] = \mathbb{E}\left[\left|W_{t_j} - W_{t_{j-1}}\right|^4 - 2\left|W_{t_j} - W_{t_{j-1}}\right|^2 (t_j - t_{j-1}) + \left(t_j - t_{j-1}\right)^2 \right].$$

For a normally distributed random variable, X, with mean zero and variance λ, $\mathbb{E}[|X|^4] = 3\lambda^2$. Therefore, we have

$$\begin{aligned}
\mathbb{E}\left[\delta_j^2\right] &= 3\left(t_j - t_{j-1}\right)^2 - 2\left(t_j - t_{j-1}\right)^2 + \left(t_j - t_{j-1}\right)^2 \\
&= 2\left(t_j - t_{j-1}\right)^2 \\
&\leq 2|\pi|\left(t_j - t_{j-1}\right).
\end{aligned}$$

Summing over j,

$$\begin{aligned}
\mathbb{E}\left[\left|S - t\right|^2\right] &\leq 2\sum_{j=1}^{n} |\pi|\left(t_j - t_{j-1}\right) \\
&= 2|\pi_n|T \\
&\to 0 \qquad \text{as } |\pi| \to 0. \quad \square
\end{aligned}$$

B8.3 Mathematical Models of Financial Derivatives

Brownian Motion and Martingales

4.3 Stochastic Integrals

Samuel Cohen

Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Oxford Mathematics

- As we've seen, we need to do something careful to define an integral 'against a Brownian motion'.
- The key idea that allows us to do this is due to Kiyosi Itô, and exploits the quadratic variation as a measure of the 'roughness' of a process.

Definition

The Itô integral of a stochastic process H against a martingale M is

$$\int_0^t H(u)\, dM_u = \lim_{|\pi|\to 0} \sum_{k=0}^{n-1} H(t_k)(M_{t_{k+1}} - M_{t_k}).$$

For fixed t this integral is a random variable and as t varies it is a stochastic process.

The limit here is taken in an L^2 sense (or in probability, see, e.g., Etheridge pp 78–85).

To show that this limit exists, we will assume (for simplicity) that H and $[M]$ are bounded. Using the tower law, writing $\delta M_k = M_{t_{k+1}} - M_{t_k}$ and $\mathbb{E}_t$ for $\mathbb{E}[\cdot|\mathcal{F}_t]$, we find that

$$\begin{aligned}\mathbb{E}\Big[\sum_{k=0}^{n-1} H(t_k)\delta M_k\Big] &= \mathbb{E}\Big[\sum_{k=0}^{n-1} \mathbb{E}_{t_k}\big[H(t_k)\delta M_k\big]\Big] \\ &= \mathbb{E}\Big[\sum_{k=0}^{n-1} H(t_k)\mathbb{E}_{t_k}[\delta M_k]\Big] \; = \; 0.\end{aligned}$$

Similarly, we have

$$\mathbb{E}\Big[\Big(\sum_{k=0}^{n-1} H(t_k)\,\delta W_k\Big)^2\Big] = \mathbb{E}\Big[\sum_{j,k=0}^{n-1} H(t_j)\,H(t_k)\,\delta M_k\,\delta M_j\Big]$$

$$= \mathbb{E}\Big[\sum_{k=0}^{n-1} \mathbb{E}_k\big[\,H(t_k)^2\,(\delta M_k)^2\,\big]\Big] + 2\mathbb{E}\Big[\sum_{j<k}^{n-1} \mathbb{E}_k\big[\,H(t_j)H(t_k)\delta M_j \delta M_k\,\big]\Big]$$

$$= \mathbb{E}\Big[\sum_{k=0}^{n-1} H(t_k)^2\,\mathbb{E}_k\big[\,(\delta M_k)^2\,\big]\Big] + 2\mathbb{E}\Big[\sum_{j<k}^{n-1} H(t_j)H(t_k)\delta W_j \mathbb{E}_k\big[\,\delta M_k\,\big]\Big]$$

$$= \mathbb{E}\Big[\sum_{k=0}^{n-1} H(t_k)^2 \mathbb{E}_{t_k}[(\delta M_k)^2]\Big] = \mathbb{E}\Big[\sum_{k=0}^{n-1} H(t_k)^2\big([M]_{t_{k+1}} - [M]_{t_k}\big)\Big]$$

As the quadratic variation is a continuous increasing process, we can define the integral

$$\int_0^t H(t)^2 d[M]_t = \lim_{|\pi|\to 0} \sum_{k=0}^{n-1} H(t_k)^2 \left([M]_{t_{k+1}} - [M]_{t_k}\right)$$

using Riemann sums.

Therefore, to define $\int f(t)dM_t$, we take a sequence of good approximations f_N to f, and writing $H = f_N - f_{N-1}$ we see that by dominated convergence

$$\mathbb{E}\Big[\Big(\sum_{k=0}^{n-1} f_N(t_k)\delta M_k - \sum_{k=0}^{n-1} f_{N'}(t_k)\,\delta M_k\Big)^2 \Big]$$
$$= \mathbb{E}\Big[\int_0^t \big(f_N(t) - f_{N'}(t)\big)^2 d[M]_t \Big] \to 0$$

- We conclude that the random variables

$$\sum_{k=0}^{n-1} f_N(t_k)\delta M_k$$

  are a Cauchy sequence in the space of square-integrable random variables, hence converge.
- We call the limit the integral $\int_0^t f(t)dM_t$.
- In fact, we can show a stronger convergence (it converges for every t simultaneously, see B8.2...)
- The integral we construct has various nice properties...

- Assuming H is sufficiently 'nice', we have

$$\mathbb{E}\Big[\int_0^t H(u)\,dM_u\Big] = 0.$$

- Furthermore, for $0 \leq s < t$

$$\mathbb{E}_s\Big[\int_0^t H(u)\,dM_u\Big] = \int_0^s H(u)\,dM_u.$$

In particular the process $Y_s = \int_0^s H(u)dM_u$ is a martingale.

- Our construction also tells us that the quadratic variation and the stochastic integral interact together nicely:

$$\text{If} \quad Y_t = \int_0^t H(u)dM_u \quad \text{then} \quad [Y]_t = \int_0^t H(u)^2 d[M]_u.$$

In particular

$$\mathbb{E}\Big[\Big(\int_0^t H(u)dM_u\Big)^2\Big] = \mathbb{E}\Big[\int_0^t H(u)^2 d[M]_u\Big]$$

- If $[Y]_t$ has finite expectation, this is enough to guarantee that Y is a martingale (you don't need H bounded).
- If $M = W$ is a Brownian motion, then $[W]_t = t$, so

$$\mathbb{E}\Big[\Big(\int_0^t H(u)dW_u\Big)^2\Big] = \int_0^t \mathbb{E}[H(u)^2]du$$

- In the very special case where H is deterministic (so it depends on time, but not on W), we see that $\int H dW$ is the limit of sums of Gaussian random variables. As sums of Gaussians are Gaussian (and L^2 limits of Gaussians are Gaussian), we know that

$$\int_0^t H(u) dW_u \sim N\Big(0, \int_0^t H(u)^2 du\Big).$$

- In general, however, the integral itself is not normally distributed. For example, $2\int_0^t W_u \, dW_u = W_t^2 - t$, which has a χ^2 distribution (shifted by t)

B8.3 Mathematical Models of Financial Derivatives

Brownian Motion and Martingales

4.4 Itô's lemma

Samuel Cohen
Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Theorem

If $f(x,\tau)$ is $C^{2,1}$ then

$$f(W_t,t) = f(0,0) + \int_0^t \frac{\partial f}{\partial \tau}(W_u,u)\,du + \int_0^t \frac{\partial f}{\partial x}(w_u,u)\,dW_u + \frac{1}{2}\int_0^t \frac{\partial^2 f}{\partial x^2}(W_u,u)d[W]_u.$$

Since $[W]_u = u$ we can replace $d[W]_u$ by du, and in practice we always do.

Proof (sketch)

UNIVERSITY OF OXFORD
Mathematical Institute

Consider the simpler case where f is independent of time and write

$$f(W_t) - f(0) = \sum_{k=0}^{n-1} \big(f(W_{k+1}) - f(W_k) \big)$$

over some partition, π, of $[0, t]$. Taylor's theorem (with remainders) shows that for each k

$$f(W_{k+1}) - f(W_k) = f'(W_k)\delta W_k + \frac{1}{2} f''(V_k)(\delta W_k)^2$$

for some V_k between W_k and W_{k+1}, where $\delta W_k = W_{k+1} - W_k$.

Thus

$$f(W_t) - f(0) = \sum_{k=0}^{n-1} f'(W_k)\delta W_k + \frac{1}{2}\sum_{k=0}^{n-1} f''(V_k)(\delta W_k)^2.$$

As we refine the partition

$$\lim_{|\pi|\to 0} \sum_{k=0}^{n-1} f'(W_k)\delta W_k \to \int_0^t f'(W_u)\, dW_u.$$

For the second sum, it can be shown that

$$\lim_{|\pi|\to 0} \sum_{k=0}^{n-1} f''(V_k)(\delta W_k)^2 \to \int_0^t f''(W_u)\, d[W]_u,$$

establishing that

$$f(W_t) - f(0) = \int_0^t f'(W_u)\, dW_u + \frac{1}{2}\int_0^t f''(W_u)\, d[W]_u.$$

□

- In practice, we usually write Itô's lemma in differential form rather than an integral form.
- If $f(W,t)$ is $C^{2,1}$ and we define $f_t = f(W_t,t)$ the differential form of Itô's lemma is

$$df_t = \left(\frac{\partial f}{\partial t}(W_t,t) + \frac{1}{2}\frac{\partial^2 f}{\partial W^2}(W_t,t) \right) dt + \frac{\partial f}{\partial W}(W_t,t)\, dW_t.$$

- This amounts to doing a regular Taylor series expansion of $f(W,t)$ then *pretending* that $dW_t^2 = dt$ (and ignoring terms of higher order than dt).

To solve the stochastic differential equation

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dW_t$$

we can proceed as follows.

- If $f(W,t) = e^{aW+bt}$ then all its partial derivatives are multiples of the function, so it makes sense to try $S_t = S_0\, e^{aW_t+bt}$.
- This gives

$$dS_t = (b\, S_t + \tfrac{1}{2}\, a^2\, S_t)\, dt + a\, S_t\, dW_t$$

or

$$\frac{dS_t}{S_t} = (b + \tfrac{1}{2}\, a^2)\, dt + a\, dW_t.$$

- If we set $a = \sigma$ and $b = \mu - \frac{1}{2}\sigma^2$ we recover our dynamics, i.e. the solution is

$$S_t = S_0 \exp\Big((\mu - \tfrac{1}{2}\sigma^2)\, t + \sigma W_t\Big).$$

- The process S_t is often called *geometric* Brownian motion.
- Note that the sign of S_t is determined by the sign of S_0.
- Taking $\mu = 0$ we see that S is a martingale, as $dS_t = \sigma S_t dW_t$, so S is an integral against W.

  This can also be checked using the moment generating function of a Gaussian distribution (which guarantees integrability).

Itô's lemma for solutions of SDEs

UNIVERSITY OF OXFORD
Mathematical Institute

Suppose that X_t is a solution of

$$X_t - X_0 = \int_0^t \mu(X_u, u)\, du + \int_0^t \sigma(X_u, u)\, dW_u,$$

and $f(x,t)$ is a $C^{2,1}$ function. Then

$$f(X_t, t) = f(X_0, 0) + \int_0^t \left(\frac{\partial f}{\partial t}(X_u, u) + \tfrac{1}{2}\sigma(X_u, u)^2 \frac{\partial^2 f}{\partial x^2}(X_u, u) \right) du + \int_0^t \frac{\partial f}{\partial x}(X_u, u)\, dX_u.$$

The proof amounts to showing the quadratic variation is given by $[X]_t = \int_0^t \sigma(X_u, u)^2\, du$.

In differential notation, which is how this result is normally used, if

$$dX_t = \mu(X_t, t)\,dt + \sigma(X_t, t)\,dW_t$$

and $f_t = f(X_t, t)$ then

$$df_t = \left(\frac{\partial f}{\partial t}(X_t, t) + \tfrac{1}{2}\sigma(X_t, t)^2\frac{\partial^2 f}{\partial x^2}(X_t, t)\right)\,dt + \frac{\partial f}{\partial x}(X_t, t)\,dX_t.$$

This can be obtained from a Taylor series expansion of $f(x, t)$ and *formally calculating* with $dX_t^2 = \sigma(X_t, t)^2\,dt$ (one approach to the proof does exactly this, but you have to check the limits work).

B8.3 Mathematical Models of Financial Derivatives

Brownian Motion and Martingales

4.5 Feynman–Kac Theorem

Samuel Cohen
Hilary Term 2021

University of Oxford
Mathematical Institute

Theorem (Feynman–Kac Theorem)

Suppose that $f(x,t)$ satisfies the terminal value problem

$$0 = \frac{\partial f}{\partial t} + \mu(x,t)\,\frac{\partial f}{\partial x} + \frac{1}{2}\sigma(x,t)^2\,\frac{\partial^2 f}{\partial x^2} - rf, \quad t < T,\ x \in \mathbb{R},$$

$$f(x,T) = F(x), \quad x \in \mathbb{R}.$$

Let X_t satisfy the stochastic differential equation

$$dX_t = \mu(X_t,t)\,dt + \sigma(X_t,t)\,dW_t$$

Then

$$f(x,t) = e^{-r(T-t)}\mathbb{E}_t\left[F(X_T)\,|\,X_t = x\right] \tag{1}$$

- This connects our SDEs to PDEs – solving a PDE gives us a solution to an SDE, and the expected values of solutions of SDEs give us solutions of PDEs
- This opens the doors to using Monte–Carlo (simulation) numerical methods to solve PDEs and pulling analytical tools from one area to the other.
- For our purposes, it means we will be able to use either PDE or stochastic tools to price derivatives, depending on what is easiest at each moment.

Proof (sketch)

University of Oxford
Mathematical Institute

Take $h = e^{r(T-t)} f$. Then

$$\frac{\partial h}{\partial t} = -rh + e^{r(T-t)} \frac{\partial f}{\partial t}, \quad \frac{\partial h}{\partial x} = e^{r(T-t)} \frac{\partial f}{\partial x}, \quad \frac{\partial^2 h}{\partial x^2} = e^{r(T-t)} \frac{\partial^2 f}{\partial x^2}$$

so h satisfies the PDE

$$0 = \frac{\partial h}{\partial t} + \mu(x,t) \frac{\partial h}{\partial x} + \frac{1}{2}\sigma(x,t)^2 \frac{\partial^2 h}{\partial x^2}, \quad t < T, \ x \in \mathbb{R},$$
$$h(x,T) = F(x), \quad x \in \mathbb{R}.$$

Note that Itô's lemma implies that

$$F(X_T) = h(X_T, T)$$
$$= h(X_t, t) + \int_t^T \sigma(X_s, s) \frac{\partial h}{\partial x}(X_s, s)\, dW_s$$
$$+ \int_t^T \left(\frac{\partial h}{\partial t}(X_s, s) + \mu(X_s, s) \frac{\partial h}{\partial x}(X_s, s) + \frac{1}{2}\sigma(X_s, s)^2 \frac{\partial^2 h}{\partial x^2}(X_s, s) \right) ds$$

By assumption, the integral on the second line vanishes. When we take expectations the integral on the first line also vanishes (as it is a martingale). Thus

$$\mathbb{E}[F(X_T)|\mathcal{F}_t] = h(X_t, t) = e^{r(T-t)} f(X_t, t)$$

The process X is a Markov process (for any reasonable μ, σ), so for $s < t$ the behaviour of X_s is independent of $\mathcal{F}_t$ given $X_t = x$.

Conditioning on $X_t = x$ (which is more restrictive than $\mathcal{F}_t$, as X is a Markov process) gives

$$f(x,t) = e^{-r(T-t)}\mathbb{E}_t[F(X_T)|X_t = x].$$

□

A simple (useful) example:

UNIVERSITY OF OXFORD
Mathematical Institute

Take

$$dX_t = adt + bdW_t$$

with $X_t = x$. Then we can see

$$X_T \sim N\big(x + a(T-t), b^2(T-t)\big)$$

.

From a stochastic perspective, we can compute

$$\begin{aligned} f(x,t) &= \mathbb{E}[f(X_T, T)|X_t = x] \\ &= \int_{\mathbb{R}} \frac{f(y,T)}{\sqrt{2\pi b^2(T-t)}} \exp\Big(-\frac{(y-x-a(T-t))^2}{2b^2(T-t)}\Big)dy. \end{aligned}$$

As $dX_t = adt + bdW_t$, the Feynman–Kac theorem connects f with the solution of the diffusion PDE

$$0 = \frac{\partial f}{\partial t} + \frac{b^2}{2}\frac{\partial^2 f}{\partial x^2} + a\frac{\partial f}{\partial x}.$$

- This PDE admits a unique solution for any smooth terminal value $f(x, T)$ with reasonable growth.
- Using integration by parts, you can verify that the formula on the previous page gives a solution .
- The Feynman–Kac theorem shows this solution must be the expected value of $f(X_T, T)$ (which we knew directly in this case!)
- This shows that the Gaussian density is the Green's function/fundamental solution for the PDE.

B8.3 Mathematical Models of Financial Derivatives

The Black–Scholes model

5.1 Basic Setup

Samuel Cohen

Hilary Term 2021

University of Oxford
Mathematical Institute

Oxford Mathematics

- Now that we have the tools of continuous time, we can build a model for derivative pricing.
- We will consider a slightly generalized version of the classic Black–Scholes model.

- We have an underlying stock S, and a bond B (representing a bank-account/interest).
- The Bond price satisfies $dB_t = rB_t dt$.
- We assume that S_t evolves as

$$\frac{dS_t}{S_t} = (\mu - q)\,dt + \sigma\,dW_t,$$

  where μ is known as the drift, q is the *continuous dividend yield* and $\sigma > 0$ is the volatility.
- W is a Brownian motion under the real probability measure $\mathbb{P}$.
- For *fixed* $T \geq 0$ the distribution of S_T is given by

$$S_T = S_0 \exp\Big(\,(\mu - q - \tfrac{1}{2}\sigma^2)\,T + \sqrt{\sigma^2\,T}\,Z\,\Big), \quad Z \sim N(0,1).$$

- Over each infinitesimal period $[t, t + dt)$ the share pays $q\, S_t\, dt$ in dividends, where for our purposes q is a constant known as the continuous dividend yield.
- This is a poor but widely used model for dividend paying shares.
- With reinvestment of dividends, one share at time zero grows to e^{qt} shares at time t and the total value at time t is $p_t = e^{qt}\, S_t$.
- Itô's lemma shows that

$$\frac{dp_t}{p_t} = \mu\, dt + \sigma\, dW_t.$$

- Our trading strategies will be given by random processes.
- If at the start of a period $[t, t+\delta t)$ we hold H_t stocks, at the end our increase in wealth will be $H_t(S_{t+\delta t} - S_t)$.
- Adding these increases together and taking a limit, we obtain a stochastic integral $\int_0^t H_u dS_u$ for the gains from trading in the stock
- By connecting our choice of H to the solution of a PDE (or using a powerful theorem from stochastic calculus), we can find trading strategies H which will eliminate the risk of holding an option, which will allow us to use no-arbitrage to identify prices.

B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

The Black–Scholes model

5.2 Hedging and replication

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

- We first consider obtaining the price through a hedging argument.
- Assume an option's payoff is give by $V_T = P_o(S_T)$ and its price $V_t = V(S_t, t)$.
- Set up a portfolio of one option and $-\Delta_t$ shares, so at t its market price at time t is

$$M_t = V_t - \Delta_t S_t.$$

- Let Π_t be the cumulative cost of executing this strategy, so

$$d\Pi_t = dV_t - \Delta_t \, dS_t - q \, \Delta_t \, S_t \, dt,$$

the final term represents payment of the dividend yield to the owner of the shares.

Itô's lemma applied to $V_t = V(S_t, t)$ gives

$$d\Pi_t = \left(\frac{\partial V}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 V}{\partial S^2}(S_t, t) - q\,\Delta_t\, S_t \right) dt + \left(\frac{\partial V}{\partial S}(S_t, t) - \Delta_t \right) dS_t,$$

which we make (instantaneously) risk-free by setting

$$\Delta_t = \frac{\partial V}{\partial S}(S_t, t).$$

A risk-free portfolio must grow at the risk-free rate, or there would be an arbitrage opportunity, so $d\Pi_t = r\,M_t\,dt$, i.e.,

$$\left(\frac{\partial V}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2\,S_t^2\,\frac{\partial^2 V}{\partial S^2}(S_t, t) - y\,S_t\,\frac{\partial V}{\partial S}(S_t, t) \right)$$

$$= r\left(V_t - S_t\frac{\partial V}{\partial S}(S_t, t) \right),$$

This gives the Black–Scholes equation

$$0 = \left(\frac{\partial V}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 V}{\partial S^2}(S_t, t) - y\, S_t \frac{\partial V}{\partial S}(S_t, t) \right) - \left(r\, V(S_t, t) - r\, S_t \frac{\partial V}{\partial S}(S_t, t) \right)$$

$$= \left(\begin{array}{c} \text{rate of return on risk-free} \\ \Delta\text{-hedged portfolio} \end{array} \right) - \left(\begin{array}{c} \text{rate of return on} \\ \text{portfolio's value} \\ \text{in bank} \end{array} \right)$$

This holds for all *attainable* S_t which, if $S_0 > 0$, is any $S_t > 0$.

Thus we obtain the Black–Scholes equation,

$$\frac{\partial V}{\partial t}(S,t) + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2}(S,t) + (r-q)\, S\, \frac{\partial V}{\partial S}(S,t) - r\, V(S,t) = 0,$$

for $S > 0$ and $t < T$.

At expiry $V_T = V(S_T, T) = P_{\rm o}(S_T)$ implies that

$$V(S,T) = P_{\rm o}(S), \quad S > 0.$$

Self-financing replication analysis

- Here we try to replicate the option's payoff using a portfolio of shares and bonds.
- The bond price, B_t, evolves as

$$\frac{dB_t}{B_t} = r\,dt.$$

- Let ϕ_t be the number of shares at t and ψ_t be the number of bonds.
- The market value of the portfolio at t is $\Phi_t = \phi_t\,S_t + \psi_t\,B_t$ and the change in the portfolio value is

$$d\Phi_t = \phi_t\,dS_t + \psi_t\,dB_t + (S_t + dS_t)\,d\phi_t + (B_t + dB_t)\,d\psi_t + q\,\phi_t\,S_t\,dt,$$

the final term coming from dividends.

- If $(S_t + dS_t)\, d\phi_t + (B_t + dB_t)\, d\psi_t = 0$ we say the portfolio is self-financing; to buy more shares we have to sell bonds and vice-versa.
- The self-financing condition is usually written as

$$d\Phi_t = \phi_t\, dS_t + \psi_t\, dB_t + q\, \phi_t\, S_t\, dt.$$

- In our case it reduces to

$$d\Phi_t = \phi_t\, dS_t + \psi_t\, r\, B_t\, dt + q\, \phi_t\, S_t\, dt.$$

If we write $\Phi_t = \Phi(S_t, t)$ and apply Itô's lemma we find

$$d\Phi_t = \left(\frac{\partial \Phi}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 \, S_t^2 \, \frac{\partial^2 \Phi}{\partial S^2}(S_t, t) \right) \, dt + \frac{\partial \Phi}{\partial S}(S_t, t) \, dS_t$$

and matching the deterministic and stochastic terms with the self-financing condition gives

$$\frac{\partial \Phi}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 \, S_t^2 \, \frac{\partial^2 \Phi}{\partial S^2}(S_t, t) = r \, \psi_t \, B_t + q \, \phi_t \, S_t,$$

$$\frac{\partial \Phi}{\partial S}(S_t, t) = \phi_t.$$

Eliminating $\psi_t B_t$ using the market value of the portfolio gives

$$\frac{\partial \Phi}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 \Phi}{\partial S^2}(S_t, t) = r\left(\Phi(S_t, t) - S_t \frac{\partial \Phi}{\partial S}(S_t, t)\right) + q\, S_t \frac{\partial \Phi}{\partial S}(S_t, t)$$

for any attainable S_t, i.e., any $S_t > 0$. Rearranging shows that any self-financing portfolio's price function must satisfy, for $S > 0$,

$$\frac{\partial \Phi}{\partial t}(S, t) + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 \Phi}{\partial S^2}(S, t) + (r - q)\, S \frac{\partial \Phi}{\partial S}(S, t) - r\, \Phi(S, t) = 0,$$

Finally, we apply the replication condition that the value of the portfolio at T always equals the payoff of the option, i.e.,

$$\Phi(S, T) = P_{\mathrm{o}}(S), \quad S > 0.$$

Then we argue that as the option and the portfolio have exactly the same cash-flows prior to expiry (in both cases here, no cash-flows) and exactly the same values at expiry they must have the same values now, i.e.,

$$V(S, t) = \Phi(S, t).$$

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD

Mathematical Institute

The Black–Scholes model

5.3 Risk-neutral measures

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

- Just as in discrete time, in the Black–Scholes model we also can use the fundamental theorem of asset pricing.
- This is the easiest way to calculate prices in some settings, while solving PDEs is easier in others.
- For simplicity, we will ignore dividends in this lecture.
- The result depends on an abstract result in stochastic calculus.

Theorem (Girsanov's Theorem)

Let X, Y be continuous martingales under a measure $\mathbb{P}$. Define M to be the solution to the SDE

$$M_t = 1 + \int_0^t M_s dX_s = \exp(X_t - [X]_t/2)$$

Then, provided M is sufficiently integrable, we can define a new measure $\mathbb{Q}$ by $\mathbb{Q}(A) = \mathbb{E}_{\mathbb{P}}[1_A M_T]$. Under this new measure,

$$Y_t - [X, Y]_t$$

is a continuous (local) martingale.

The integrability assumptions to make everything work are a little technical, but a sufficient condition (Novikov's condition) is that $\mathbb{E}_{\mathbb{P}}[\exp([X]_t/2)] < \infty$ for all t.

In practice, this theorem is typically used to show:

- If $dX_t = H_s dW_t$ and $Y_t = W_t$, for W a Brownian motion under probabilites $\mathbb{P}$,
- then there is a measure $\mathbb{Q}$ such that $W_t - \int_0^t H_s ds$ is a martingale.

In a Black–Scholes world, we have

$$\begin{aligned} dS_t &= \mu S_t dt + \sigma S_t dW_t \\ &= r S_t dt + \sigma S_t dW_t^{\mathbb{Q}} \end{aligned}$$

where W is a BM under $\mathbb{P}$, $W^{\mathbb{Q}}$ is a BM under $\mathbb{Q}$. It follows that $e^{-rt} S_t$ is a $\mathbb{Q}$-martingale.

We can join this with another theorem from stochastic calculus:

Theorem (Martingale representation theorem)

Let W be a Brownian motion, X be a local martingale in the filtration generated by W. Then there exists a process H such that

$$X_t = X_0 + \int_0^t H_s dW_s$$

- Extensions of this theorem, to non-Brownian settings, are possible.
- The Black–Scholes model is essentially Brownian (after rearrangement, you can see that S and W generate the same filtration)

By combining these results, we can establish a fundamental theorem of asset pricing (at least in a Brownian setting)

Theorem

Suppose we have a Black–Scholes model. Then there exists a (unique) measure $\mathbb{Q}$ such that every traded European claim (with no intermediate payments) has no-arbitrage time-t price

$$X_t = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[X_T|\mathcal{F}_t],$$

that is, the discounted price $e^{-rt}X_t$ is a martingale under $\mathbb{Q}$.

Proof

UNIVERSITY OF OXFORD
Mathematical Institute

- Find the measure $\mathbb{Q}$ as above.
- Use the martingale representation theorem to show that the martingale $e^{-rt}X_t := \mathbb{E}_{\mathbb{Q}}[e^{-rT}X_T|\mathcal{F}_t]$ can be written

$$e^{-rt}X_t = X_0 + \int_0^t H_s dW_s^{\mathbb{Q}} = X_0 + \int_0^t \frac{H_s}{\sigma S_s}(dS_s - rS_s ds).$$

- Writing $\Delta_t = H_t/(\sigma S_t)$, and applying Itô's lemma, we see that X_t satisfies the self-financing condition

$$dX_t = r(X_t - \Delta_t S_t)dt + \Delta_t dS_t.$$

- Therefore, X_t is the value of a self financing portfolio with terminal value X_T, so by no-arbitrage is the price of the asset.

□

B8.3 Mathematical Models of Financial Derivatives

The Black–Scholes model

5.4 General solution of the Black–Scholes problem

Samuel Cohen
Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Oxford Mathematics

- Given we have described the price of a claim (both in terms of PDEs and in terms of the risk-neutral measure), we can now give an explicit description of the price of a European claim with payoff $P_o(S_T)$, for P_o some function
- The price we derive can be seen both from a PDE and a probabilistic perspective.

The Black–Scholes problem for the price function of a European option with payoff given by $V_T = P_o(S_T)$ is

$$0 = \frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r-q)\, S \frac{\partial V}{\partial S} - r\, V, \quad S > 0,\ t < T,$$

$$V(S,T) = P_o(S), \qquad S > 0.$$

If we set $V(S,t) = e^{-r(T-t)}\, U(S,t)$ then

$$0 = \frac{\partial U}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 U}{\partial S^2} + (r-q)\, S \frac{\partial U}{\partial S}, \quad S > 0,\ t < T,$$

$$U(S,T) = P_o(S), \qquad S > 0.$$

If we start with the PDE, the Feynman–Kac formula shows that

$$U(S,t) = \mathbb{E}_{\mathbb{Q}}\left[\, P_{\mathrm{o}}(S_T) \,|\, S_t = S \,\right],$$

where S_t evolves according to

$$\frac{dS_t}{S_t} = (r-q)\,dt + \sigma\,dW_t^{\mathbb{Q}}.$$

This means that the option's price can be written as

$$V(S,t) = e^{-r(T-t)}\,\mathbb{E}_{\mathbb{Q}}\left[\, P_{\mathrm{o}}(S_T) \,|\, S_t = S \,\right].$$

- We know that if $S_t = S$ then

$$S_T = S \exp\big((r - q - \tfrac{1}{2}\sigma^2)\tau + \sigma\, W_\tau^{\mathbb{Q}} \big), \quad \tau = T - t$$

- As $W_\tau^{\mathbb{Q}} \sim N(0, \tau)$ we can write $\sigma\, W_\tau^{\mathbb{Q}} = \sqrt{\sigma^2 \tau}\, Z$ where $Z \sim N(0, 1)$.
- We now have a choice: either we calculate the density of S_T, or we use the density of Z directly.

Option 1: Calculate density of S

We compute the cumulative distribution function for S_T, for $x > 0$, as follows

$$\begin{aligned}F_T(x) &= \mathrm{prob}(S_T < x)\\ &= \mathrm{prob}\big(\log(S_T) < \log(x)\big)\\ &= \mathrm{prob}\Big(\sigma\, W_\tau^{\mathbb{Q}} < \log(x/S) - (r - q - \frac{1}{2}\sigma^2)\tau\Big)\\ &= \mathrm{prob}\big(Z < d_*\big) = \mathrm{N}(d_*),\end{aligned}$$

where

$$d_* = \frac{\log(x/S) - (r - q - \frac{1}{2}\sigma^2)(T-t)}{\sqrt{\sigma^2\,(T-t)}},$$

$$\mathrm{N}(d_*) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{d_*} e^{-p^2/2}\,dp.$$

Differentiating $F_T(x)$ with respect to x gives the probability density function for S_T, conditional on $S_t = S$,

$$f_T(x) = \frac{\exp\left(-\frac{1}{2}d_*^2\right)}{x\sqrt{2\pi\,\sigma^2\,(T-t)}}, \quad x > 0,$$

and so we arrive at an explicit formula for the option price,

$$V(S,t) = e^{-r(T-t)}\frac{1}{\sqrt{2\pi\,\sigma^2\,(T-t)}}\int_0^\infty P_\mathrm{o}(x)\,\exp\left(-\tfrac{1}{2}d_*^2\right)\frac{dx}{x},$$

where d_* depends on x (as well as S, $r-q$, σ and $(T-t)$).

This approach also shows us that the Black–Scholes equation admits a fundamental solution/Green's function

$$G(x) = \frac{1}{x\sqrt{2\pi\,\sigma^2\,(T-t)}} \exp\left(-r(T-t) - \tfrac{1}{2} d_*^2 \right)$$

where

$$d_* = \frac{\log(x/S) - (r - q - \frac{1}{2}\sigma^2)(T-t)}{\sqrt{\sigma^2\,(T-t)}}$$

It is possible (but tedious) to derive this from the fact that $\frac{1}{\sqrt{2\pi t}} e^{-\frac{x^2}{2t}}$ is a fundamental solution to the Heat equation $\partial_t u = (1/2)\partial_{xx} u$, using repeated changes of variables, thereby avoiding all probability.

Option 2: Use explicit formula for S

UNIVERSITY OF OXFORD
Mathematical Institute

We know that

$$S_T = S \exp\big((r - q - \tfrac{1}{2}\sigma^2)\tau + \sigma\sqrt{\tau} Z \big)$$

Hence, as $Z \sim N(0,1)$, the price is given by

$$\begin{aligned} & e^{-r\tau}\mathbb{E}_{\mathbb{Q}}[P_o(S_T)|S_t = S] \\ &= e^{-r\tau}\int_{-\infty}^{\infty} P_o\Big(Se^{(r-q-\frac{1}{2}\sigma^2)\tau+\sigma\sqrt{\tau}z}\Big)\frac{1}{\sqrt{2\pi}}e^{-z^2/2}dz \end{aligned}$$

This can be reduced (by changes of variables) to the formula obtained above, and may be simpler to use in some problems.

UNIVERSITY OF OXFORD

Mathematical Institute

B8.3 Mathematical Models of Financial Derivatives

The Black–Scholes model

5.5 Call and Put option prices

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- We now seek to give an explicit formula for Call and Put option prices, in a Black–Scholes model.
- By Put–Call parity, it's enough to consider Call options
- Various approaches are possible, but the simplest is arguably to notice that

$$
\begin{aligned}
(S-K)^+ = \max\{S-K, 0\} &= (S-K)1_{\{S>K\}} \\
&= S1_{\{S>K\}} - K1_{\{S>K\}}.
\end{aligned}
$$

We will find the value of each of the terms on the right-hand side.

In order to use our second option from the previous lecture, without overwhelming notation, it's convenient to first show that if

$$X = \exp(a + bZ) \qquad \text{for } Z \sim \mathcal{N}(0,1), \quad b > 0,$$

we can compute

$$\begin{aligned}\mathbb{E}[1_{\{X>K\}}] &= \int 1_{\{\exp(a+bZ)>K\}} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz \\ &= \int 1_{\left\{z>\frac{\log(K)-a}{b}\right\}} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz \\ &= 1 - \mathrm{N}\Big(\frac{\log(K)-a}{b}\Big) = \mathrm{N}\Big(\frac{a-\log(K)}{b}\Big).\end{aligned}$$

A similar calculation, made easier by noticing that

$$a + bz - \frac{z^2}{2} = a + \frac{b^2}{2} - \frac{(z-b)^2}{2},$$

gives that for $X = \exp(a + bZ)$,

$$\begin{aligned}
\mathbb{E}[X\mathbb{1}_{\{X>K\}}] &= \int e^{a+bz}\mathbb{1}_{\{\exp(a+bz)>K\}}\frac{e^{-z^2/2}}{\sqrt{2\pi}}dz \\
&= e^{a+\frac{b^2}{2}}\int \mathbb{1}_{\left\{z-b>\frac{\log(K)-a}{b}-b\right\}}\frac{e^{-(z-b)^2/2}}{\sqrt{2\pi}}dz \\
&= e^{a+\frac{b^2}{2}}\mathrm{N}\left(\frac{a-\log(K)}{b}+b\right)
\end{aligned}$$

Hence, as we know that for some $Z \sim \mathcal{N}(0,1)$,

$$S_T = S_t \exp\left(\left(r + \frac{\sigma^2}{2}\right)(T-t) + \sigma(W_T - W_t)\right)$$
$$= \exp\left(\log(S_t) + \left(r + \frac{\sigma^2}{2}\right)\tau + \sigma\sqrt{\tau}Z\right)$$

The prices of our two components simplify down to

$$e^{-r\tau}\mathbb{E}_{\mathbb{Q}}[K1_{\{S_T>K\}}|S_t = S] = e^{-r\tau}K\mathrm{N}(d_-)$$
$$e^{-r\tau}\mathbb{E}_{\mathbb{Q}}[S_T1_{\{S_T>K\}}|S_t = S] = Se^{-q\tau}\mathrm{N}(d_+)$$

with $d_\pm = \frac{\log(S/K)+(r-q\pm\frac{1}{2}\sigma^2)\tau}{\sigma\sqrt{\tau}}$.

Combining these calculations, we obtain the celebrated Black–Scholes formula for the price of a European call option:

$$C(S,t) = S\,e^{-q(T-t)}\,\mathrm{N}(d_+) - K\,e^{-r(T-t)}\,\mathrm{N}(d_-),$$

where

$$d_\pm = \frac{\log(S/K) + (r - q \pm \frac{1}{2}\sigma^2)(T-t)}{\sqrt{\sigma^2(T-t)}},$$

$$\mathrm{N}(x) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{x} e^{-\frac{1}{2}p^2}\,dp.$$

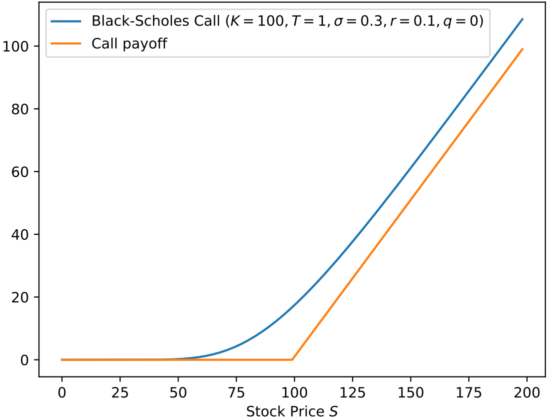

A Call option price (and payoff)

By differentiating with respect to S, we see that the Δ for a call option is

$$\Delta_c(S,t) = (\partial C/\partial S) = e^{-q(T-t)}\,\mathrm{N}(d_+).$$

(Note, this is not trivial, as $d_\pm$ depend on S).

Using Put–Call parity, we see that

$$C(S,t) - P(S,t) = e^{-q(T-t)}S - e^{-r(T-t)}K$$

so

$$\begin{aligned} P(S,t) &= C(S,t) - e^{-q(T-t)}S + e^{-r(T-t)}K \\ &= S\,e^{-q(T-t)}(\mathrm{N}(d_+) - 1) - K\,e^{-r(T-t)}(\mathrm{N}(d_-) - 1) \\ &= K\,e^{-r(T-t)}\mathrm{N}(-d_-) - S\,e^{-q(T-t)}\mathrm{N}(-d_+) \end{aligned}$$

The Δ for a put option is

$$\Delta_p(S,t) = -e^{-q(T-t)}\,\mathrm{N}(-d_+) = \Delta_c(S,t) - e^{-q(T-t)}$$

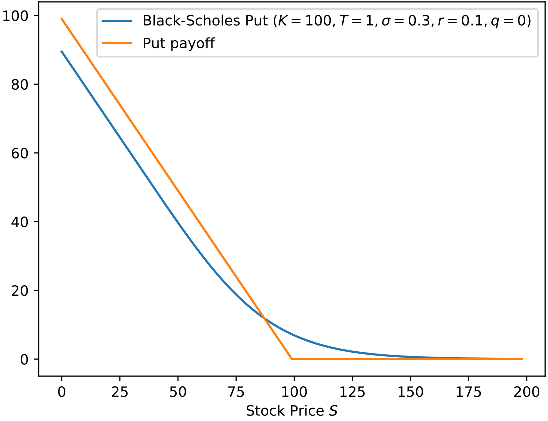

A Put option price (and payoff)

- For a call option, as $S \to \infty$, we have

$$d_+ \to \infty \qquad \text{and} \qquad d_- \to \infty,$$

  so $C^{BS} \approx S_t e^{q(T-t)} - Ke^{r(T-t)}$. This is natural, as the call we can be sure that the call will be exercised, so its value is similar to the value of a forward with the corresponding strike.
- Conversely, the put option's price converges to zero, for the same reason.
- As $S \to 0$ we have $d_\pm \to -\infty$, so the call price converges to zero, and the put price converges to $Ke^{-r(T-t)} - S_t e^{-q(T-t)}$, which is the price of the short forward with strike K.

As $t \to T$, we have

$$d_\pm \to \begin{cases} \infty & \text{if } S > K \\ -\infty & \text{if } S < K \end{cases}$$

so $C^{BS} \to (S-K)^+$, that is, the price of the option converges to its payoff.

UNIVERSITY OF OXFORD

Mathematical Institute

B8.3 Mathematical Models of Financial Derivatives

The Black–Scholes model

5.6 Properties of the Black–Scholes equation

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- As we will use it regularly, it's worth thinking a little about the Black–Scholes PDE and its solutions
- The Black-Scholes equation is

$$\mathcal{L}_{BS} V := \frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r-q)\, S \frac{\partial V}{\partial S} - r\, V = 0,$$

- This equation should hold for all $S > 0$ and $t < T$.

Mathematical Institute

The Black–Scholes equation has the following properties:

1. it is linear;
2. it is solved *backwards* in time, for $t < T$;
3. if $V(S,t)$ is a solution so too is $V(\lambda S, t)$ for any $\lambda > 0$;
4. $V(S,t)$ depends on t and T only through the combination $T - t$;
5. if $V(S,t)$ is a solution so too is $S\left(\partial V/\partial S\right)$ (and, by induction, so too are $S^n\left(\partial^n V/\partial S^n\right)$ for $n = 2, 3, \ldots$);
6. if $V(S,t)$ is a solution so too is

$$\hat{V}(S,t) = (S/A)^{2\alpha}\, V\left(B^2/S, t\right), \quad 2\alpha = 1 - 2(r-q)/\sigma^2,$$

for any constants $A > 0$, $B^2 > 0$.

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD
Mathematical Institute

Implied volatility

6.1 Calculating Implied Volatility

Samuel Cohen
Hilary Term 2021

The Black–Scholes formula gives us the value of an option as a function of the following inputs:

- The current price of the underlying S_t
- The time to maturity $\tau = T - t$
- The risk-free interest rate r
- The volatility σ

How do we choose the value of r and σ?

UNIVERSITY OF OXFORD
Mathematical Institute

How to choose r?

- The central bank sets a 'risk free' interest rate, which might suggest a good value of r.
- However, this rate is not representative of the cost of *borrowing* (as we cannot borrow at the rate r in reality).
- Practically, to get a good value of r from market data, we can look at the forward price for S.
- Remember that the forward price (assuming constant interest) is given by $F_t = e^{r(T-t)}S_t$. Hence $\log(S_t/F_t) = r(T-t)$, which gives us a representative interest rate for participants in the forward market.

How to choose σ?

- One idea: estimate from historical data. Split the interval $[t - N\delta, t]$ with a partition $\{t - N\delta = t_0, ..., t_N = t\}$ with $|t_n - t_{n+1}| = \delta$. Then an unbiased and efficient estimator of σ^2 is

$$\sigma^2 = \frac{1}{(N-1)\delta t} \sum_{i=1}^{N} \left(\frac{S_{t_{i+1}} - S_{t_i}}{S_{t_i}} \right)^2 .$$

- This is natural. However this method has certain drawbacks due to the fact that, in reality, σ varies significantly over time, and it is rather hard to capture its "most recent" value: the estimator becomes less reliable as we decrease t, while, if we increase it, we obtain an "averaged" value of σ over time.
- It also has the problem that it only looks at the past, while the price should reflect what the market expects to happen through the future.

Another way is to deduce σ from the prices currently observed in the market. This gives rise to **Implied Volatility**:

Definition

Given a market price of an option V^{mkt_t}, the implied volatility σ_{imp} is the value such that

$$V_t^{\text{mkt}} = V^{BS}(S_t, t; \sigma_{\text{imp}}).$$

Implied volatility is very common for call and put option (in many cases, markets quote prices in terms of implied volatility, rather than the price of the option).

Implied Volatility is well defined, provided

- $V_t^{\text{call,mkt}}$ is **not impossible**, that is,

$$V_t^{\text{call,mkt}} \in \left[\left(S - Ke^{-r(T-t)} \right)^+, S \right].$$

- And $\frac{\partial}{\partial \sigma} V^{\text{call,BS}}$ **does not change its sign**, as a function of σ.

The first condition is satisfied in practice (up to model error), since otherwise there is a **model-independent arbitrage**.

The second condition is satisfied, since, as we'll see, the BS Vega is always nonnegative. This also makes a simple numerical search algorithm work well.

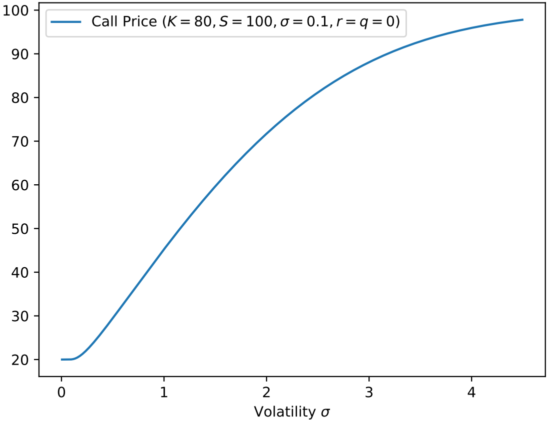

Black-Scholes price of a call option with strike $K = 80$ and underlying $S = 100$, as a function of the volatility σ

B8.3 Mathematical Models of Financial Derivatives

Implied volatility

6.2 Implied Smile and volatility surfaces

Samuel Cohen

Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Oxford Mathematics

- If the BS model was true, there would exist one value of implied volatility σ_{imp} for call options of all strikes and maturities.
- However, **this is not true** in practice. Typically, for each pair (T, K), we have a different value of implied vol $\sigma_{imp}(T, K)$.
- Plotted as a function of *negative log-moneyness* $x = \log(K/S)$, this function is typically *convex* around $x = 0$, and, hence, is often referred to as the **implied smile**.

(Recall that log-moneyness of a call or put option is defined as $\log(S/K)$.)

- In equities (where S is the price of a stock or stock index), the implied smile typically has a **negative skew**, assigning higher values to negative $x = \log(K/S)$ (i.e. $K < S$).
- The implied volatility is often higher for long and short maturity options, and lower for intermediate maturities
- By plotting the implied volatility for all traded values of (T, K), we obtain the implied volatility surface.
- This is commonly used to give an understanding of market perspectives on the riskiness of options trades, and as an input to risk management and hedging decisions.

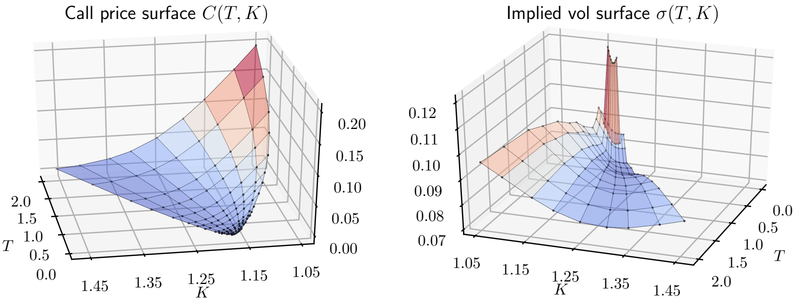

End of Day EURUSD Call options listed on Bloomberg on 31st May 2018 (Over the Counter options)

Spot: 1 EUR = 1.1694 USD

UNIVERSITY OF OXFORD

Mathematical Institute

B8.3 Mathematical Models of Financial Derivatives

Greeks

7.1 Δ (Delta)

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- Assume the market is described by a BS model, and denote the price function of an option by $V(S,t)$.
- **Sensitivities** of the option price V with respect to the input **variables** (S and t) and **parameters** (σ and r) are called the **Greeks**.
- These sensitivities are very important for **hedging** and **risk management**, as they show how the value of the option changes with small changes in the **uncertain** input!

We have already encountered the first (and most important) of the 'Greeks'

Definition

Delta is defined as

$$\Delta = \frac{\partial}{\partial S} V,$$

and it is the primary sensitivity, as, even if the model is true, the value of underlying will change, and its change is likely to be of a **higher magnitude than the time increment**.

Notice that

$$S_{t+\delta t} \approx S_t \mu \delta t + S_t \sigma \xi \sqrt{\delta t},$$

where ξ is a standard normal.

- For a call option, $\Delta^{\text{call}} = e^{-q(T-t)}\mathrm{N}(d_+)$
- By put-call parity, $\Delta^{\text{put}} = \Delta^{\text{call}} - 1 = -e^{-q(T-t)}\mathrm{N}(-d_+)$.
- The Δ-**hedging portfolio**, consisting of Δ^{call} units of S and $V_t - \Delta^{\text{call}} S$ units in bonds is an "instantaneous perfect hedge' of a short position in the option, assuming the Black–Scholes Model,
- In more complicated models, and in practice, it is a reasonably good hedge.

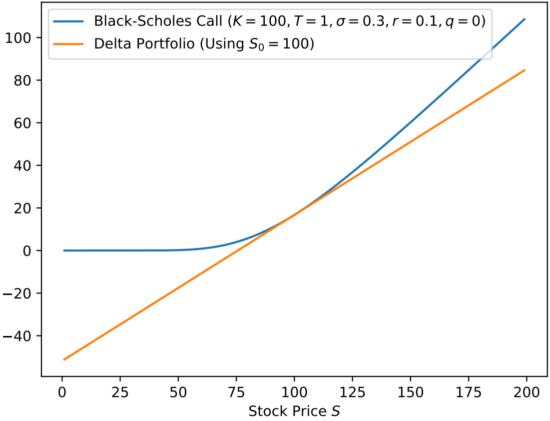

Option price as a function of S, with the tangent line given by the Delta-hedging portfolio value.

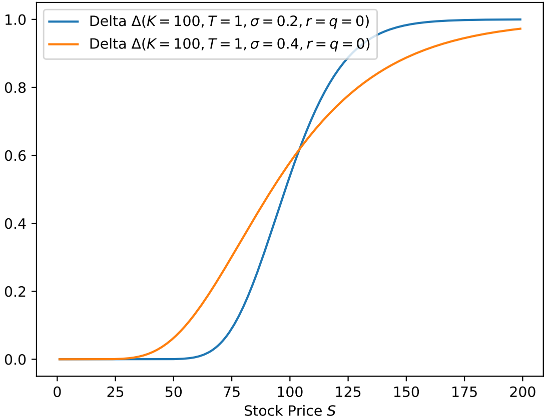

Black–Scholes Δ^{call} as a function of S.

- For example, suppose we are short a call option with $r = 0.05, \sigma = 0.2, T = 1, K = S_0 = 100.$
- Then $V^{\text{call}} = 10.4506,\ \Delta^{\text{call}} = 0.63683$
- We buy a portfolio consisting of
  - $\Delta S = \$63.683$ invested in stocks (buy $\Delta = 0.63683$ stocks)
  - $-V_t = \$10.4506$ in the option (our short position)
  - $V_t - \Delta S = -\$53.2325$ in bonds (i.e. borrow \$53.2325 risk free)
- Our total position is zero.

- Suppose the next day $S = 101$. Now $T \approx 249/250$, and the option price is \$11.0708.
- Our borrowing is now \$53.2431 due to interest, so our total portfolio is worth

$$0.63683 \times 101 - 11.0708 - 53.2431 = 0.006$$

- By using this 'Δ-neutral porfolio, we have eliminated the loss due to the increase in the option price.

- Suppose instead the next day $S = 110$. Now $T \approx 249/250$, and the option price is \$17.6365.
- Our borrowing is still now \$53.2431 due to interest, so our total portfolio is worth

$$0.63683 \times 101 - 17.6365 - 53.2431 = -0.8283$$

- We have partly minimized the loss due to the increase in the option price, but not as effectively as for a smaller change.

- We have so far assumed we can trade continuously, but in practice, we can only trade finitely often.
- As a result, we encounter the **discretization error** – the price of the hedging (replicating) portfolio no longer coincides with the option price at all times.
- In particular, we **cannot keep both positions** – in the stock and in bonds – **as prescribed by the Black–Scholes model**.
- Therefore, at each moment of rebalancing, we have to choose whether
  - we keep Δ (the amount of shares of stock) as prescribed by the model, and invest the rest of the available capital in bonds (or borrow the required amount by shorting),
  - or keep the amount of money in bonds as prescribed by the model, and invest the rest in the stock.

- Typically traders choose to keep the value of Δ as prescribed by the model, because changes in the stock price are more significant than changes in the value of bonds.
- This strategy is called Δ-**hedging**.

B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

Greeks

7.2 Γ (Gamma)

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

Definition

Gamma is the sensitivity of Δ with respect to changes in S:

$$\Gamma = \frac{\partial}{\partial S}\Delta = \frac{\partial^2}{\partial S^2}V.$$

- Γ measures how fast the hedging weight Δ changes with the changes in the underlying.
- This is important since, as mentioned before, in practice we only trade at discrete times.
- The smaller the curvature of the price of an option, as a function of S, the smaller the error of the discretization hedge – the difference between the price function and its tangent line around the point of tangency.
- Thus, **Gamma measures the discretization error**.

Let's make this statement more precise. Assume we have short-sold an option and set up the hedging portfolio at time t:

$$\gamma_t = V_t - \Delta_t S_t, \qquad \Delta_t = \frac{\partial}{\partial S} V(S_t, t)$$

Then the hedging error at time $t + \delta t$ is given by

$$\begin{aligned}
&\gamma_t B_{t+\delta t} + \Delta_t S_{t+\delta t} - V(S_{t+\delta t}, t + \delta t) \\
&= (V(S_t, t) - \Delta_t S_t)(1 + r\delta t) + \Delta_t S_{t+\delta t} - V(S_{t+\delta t}, t + \delta t)
\end{aligned}$$

Recall that $S_{t+\delta t} \approx S_t + r\delta t S_t + \sigma S_t (W_{t+\delta t} - W_t)$ and that $(W_{t+\delta t} - W_t)^2 \approx \delta t$.

Expanding using Taylor series, our error becomes

$$\begin{aligned}
&V_t - S_t \frac{\partial}{\partial S} V_t + r\delta t \left(V_t - S_t \frac{\partial}{\partial S} V_t\right) \\
&\quad - V_t - \delta t \frac{\partial}{\partial t} V_t - (S_{t+\delta t} - S_t) \frac{\partial}{\partial S} V_t \\
&\quad - \frac{1}{2}\sigma^2 S_t^2 (W_{t+\delta t} - W_t)^2 \frac{\partial^2}{\partial S^2} V_t + S_{t+\delta t} \frac{\partial}{\partial S} V_t + o(\delta t) \\
&= -\frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2}{\partial S^2} V_t ((W_{t+\delta t} - W_t)^2 - \delta t) + o(\delta t) \\
&= -\frac{1}{2}\sigma^2 S_t^2 \Gamma_t ((W_{t+\delta t} - W_t)^2 - \delta t) + o(\delta t),
\end{aligned}$$

where $o(\delta t)$ is a function satisfying

$$\frac{o(\delta t)}{\delta t} \to 0, \qquad \text{as } \delta t \to 0.$$

- Notice that $(\delta W_t)^2 - \delta t$ is a random variable with zero mean and variance $2(\delta t)^2$.
- We conclude Γ_t scales the main term in the discretization error of the hedge.
- If $\Gamma_t < 0$ (short Gamma), the hedged portfolio benefits from large market moves, and loses on small ones.
- If $\Gamma_t > 0$ (long Gamma) – vice versa.
- If we hedge a long position in the option, the opposite conclusions hold.

The Γ of a call is

$$\Gamma^{call} = \frac{e^{-\frac{1}{2}d_+^2}}{S\sigma\sqrt{2\pi(T-t)}}$$

As $t \to T$, $\Gamma^{call} \to \frac{\partial^2}{\partial S^2}(S-K)^+ = \delta(S-K)$

Due to put–call parity, the put Γ is the same.

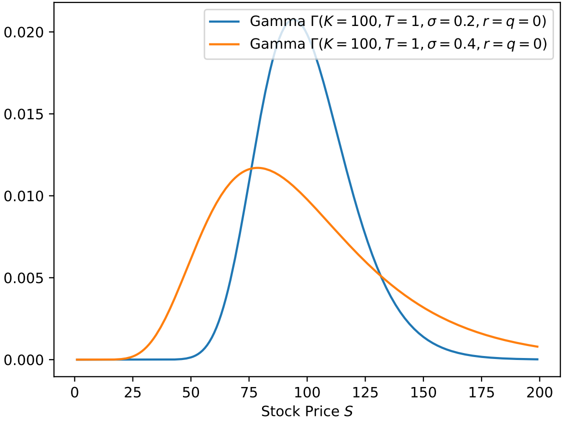

Gamma Γ(K = 100, T = 1, σ = 0.2, r = q = 0)
Gamma Γ(K = 100, T = 1, σ = 0.4, r = q = 0)
0.020
0.015
0.010
0.005
0.000
0
25
50
75
100
125
150
175
200
Stock Price S

- We can reduce the discretization error of the hedge over the first time step by **Gamma-hedging**.
- We cancel the current (instantaneous) Gamma of our option V by opening a position in another option V^1. Typically, we hedge an exotic option with underlying and a vanilla call or put. The current value of the resulting portfolio is given by

$$-V_t + \Delta^1 V_t^1 + \Delta S_t + \gamma B_t = 0,$$

due to self-financing. We would like it to stay close to zero at time $t + \delta t$.

- The above portfolio is **Gamma-neutral** if

$$\frac{\partial^2}{\partial S^2} V_t - \Delta^1 \frac{\partial^2}{\partial S^2} V_t^1 = 0.$$

So

$$\Delta^1 = \frac{\frac{\partial^2}{\partial S^2} V_t}{\frac{\partial^2}{\partial S^2} V_t^1} = \frac{\Gamma_t}{\Gamma_t^1}.$$

- Thus, we obtain a new option which is a linear combination of V and V^1. We then Delta-hedge this new option:

$$\Delta = \frac{\partial}{\partial S} V_t - \Delta^1 \frac{\partial}{\partial S} V_t^1$$

- As before, this is only an instantaneous hedge, as Δ and Γ will change through time.

Example:

- Suppose $S = 100, \sigma = 0.3, T = 1, r = q = 0$.
- We wish to hedge a short position in a put option with $K = 100$, using the stock and a call option with $K = 80$.
- We have the initial prices, Δs and Γs,

	Price	Δ	Γ
Stock	100	1	0
Put	11.9235	-0.4404	0.01315
Call	23.5344	0.8143	0.0089

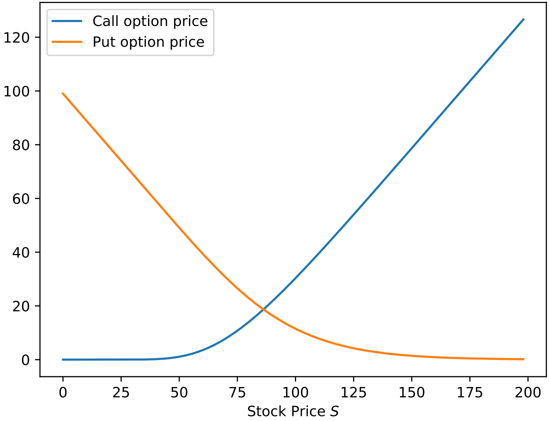
Call option price
Put option price
120
100
80
60
40
20
0
0
25
50
75
100
125
150
175
200
Stock Price S

- To get a Γ-neutral portfolio, as we are short one put option, we need to purchase

$$\Delta_1 = \frac{0.01315}{0.0089} = 1.4743$$

call options.

- To get a Δ-neutral portfolio, we then need to purchase

$$\Delta = -0.4404 - 1.4743 \times 0.8143 = -1.6409$$

stocks (i.e. a short position).

- These trades give us 141.3170 in cash, which we invest.

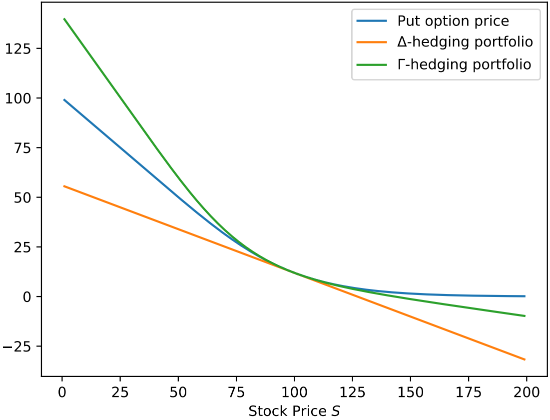

Values of put option, Δ-hedging and Γ-hedging portfolios

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD
Mathematical Institute

Greeks

7.3 Vega (ν ??)

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

- Volatility σ is the only parameter in the Black-Scholes model that is **not directly observed** in the market.
- It is, therefore, important to be able to evaluate the dependence of option price on volatility.

Definition

Vega (commonly written ν), is the sensitivity of the option price to changes in the volatility σ.

$$\nu = \frac{\partial}{\partial \sigma} V$$

- In the BS model, σ is constant, so hedging with respect to changes in σ doesn't make sense.
- However, one can ask: *what if my estimate of σ is wrong*?
- To estimate how far off, in this case, the computed option price is from the "true" price, we need to find Vega.
- For a call option, we have

$$\nu^{call} = \sqrt{\frac{T-t}{2\pi}} S e^{-\frac{d_+^2}{2}}.$$

  And it tends to zero as $t \to T$, since the payoff is independent of σ.
- It is the same for a put, due to put-call parity.

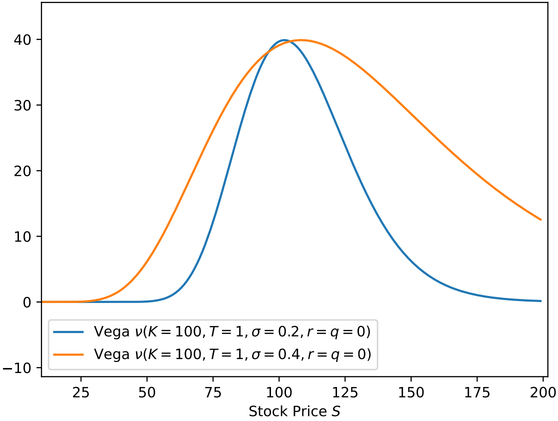
40
30
20
10
0
−10
25
50
75
100
125
150
175
200
Stock Price S
Vega $\nu(K = 100, T = 1, \sigma = 0.2, r = q = 0)$
Vega $\nu(K = 100, T = 1, \sigma = 0.4, r = q = 0)$

- **Vega-hedging** can be defined in the same way as Gamma-hedging, however, its purpose is different:
  - rather than reducing the discretization error, it is meant to reduce the **model error** – a misspecification of σ.
- Given an additional derivative with price V^1, the Vega-hedge of a short position in the original option prescribes to hold Δ^1 units of V^1 and Δ units of S.

- In order to make the portfolio instantaneously **Vega-neutral**, we need

$$-\frac{\partial}{\partial \sigma} V_t + \Delta^1 \frac{\partial}{\partial \sigma} V_t^1 = 0 \qquad \Rightarrow \qquad \Delta^1 = \frac{\frac{\partial}{\partial \sigma} V_t}{\frac{\partial}{\partial \sigma} V_t^1} = \frac{\nu_t}{\nu_t^1}$$

- As before, Δ is determined as the corresponding S-derivative of the portfolio of options $V - \Delta^1 V^1$, assuming Δ^1 is fixed:

$$\Delta = \frac{\partial}{\partial S} V_t - \Delta^1 \frac{\partial}{\partial S} V_t^1$$

- Of course, in order to keep the portfolio Vega-neutral at the next moment in time $t + \delta t$, the value of Δ^1 (as well as Δ) will need to be changed at that time.

- Vega hedging makes sense if one believes that the "true" volatility is constant, but we may be mistaken about its true value.
- However, sometimes, you may see a description of Vega-hedging as "hedging the non-constant volatility".
- If the volatility is believed to be changing dynamically, thenmore complicated **stochastic volatility models** have to be used
- Using a **constant volatility model**, to design a **hedge against stochastic volatility** is, clearly, **self-contradictory**. It may sometimes be used in practice, if other options are too hard to implement, however, then, one has to be very careful and make sure that the side effects of such hedging do not overweigh the positive impact.

- There is a **universal relation between Vega and Gamma** which holds for **all European options**, because their prices are **functions of time and the value of the underlying** and these functions **satisfy the BSPDE**.
- Consider the BSPDE

$$\mathcal{L}_{BS}V = \frac{\partial}{\partial t}V + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2}{\partial S^2}V + rS\frac{\partial}{\partial S}V - rV = 0$$

and differentiate it with respect to σ.

$$\frac{\partial}{\partial \sigma}V = \nu, \qquad \frac{\partial^2}{\partial \sigma \partial t}V = \frac{\partial}{\partial t}\nu, \qquad \cdots$$

- As a result, we obtain

$$\begin{aligned}\mathcal{L}_{BS}\nu &= \frac{\partial}{\partial t}\nu + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2}{\partial S^2}\nu + rS\frac{\partial}{\partial S}\nu - r\nu = -\sigma S^2 \frac{\partial^2}{\partial S^2} V \\ &= -\sigma S^2 \Gamma,\end{aligned}$$

  with $\nu(S, T) = 0$.
- In PDE language, the above equation means that "$-\sigma S^2\Gamma$" is a *source* for ν.
- Notice also that $\sigma S^2 \Gamma = \sigma S^2 \frac{\partial^2}{\partial S^2} V$ satisfies the BSPDE (recall homework exercise).

- Therefore, it is easy to check that

$$\nu(S,t) = (T-t)\sigma S^2 \Gamma(S,t)$$

satisfies the desired PDE:

$$\begin{aligned}\mathcal{L}_{BS}\left[(T-t)\sigma S^2\Gamma(S,t)\right] &= -\sigma S^2\Gamma + (T-t)\mathcal{L}_{BS}\left[\sigma S^2\Gamma(S,t)\right] \\ &= -\sigma S^2\Gamma\end{aligned}$$

- This is a useful trick, and a good example of how the **PDE techniques** may help in establishing certain **non-trivial** relations between various quantities in mathematical finance.

- From the fact that

$$\nu(S,t) = (T-t)\sigma S^2 \Gamma(S,t),$$

  we see that for European options, Γ-hedging and Vega-hedging amount to the same thing.
- The motivations however, are quite different!
- This relationship does not necessarily hold for other types of options.

B8.3 Mathematical Models of Financial Derivatives

Greeks

7.4 Other Greeks

Samuel Cohen

Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Oxford Mathematics

There are also sensitivities to other parameters

- Theta $\Theta = \partial V/\partial t$ is sensitivity with respect to t, and measures the maturity sensitivity of our portfolio.
- Rho $\rho = \partial V/\partial r$ is sensitivity with respect to the interest rate r, and measures the potential impact of interest rate changes on the value of the portfolio (more important over the long-term).
- Epsilon $\epsilon = \partial V/\partial q$ is the sensitivity with respect to the dividend yield

You can have higher order sensitivities:

- Vanna $\partial\Delta/\partial\sigma = \partial^2 V/\partial\sigma\partial S$ is the sensitivity of Δ with respect to changes in volatility σ. This is a measure of model dependence of the Delta-hedging strategy itself.
- Charm $\partial\Delta/\partial t = \partial\Theta/\partial S = \partial^2 V/\partial t\partial S$ is the sensitivity of Δ to time
- Vomma $\partial\nu/\partial\sigma = \partial^2 V/\partial\sigma\partial\sigma$ is the second order sensitivity to the volatility.
- Veta $-\partial\nu/\partial t$
- Vera $\partial\rho/\partial\sigma = \partial^2 V/\partial\sigma\partial r$
- Speed $\partial\Gamma/\partial S = \partial^3 V/\partial S^3$
- Zomma $\partial\Gamma/\partial\sigma$
- Color $-\partial\Gamma/\partial t$

- Given enough trading instruments (assets, options), we can cancel the higher order sensitivities as well.
- However, decreasing the risk associated with a wrong choice of parameters can increase model risk: our family of models (the Black–Scholes models, parameterized by r and σ) is not the true model!
- We also may have worse performance for larger moves (as we are using local sensitivities)
- Therefore, one should find an optimal trade-off, and shouldn't go too far with "matching the Greeks".

Timeo Danaos et dona ferentes (I fear the Greeks, even those bearing gifts) —Virgil, Aeneid

B8.3 Mathematical Models of Financial Derivatives

Mathematical Institute

Dividends

8.1 Discrete Dividends

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- So far we have only considered the case where the stock pays a continuous dividend of $qS_t dt$
- This is not a particularly convincing model for dividends over short–medium horizons
- Over long horizons (ie. decades) it is often a reasonable model, but most equity options have much shorter horizons.
- We need to handle the discrete nature of dividends in practice.

- Suppose that a share pays a deterministic dividend D at time t_D.
- If both D and t_D are known in advance we must have

$$S_{t_D^-} = S_{t_D^+} + D \quad \Longleftrightarrow \quad S_{t_D^+} = S_{t_D^-} - D$$

otherwise there is an arbitrage opportunity.
- Dividends are generally announced well in advance (say 1 month), so this can be a useful model.

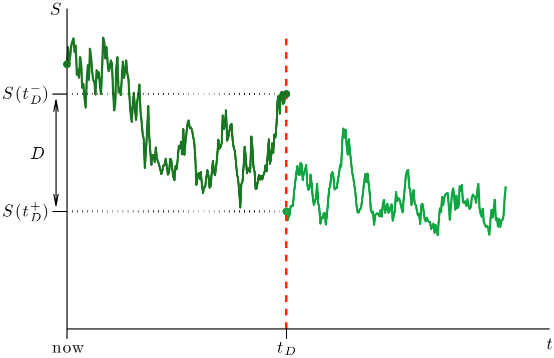

A jump in share price across a discrete dividend date.

- If we have an option on this share then we *don't* get the dividend and so we must have the jump condition

$$V\left(S_{t_D^-}, t_D^-\right) = V\left(S_{t_D^+}, t_D^+\right) = V\left(S_{t_D^-} - D, t_D^+\right).$$

- As this is true for any $S_{t_d^-}$ and we solve the Black-Scholes equation backwards in time, we generally write this *jump condition* as

$$V(S, t_D^-) = V(S - D, t_{D^+}).$$

The strategy is to

- solve the Black-Scholes equation back from expiry, T, until the dividend date t_D^+, then
- apply the jump condition to find $V(S, t_D^-)$ and then
- solve the Black-Scholes equation backwards from t_D^- to the present time, using $V(S, t_D^-)$ as a “payoff” at t_D^-.

Note that D can be a function of S and t. Indeed, if we want the share price to remain positive, it must be.

Modelling discrete dividend payments for a share price that follows geometric Brownian motion is problematic to this day.

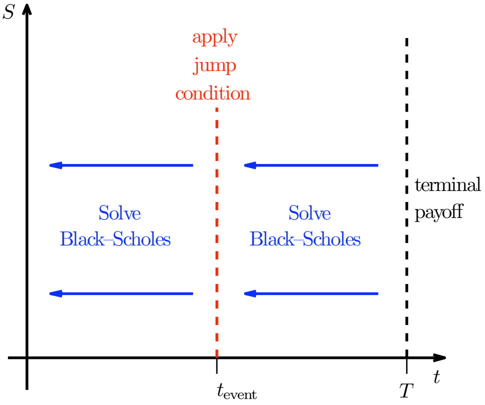

S
apply
jump
condition
terminal
payoff
Solve
Black–Scholes
Solve
Black–Scholes
t_{event}
T
t

If we assume a discrete dividend of the form

$$D = d_y \, S_{t_d^-} \, ,$$

where the discrete dividend yield $d_y < 1$, i.e., the dividend is proportional to the share price immediately before the dividend is paid then we find that

$$S_{t_d^-} = S_{t_d^+} + d_y \, S_{t_d^-} \quad \Longleftrightarrow \quad S_{t_d^+} = (1 - d_y) \, S_{t_d^-}$$

and the jump condition for the option becomes

$$V(S, t_d^-) = V\big((1 - d_y)S, t_d^+\big).$$

- We can then use the fact that if $V(S,t)$ is a solution of the Black-Scholes equation then so too is $V(\lambda S,t)$.
- With $\lambda = (1-d_y)$ in this case, we see that the solution for $t < t_d$ is simply

$$V\big((1-d_y)S,t\big),$$

as it is a solution of the Black-Scholes equation and obviously satisfies the "payoff" condition at t_d^-.
- Let's make this explicit for a European Call.

- Let $C_{\mathrm{v}}(S,t)$ be the Black–Scholes price function for a vanilla call, i.e.

$$C_{\mathrm{v}}(S,t) = S\,\mathrm{N}(d_+) - K\,e^{-r(T-t)}\,\mathrm{N}(d_-),$$
$$d_\pm = \frac{\log(S/K) + (r \pm \frac{1}{2}\sigma^2)(T-t)}{\sqrt{\sigma^2(T-t)}.}$$

- Suppose the stock pays a discrete dividend yield of d_y at time $0 < t_d < T$.
- Let $C(S,t)$ be the price of the option on this stock. Then for $t_d < t < T$ we have

$$C(S,t) = C_{\mathrm{v}}(S,t).$$

- Across the dividend date t_d, we apply the jump condition to get
$$C(S, t_d^-) = C_{\mathrm{v}}\big((1 - d_y)S, t_d\big).$$
- As $1 - d_y > 0$ is a constant, the function $C_{\mathrm{v}}\big((1 - d_y)S, t_d\big)$ is itself a solution of the Black-Scholes equation.
- So for all $t < t_d$ we have
$$C(S, t) = C_{\mathrm{v}}\big((1 - d_y)S, t\big).$$

Jump condition across a discrete dividend yield date

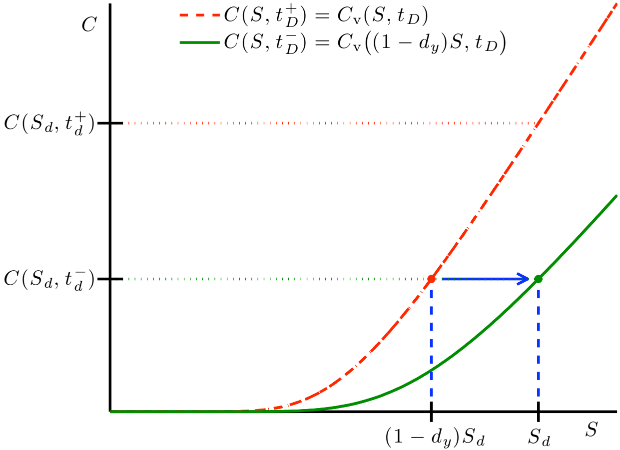

- The same reasoning shows that if there are n discrete dividend yields at times

$$t < t_1 < t_2 < \cdots < t_n < T$$

between now and expiry with dividend yields

$$d_1,\ d_2,\ \ldots\ ,\ d_n,$$

where each $d_k < 1$, then

$$C(S,t) = C\big(\alpha_n S, t\big), \quad \text{where} \quad \alpha_n = \prod_{k=1}^{n}(1 - d_k).$$

- Clearly this result generalises to any European option, regardless of the its payoff.

B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

American Options

9.1 Basic observations

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

- An American option is an option which can be exercised at any time between being initiated and expiring (inclusive).
- The key observation is that it is the holder (rather than the writer) of the option who decides whether the option should be exercised.
- Most traded options are of this type in practice.

Suppose the option has a payoff $P_o(S, t)$, which may depend on t. By no-arbitrage, it is clear that

- The American option cannot be worth less than $P_o(S_t, t)$, because the option can be exercised at any time $0 \le t \le T$.
- The American option can't be worth less than an otherwise equivalent European option (with payoff $P_o(S_T, T)$).
- If the European option is always worth at least $P_o(S_t, t)$, then it is no worse to hold the option to expiry than to exercise it, so the price of the European and American options will agree.

Check you can construct an arbitrage if any of these are violated!

- For a call option, if $q = 0$ we know that the equivalent European call has value

$$C(S_t, t) = e^{-r(T-t)}\mathbb{E}_{\mathbb{Q}}[(S_T - K)^+ | \mathcal{F}_t]$$

- As $x \mapsto (x-K)^+$ is convex, by Jensen's inequality we know

$$C(S_t, t) \geq e^{-r(T-t)}\Big(\mathbb{E}_{\mathbb{Q}}[S_T|\mathcal{F}_t] - K\Big)^+ = (S_t - e^{-r(T-t)}K)^+.$$

- Therefore, if $r \geq 0$, we have

$$C(S_t, t) \geq (S_t - K)^+$$

  so the European option is always worth more than its exercise value.
- Hence European and American Call options agree (for non-dividend paying assets).

- Conversely, if $r > 0$ then for a European put we have

$$\lim_{S\to 0} P(S,t) = K\, e^{-r(T-t)} < K.$$

- Since the European put price is differentiable, it is also continuous. Therefore, prior to expiry, a European put is less valuable than its payoff, for small enough S.
- As an American put can't be less valuable than the payoff, the values of American and European puts must be different.
- As they both have the same payoff, $(K-S)^+$, the American put can't satisfy the Black–Scholes equation for all $S > 0$.

B8.3 Mathematical Models of Financial Derivatives

Mathematical Institute

American Options

9.2 Linear complementarity formulation

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- There are a number of ways of formulating the American option problem. One is the *linear complementarity formulation*, which we give here.
- Let $V(S,t)$ be the value (function) of the option and $P_0(S,t)$ be the payoff (function).
- No arbitrage implies that

$$V(S,t) \geq P_o(S,t), \quad S > 0, \ t \leq T.$$

- Go back to the derivation of the Black-Scholes pricing equation so at any time we hold one *long* position in the American option, V, and Δ_t short positions in the underlying asset.

- As for the European option, the market value is

$$M_t = V(S_t, t) - \Delta_t \, S_t$$

and the change in the value of the portfolio is

$$d\Pi_t = dV_t - \Delta_t \, dS_t.$$

- Using Itô's lemma we get

$$d\Pi = \left(\frac{\partial V}{\partial t}(S_t, t) + \tfrac{1}{2}\sigma^2 \, S_t^2 \, \frac{\partial^2 V}{\partial S^2}(S_t, t) \right) dt + \left(\frac{\partial V}{\partial S}(S_t, t) - \Delta_t \right) dS_t$$

and so taking $\Delta_t = (\partial V / \partial S)(S_t, t)$ makes the change in portfolio value (instantaneously) risk-free.

- For the European option we then argued that $d\Pi_t > r\,M_t\,dt$ and $d\Pi_t < r\,M_t\,dt$ both represented arbitrage opportunities and hence $d\Pi_t = r\,M_t\,dt$, which gives the Black-Scholes equation.
- For the American option it is still true that $d\Pi_t > r\,M_t\,dt$ gives a clear arbitrage: borrow the price of the portfolio, M_t, set up the portfolio with the correct value of Δ_t. At time $t + dt$ the portfolio's risk-free value is $M_t + d\Pi_t$ which is greater than $(1 + r\,dt)\,M_t$ than you owe.

- Therefore we must have $d\Pi_t \le r\, M_t\, dt$ which is equivalent to the partial differential *inequality*

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2\, S^2\, \frac{\partial^2 V}{\partial S^2} + r\, S\, \frac{\partial V}{\partial S} - r\, V \le 0.$$

- The problem comes with showing that $d\Pi_t < r\, M_t\, dt$ is an arbitrage if the option is American.
- This is because it involves short-selling or writing the option and, unlike a European option, an American option can be exercised at *any* time, not just at expiry.
- Indeed, the only reason for exercising an American option before expiry is that the return on the delta-hedged portfolio is less than the return on the bank account.

- Now suppose that $V(S_t, t) > P_{\rm o}(S_t, t)$.
- Then it would be absurd to exercise the American option early as you could sell it for more.
- You could also short-sell it knowing that it wouldn't be exercised immediately.
- Therefore you *can* make an arbitrage if $d\Pi_t < r\,M_t\,dt$ and $V(S_t, t) > P_{\rm o}(S_t, t)$ and so

$$V(S, t) > P_o(S, t) \implies \mathcal{L}_{\rm bs}(V) = 0,$$

where $\mathcal{L}_{\rm bs}(V)$ is the Black-Scholes operator

$$\mathcal{L}_{\rm bs}(V) = \frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2\,S^2\,\frac{\partial^2 V}{\partial S^2} + r\,S\,\frac{\partial V}{\partial S} - r\,V.$$

- Now if $\mathcal{L}_{\mathrm{bs}}(V) < 0$ we can't have $V > P_{\mathrm{o}}$ for the reason immediately above.
- We can't have $V < P_{\mathrm{o}}$ as this represents an arbitrage and so the only possibility is $V(S,t) = P_{\mathrm{o}}(S,t)$.
- Thus
$$\mathcal{L}_{\mathrm{bs}}(V) < 0 \implies V(S,t) = P_{\mathrm{o}}(S,t).$$
In total we can write this as the *linear complementarity problem*
$$\mathcal{L}_{\mathrm{bs}}(V) \le 0, \quad V(S,t) \ge P_{\mathrm{o}}(S,t),$$
$$\big(V(S,t) - P_{\mathrm{o}}(S,t) \big)\, \mathcal{L}_{\mathrm{bs}}(V) = 0.$$
- At expiry we have $V(S,T) = P_{\mathrm{o}}(S,T)$. No arbitrage implies that $V(S,t)$ is continuous in S for $t < T$.

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD

Mathematical Institute

American Options

9.3 Smooth pasting

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- We need another condition to uniquely determine $V(S,t)$ and it is that the *holder chooses the early exercise strategy in order to maximize the option's value.*
- The fact that the holder chooses the early exercise strategy to maximize the option's value is *often* equivalent to the *smooth pasting conditions*, there is some $\hat{S}(t)$ such that

$$V(\hat{S}(t),t) = P_{\mathrm{o}}(\hat{S}(t),t), \quad \frac{\partial V}{\partial S}(\hat{S}(t),t) = \frac{\partial P_{\mathrm{o}}}{\partial S}(\hat{S}(t),t)$$

  and on one side of $\hat{S}(t)$ we have $\mathcal{L}_{\mathrm{bs}}(V) = 0$ and on the other we have $V(S,t) = P_{\mathrm{o}}(S,t)$.
- The function $\hat{S}(t)$ is called the *optimal exercise* boundary. It is part of the problem to find $\hat{S}(t)$, hence the two conditions applied there.

- The first smooth pasting condition (continuity) can be seen by no arbitrage. The second is more difficult.
- Smooth pasting is *not* universally true.
- There are American options for which it is always true, there are some American options for which it is always false and there are other American options where it is sometimes true and sometimes false.
- For American puts and calls (calls with positive dividends $q > 0$) it is always true.

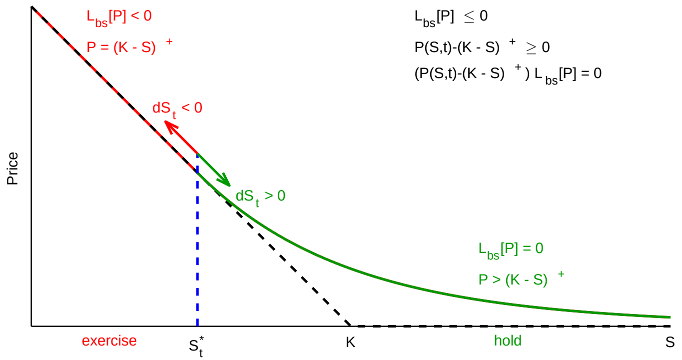

L_bs[P] < 0
P = (K - S) +
L_bs[P] ≤ 0
P(S,t)-(K - S) + ≥ 0
(P(S,t)-(K - S) +) L bs[P] = 0
dS t < 0
dS t > 0
Price
L_bs[P] = 0
P > (K - S) +
exercise
S*t
K
hold
S

- Consider an American put option (for simplicity, with no dividends) where the share price at time t is equal to the optimal exercise price, S_t^*.
- If $dS_t < 0$, so the share price goes down, then the put's value equals the payoff (and the option is exercised).
- If $dS_t > 0$, so the share price goes up, then the put's value is above the payoff (and the option is held). Thus we have

$$P(S_t^* + dS_t, t + dt) = \begin{cases} K - S^* - dS_t & \text{if } dS_t < 0, \\ P(S^* + dS_t, t + dt) & \text{if } dS_t > 0. \end{cases}$$

- Now assuming that S_t follows a geometric Brownian motion,

$$\frac{dS_t}{S_t} = \mu\, dt + \sigma\, dW_t,$$

where $dW_t \sim N(0, dt)$, it follows that $dW_t = \sqrt{dt}\, Z$ where $Z \sim N(0, 1)$.

- Thus $dW_t = \mathcal{O}(\sqrt{dt})$ and since dt is infinitesimally small we have $dW_t \gg dt$.
- This in turn implies that

$$dS_t = \sigma\, S_t^*\, dW_t + \mathcal{O}(dt)$$

(recall that $S_t = S_t^*$)

- Substituting, we have the expansion

$$P(S_t^* + dS_t, t + dt) = P(S^*, t) + \sigma\, S_t^* \, \frac{\partial P}{\partial S}(S_t^*, t)\, dW_t + \mathcal{O}(dt).$$

- Thus, with $P_t = P(S_t^*, t)$, we have

$$dP_t^* = \begin{cases} -\sigma\, S_t^*\, dW_t & \text{if } dW_t < 0, \\ \sigma\, S_t^* \, \dfrac{\partial P}{\partial S}(S^*, t)\, dW_t & \text{if } dW_t > 0. \end{cases}$$

- Consider a portfolio with a long put and a long share, $\Pi_t = P(S_t^*, t) + S_t^*$, (recall again that $S_t = S_t^*$).
- From the above we see that

$$d\Pi_t = \begin{cases} 0 & \text{if } dW_t < 0, \\ \sigma S_t^* \left(\frac{\partial P}{\partial S}(S^*, t) + 1 \right) & \text{if } dW_t > 0. \end{cases}$$

- Now suppose that

$$\frac{\partial P}{\partial S}(S_t^*, t) + 1 > 0 \quad \text{or} \quad \frac{\partial P}{\partial S}(S_t^*, t) + 1 < 0.$$

Both of these cases lead to an arbitrage in which $d\Pi_t$ is either non-negative with a non-zero probability of being strictly positive (the first case) or non-positive with a non-zero probability of being strictly negative (the second case).

Therefore, to avoid an arbitrage we must have

$$\frac{\partial P}{\partial S}(S_t^*, t) = -1,$$

which is the (second) smooth pasting condition.

B8.3 Mathematical Models of Financial Derivatives

UNIVERSITY OF OXFORD
Mathematical Institute

American Options

9.4 Perpetual American options

Samuel Cohen
Hilary Term 2021

- We consider the case where the option never expires, $T \to \infty$.
- In this case there is no difference in the option pricing problem between the spot/time points (S, t_1) and (S, t_2) when $t_1 \neq t_2$, so, we can assume that $V = V(S)$.
- In this case, provided the option hasn't already been exercised, it satisfies the ordinary differential equation

$$\frac{1}{2}\sigma^2 S^2 V''(S) + (r - q) S V'(S) - r V(S) = 0.$$

This equation is sometimes called an Euler equation.

- One way to solve this equation is to look for solutions in terms of the eigenfunctions of $S\,\partial/\partial S$,

$$V(S) = S^m, \quad S\,V'(S) = m\,S^m, \quad S^2\,V''(S) = m(m-1)\,S^m.$$

- With this choice, we see that m must satisfy the quadratic equation

$$\tfrac{1}{2}\sigma^2\,m(m-1) + (r-q)\,m - r = 0.$$

- If we assume $\sigma > 0$, $r > 0$ and $q \geq 0$ and set

$$p(m) = \tfrac{1}{2}\sigma^2\,m(m-1) + (r-q)\,m - r$$

then $p(m)$ has apositive coefficient for the quadratic term m^2 and at the points $m = 0$ and $m = 1$ we have $p(0) = -r < 0$ and $p(1) = -q \leq 0$.

- From these facts it follows that if $m^{\pm}$ are the roots of the quadratic then
$$m^{-} < 0, \quad m^{+} \geq 1.$$
- Thus the general solution is
$$V(S) = A\,S^{m^{-}} + B\,S^{m^{+}}, \quad m^{-} < 0,\; m^{+} \geq 1.$$
- In particular, if we know $V(\infty) = 0$, then $B = 0$, and $V(S) = AS^{m^{-}}$ for some A.

- With $y \geq 0$ and $r > 0$ we find that the problem for the American put is

$$0 = \tfrac{1}{2}\sigma^2 S^2 P''(S) + (r-q)\, S\, P'(S) - r\, P(S), \quad S > \hat{S},$$
$$P(\hat{S}) = K - \hat{S}, \qquad P'(\hat{S}) = -1, \qquad P(\infty) = 0.$$

- The two conditions $P(\hat{S}) = K - \hat{S}$ and $P(\infty) = 0$ give

$$P(S) = (K - \hat{S})\left(\frac{S}{\hat{S}}\right)^{m^-}.$$

- The remaining boundary condition, $P'(\hat{S}) = -1$, then gives

$$0 < \hat{S} = \frac{m^- K}{m^- - 1} < K$$

since $m^- < 0$ (which implies $m^- - 1 < m^- < 0$).

- For the perpetual Amerian put, we can also find the value directly by maximizing over exercise values.
- Again assume that $q \geq 0$ and $r > 0$.
- Choose an arbitrary $0 < \bar{S} < K$ and exercise as soon as S falls to $\bar{S}$.
- Then

$$\tfrac{1}{2}\sigma^2 S^2 P''(S) + (r - q)\, S\, P'(S) - r\, P(S) = 0, \quad S > \bar{S},$$
$$P(\bar{S}) = K - \bar{S}, \quad P(\infty) = 0.$$

- As above, the solution is

$$P(S; \bar{S}) = (K - \bar{S}) \left(\frac{S}{\bar{S}} \right)^{m^-},$$

where $m^- < 0$.

- Now (formally) set

$$\frac{\partial P}{\partial \bar{S}}(S;\bar{S}) = \left(\frac{S}{\bar{S}}\right)^{m^-} \left(-1 - m^- \frac{K - \bar{S}}{\bar{S}}\right) = 0.$$

- This gives

$$-1 - m^- \left(\frac{K - \bar{S}}{\bar{S}}\right) = 0.$$

- In turn, this implies the optimal value of $\bar{S}$, $\hat{S}$, is

$$0 < \hat{S} = \frac{m^- K}{m^- - 1} < K.$$

- This is the same as the smooth pasting version.

- For an example without smooth pasting, consider the perpetual American digital put option.
- Assume $r > 0$ and $q \geq 0$.
- Given $r > 0$ and the payoff is constant for $S < K$, it is clear that the optimal strategy is to exercise the first time $S = \hat{S} = K$.
- For $S > K$, the problem for the perpetual American digital put option is

$$0 = \tfrac{1}{2}\sigma^2 S^2 P_d''(S) + (r-q)\, S\, P_d'(S) - r\, P_d(S), \quad 0 < K < S,$$
$$P_d(K) = 1, \quad P_d(\infty) = 0.$$

- The solution is

$$P_d(S) = \left(\frac{S}{K}\right)^{m^-}.$$

- It is *not* possible to make $P_d'(K)$ continuous at $S = K$.
- Thus, the second smooth pasting condition (involving $P_d'(\hat{S})$) *does not* apply in this case!

Note that you *cannot* adapt the smooth pasting argument used for the American put option above, so that it works for an American digital put (perpetual or with finite expiry date T).

B8.3 Mathematical Models of Financial Derivatives

Mathematical Institute

Exotic Options

10.1 Forward start options

Samuel Cohen

Hilary Term 2021

Oxford Mathematics

- *Exotic option* is a catch-all term that includes options which are not generally traded on markets, or not widely traded.
- They often occur embedded in other more complex financial products, as structured products or as over-the-counter options created for clients with special financial needs.
- For some asset classes (e.g. electricity, interest rates), these may be the main way the commodity is traded.

- We begin with a simple case – a forward start option
- These involve a payoff with a strike, e.g., a call, where the strike is determined by the share price at some time T_1, where $0 < T_1 < T_2$.
- For example, we could have $K = S_{T_1}$.
- The trick to pricing them is to note that for $T_1 < t \leq T_2$ the strike K is known.

- For a call, this means that for $T_1 < t \le T_2$ we can write

$$C_{\mathrm{fs}}(S,t) = S\,e^{-q(T_2-t)}\,\mathrm{N}(d_+) - K\,e^{-r(T_2-t)}\,\mathrm{N}(d_-), \quad K = S_{T_1}.$$

- At time T_1, $K = S$ by definition and so

$$C_{\mathrm{fs}}(S,T_1) = A(r,q,\sigma,T_1,T_2)\,S$$

where $A(r,q,\sigma,T_1,T_2)$ is independent of S and t and is given by

$$A(r,q,\sigma,T_1,T_2) = e^{-q(T_2-T_1)}\,\mathrm{N}(d_+^*) - e^{-r(T_2-T_1)}\,\mathrm{N}(d_-^*),$$

$$d_\pm^* = \frac{(r-q\pm\frac{1}{2}\sigma^2)(T_2-T_1)}{\sqrt{\sigma^2(T_2-T_1)}}.$$

- As the price at time T_1 is simply a (known) multiple of S, a no arbitrage argument (or the Black–Scholes equation) shows that for $t \le T_1$ the price must be given by

$$C_{\mathrm{fs}}(S,t) = A(r,q,\sigma,T_1,T_2)\, S\, e^{-q(T_1-t)}.$$

- More complicated forward start values will involve solving the Black–Scholes equation for $t \le T_1$.

B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

Exotic Options

10.2 Barrier options

Samuel Cohen
Hilary Term 2021

- A down and out barrier call option becomes worthless (colloquially refered to as "knocking out") if the share price falls to or below a barrier level, $B > 0$, at any time during the option's life.
- For simplicity, we take B to be a constant.
- If $S_t > B$ for all $t \in [0, T]$ then it has payoff $(S_T - K)^+$.
- The pricing problem, assuming that the option has not already knocked out, is

$$\begin{aligned} \mathcal{L}_{BS} C_{\mathrm{do}} &= \frac{\partial C_{\mathrm{do}}}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 C_{\mathrm{do}}}{\partial S^2} + (r-q)\, S\, \frac{\partial C_{\mathrm{do}}}{\partial S} - r\, C_{\mathrm{do}} \\ &= 0, \qquad S > B,\ \ t < T, \\ C_{\mathrm{do}}(S, T) &= (S-K)^+, \qquad S > B, \\ C_{\mathrm{do}}(B, t) &= 0, \qquad t \le T. \end{aligned}$$

We begin with the case $0 < B \leq K$.

- The trick here is to recall that if $V(S,t)$ is a solution of the Black-Scholes equation then so too is

$$\hat{V}(S,t) = (S/B)^{2\alpha}\, V\left(B^2/S, t\right),$$

where $2\alpha = 1 - 2(r-q)/\sigma^2$, and that

$$\hat{V}(B,t) = (B/B)^{2\alpha}\, V\left(B^2/B, t\right) = V\left(B, t\right).$$

- So if $C_{\mathrm{bs}}(S,t)$ is the price of a vanilla call option we see that

$$C_{\mathrm{do}}(S,t;B) = C_{\mathrm{bs}}(S,t) - (S/B)^{2\alpha} C_{\mathrm{bs}}\left(B^2/S, t\right)$$

is a solution of the Black-Scholes equation which satisfies $C_{\mathrm{do}}(B,t;B) = 0$.

- Then we notice that as $B < S$ and $B \leq K$ we have $B^2/S < B \leq K$ so that

$$C_{\text{bs}}\left(B^2/S, T\right) = \left(B^2/S - K\right)^+ = 0$$

which shows that for $S > B$

$$C_{\text{do}}(S, T; B) = C_{\text{bs}}(S, T) = (S - K)^+.$$

- Thus $C_{\text{do}}(S, t; B)$ satisfies the pricing problem and is the Black–Scholes value of the barrier option (before the Barrier is hit)

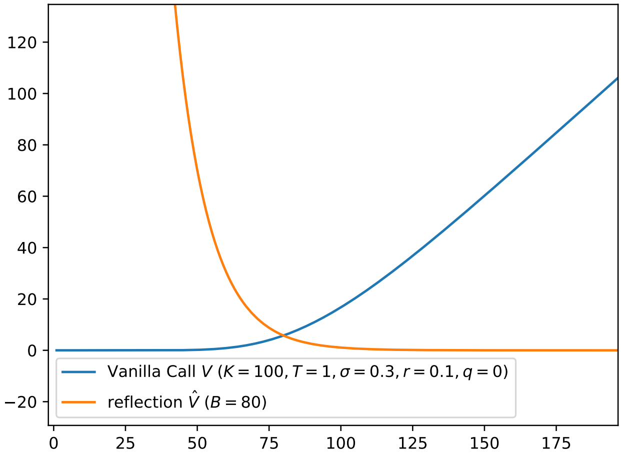
120
100
80
60
40
20
0
−20
0
25
50
75
100
125
150
175
Vanilla Call V ($K = 100, T = 1, \sigma = 0.3, r = 0.1, q = 0$)
reflection $\hat{V}$ ($B = 80$)

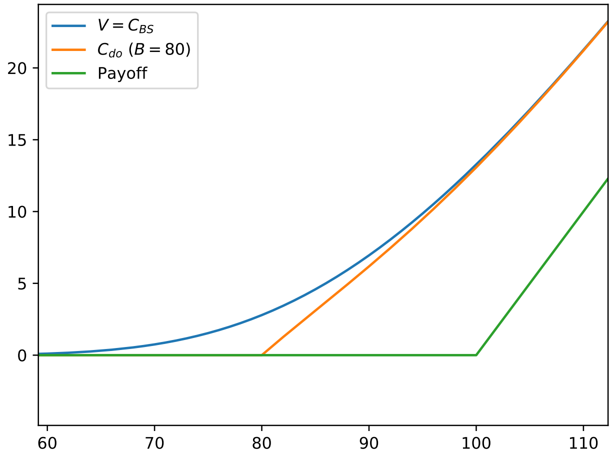
$V = C_{BS}$
C_{do} $(B = 80)$
Payoff
20
15
10
5
0
60
70
80
90
100
110

- We can now consider the more difficult case $B > K > 0$.
- In this case the trick above fails because we find that $C(B^2/S, T) \neq 0$ for all $S > B$.
- The way to deal with it is to truncate the payoff of the call so that it becomes equal to zero if $S \leq B$ but remains unchanged if $S > B$, i.e., replacing the vanilla call above with an option whose payoff is

$$V(S,T) = \begin{cases} 0 & \text{if } 0 < S \leq B, \\ S - K & \text{if } S > B > K. \end{cases}$$

- This payoff can be achieved by using a vanilla call with strike B plus $(B - K)$ digital calls, also with strike B.

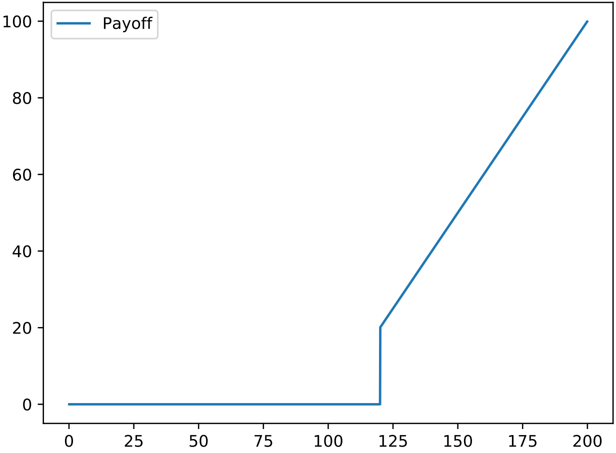
Payoff
100
80
60
40
20
0
0
25
50
75
100
125
150
175
200

- So instead of using the vanilla call price as above, we work with

$$V(S,t;B) = C_{\mathrm{bs}}(S,t;K=B) + (B-K)\, C_{\mathrm{d}}(S,t;K=B).$$

- The Black-Scholes price function is given by

$$C_{\mathrm{do}}(S,t;B) = V(S,t;B) - (S/B)^{2\alpha} V\left(B^2/S, t\right)$$

  since this satisfies the Black-Scholes equation and boundary condition at $S = B$ and if $S > B$ then $B^2/S < B$.
- So $V\left(B^2/S, T\right) = 0$, giving us the correct payoff at T.

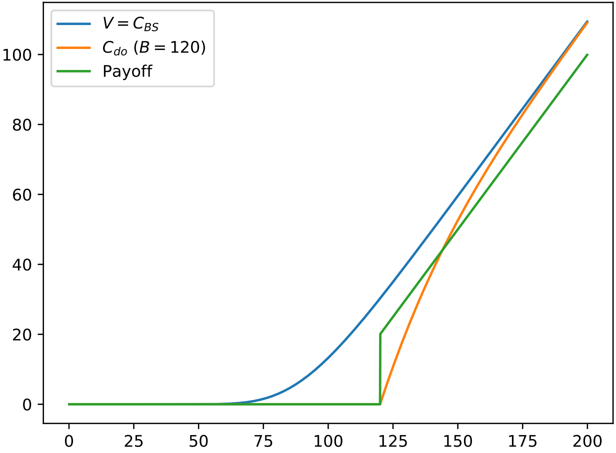
$V = C_{BS}$
C_{do} $(B = 120)$
Payoff
100
80
60
40
20
0
0
25
50
75
100
125
150
175
200

- Using our Down-and-Out formula, its easy to obtain the price of the 'down-and-in' option as well.
- This option remains worthless if the share price does *not* fall *below* the barrier $B > 0$ during the life of the option.
- If at some point during the life of the option we have $S_t < B$ then the option turns into a vanilla call with payoff $(S_T - K)^+$; this is often referred to as "knocking in".

- If we hold both a down-and-out and a down-and-in call option then we are guaranteed the payoff $(S_T - K)^+$ and so there is a down-and-in / down-and-out parity relation,

$$C_{\mathrm{do}}(S,t;B) + C_{\mathrm{di}}(S,t;B) = C_{\mathrm{bs}}(S,t)$$

  and hence

$$C_{\mathrm{di}}(S,t;B) = C_{\mathrm{bs}}(S,t) - C_{\mathrm{do}}(S,t;B).$$

- In the case that $B < K$ this simplifies to

$$C_{\mathrm{di}}(S,t;B) = (S/B)^{2\alpha}\, C_{\mathrm{bs}}\left(B^2/S,t\right).$$

- Note that these formula for C_{di} are only valid if $S \geq B$ for all time up to the present. As soon as $S < B$ the option turns into a vanilla call and remains so until expiry.

B8.3 Mathematical Models of Financial Derivatives

University of Oxford
Mathematical Institute

Exotic Options

10.3 Arithmetic Asian options

Samuel Cohen
Hilary Term 2021

Oxford Mathematics

Mathematical Institute

- Asian options are options which depend on the average share price over the life of the option.
- In practice, it is usually the arithmetic average which we can define using the running sum of the share price

$$R_t = \int_0^t S_u \, du, \quad A_T = R_T/T = \frac{1}{T}\int_0^T S_u \, du,$$

where A_T is the average price at T.
- The option's price is a function of S_t, R_t and t, $V_t = V(S_t, R_t, t)$ for some function $V(S, R, t)$.

If we note that

$$dR_t = S_t \, dt$$

and assume that $dW_t^2 = dt$ (which really means $d[W]_t = dt$) and perform a formal Taylor series expansion then we find

$$\begin{aligned} dV_t &= \left(\frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 \, S_t^2 \, \frac{\partial^2 V}{\partial S^2} \right) dt + \frac{\partial V}{\partial R} \, dR_t + \frac{\partial V}{\partial S} \, dS_t \\ &= \left(\frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 \, S_t^2 \, \frac{\partial^2 V}{\partial S^2} + S_t \, \frac{\partial V}{\partial R} \right) dt + \frac{\partial V}{\partial S} \, dS_t, \end{aligned}$$

where all partial derivatives are evaluated at (S_t, R_t, t).

Applying the usual hedging (or self-financing replication) argument(s) shows that to eliminate risk we must hold

$$\Delta_t = \Delta(S_t, R_t, t) = \frac{\partial V}{\partial S}(S_t, R_t, t)$$

shares at time t and that the pricing equation is

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r - q)\, S \frac{\partial V}{\partial S} + S \frac{\partial V}{\partial R} - r\, V = 0.$$

- This holds for all $S > 0$, $t < T$ and $R > 0$.
- As $S_u > 0$, $R_t = \int_0^t S_u \, du$ can take only positive values.

If the option is what is known as a floating-strike asian call, where the average plays the role of the strike, so the payoff is

$$V(S, R, T) = \left(S - R/T\right)^+,$$

then we can simplify the problem by working with the variables

$$x = R/S, \quad V(S, R, t) = S\,u(x, t),$$

i.e., by pricing relative to the share price rather than a unit of currency (in finance this is usually called *a change of numeraire*).

We find that

$$\frac{\partial V}{\partial t} = S\,\frac{\partial u}{\partial t}, \qquad S\,\frac{\partial V}{\partial R} = S\,\frac{\partial u}{\partial x},$$

$$S\,\frac{\partial V}{\partial S} = S\,\left(u - x\,\frac{\partial u}{\partial x}\right), \qquad S^2\,\frac{\partial^2 V}{\partial S^2} = S\,x^2\,\frac{\partial^2 u}{\partial x^2}.$$

Substituting these into the pricing PDE gives

$$\frac{\partial u}{\partial t} + \tfrac{1}{2}\sigma^2\,x^2\frac{\partial^2 u}{\partial x^2} + \big(1 + (q - r)x\big)\,\frac{\partial u}{\partial x} - q\,u = 0.$$

At expiry, $t = T$, we have $V(S, R, T) = (S - R/T)^+ = S\,u(x, T)$ and so

$$u(x, T) = (1 - x/T)^+.$$

The Feynman–Kac formula shows that the solution can be expressed as

$$u(x, t) = e^{-q(T-t)}\mathbb{E}\Big[\big(1 - x_T/T\big)^+ \,|\, x_t = x\Big],$$

where x_τ evolves as

$$dx_\tau = \big(1 + (q - r)\,x_\tau\big)\,d\tau + \sigma\,x_\tau\,dW_\tau,$$

for $\tau > t$, with $x_t = x$. Sadly this doesn't have a nice closed-form solution.

B8.3 Mathematical Models of Financial Derivatives

Exotic Options

10.4 Geometric Asian options

Samuel Cohen

Hilary Term 2021

UNIVERSITY OF OXFORD
Mathematical Institute

Oxford Mathematics

- An alternative type of Asian option is where the Geometric average is used, instead of the arithmetic average.
- This leads to a much simpler analysis, given we have a Geometric Brownian Motion as our fundamental model.
- A PDE approach is possible, but we will consider a probabilistic argument for the sake of variety.

- We will consider the simplest case, of option with payoff

$$G(S) = \left(\exp \frac{1}{T} \int_0^T \log(S_u) du - K \right)^+$$

- We ignore dividends ($q = 0$), and recall that the usual argument shows that the price must be given by

$$e^{-r(T-t)} \mathbb{E}_{\mathbb{Q}}[G(S) | \mathcal{F}_t]$$

where, under $\mathbb{Q}$, the stock follows the GBM

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

or equivalently for $t < u$

$$S_u = S_t \exp\left(\left(r - \frac{\sigma^2}{2} \right)(u - t) + \sigma(W_u - W_t) \right).$$

We now observe that

$$\log(S_u) = \log(S_t) + \left(r - \frac{\sigma^2}{2}\right)(u - t) + \sigma(W_u - W_t)$$

and hence

$$\begin{aligned}
&\int_t^T \log(S_u)du - (T-t)\log(S_t) \\
&= \left(r - \frac{\sigma^2}{2}\right)\int_t^T (u-t)du + \sigma\int_t^T (W_u - W_t)du \\
&= \left(r - \frac{\sigma^2}{2}\right)\frac{(T-t)^2}{2} + \sigma\int_t^T (W_u - W_t)du
\end{aligned}$$

We know that $\{W_u - W_t\}_{u \geq t}$ are mean-zero jointly Gaussian random variables, so

$$\mathbb{E} \int_t^T (W_u - W_t) du = 0$$

and, as $\text{cov}(W_t, W_s) = \min(t, s)$, (as in PS2 Q3)

$$\begin{aligned}
\mathbb{E}\left[\left(\int_t^T (W_u - W_t) du\right)^2\right] &= \int_t^T \int_t^T \mathbb{E}[(W_u - W_t)(W_s - W_t)] ds du \\
&= \int_t^T \int_t^T \min(u - t, s - t) ds du \\
&= \int_t^T \int_t^u s - t \, ds du + \int_t^T \int_u^T u - t \, ds du \\
&= (T - t)^3/3
\end{aligned}$$

From this, we conclude that, under the risk-neutral measure $\mathbb{Q}$,

$$\int_t^T (W_u - W_t)du \quad \sim \quad \mathcal{N}(0, (T-t)^3/3)$$

and in particular, given $\mathcal{F}_t$,

$$\frac{1}{T}\int_0^T \log(S_u)du \quad \sim \quad \mathcal{N}(\bar{\mu}, \bar{\sigma}^2)$$

where

$$\bar{\mu} = \frac{1}{T}\int_0^t \log(S_u)du + \frac{T-t}{T}\log(S_t) + \left(r - \frac{\sigma^2}{2}\right)\frac{(T-t)^2}{2T}$$

$$\bar{\sigma}^2 = \frac{(T-t)^3}{3T^2}$$

We know from the derivation of the Black–Scholes formula that if $X = \exp(a + bZ)$ for $Z \sim \mathcal{N}(0,1)$, then

$$\mathbb{E}[X\mathbf{1}_{X>K}] = \exp(a - b^2/2)\mathrm{N}\Big(\frac{a - \log(K)}{b} + b\Big)$$

$$\mathbb{E}[\mathbf{1}_{X>K}] = \mathrm{N}\Big(\frac{a - \log(K)}{b}\Big)$$

and so we can price our geometric Asian option,

$$e^{-r(T-t)}\mathbb{E}\Big[\Big(\exp\frac{1}{T}\int_0^T \log(S_u)du - K\Big)^+\Big|\mathcal{F}_t\Big]$$

$$= e^{-r(T-t)+\bar{\mu}-\bar{\sigma}^2/2}\mathrm{N}\Big(\frac{\bar{\mu} - \log(K)}{\bar{\sigma}} + \bar{\sigma}\Big) - e^{-r(T-t)}K\mathrm{N}\Big(\frac{\bar{\mu} - \log(K)}{\bar{\sigma}}\Big).$$