S&P 500 Binomial Option Pricing

March 2026

Group project. Collaborator names and student identifiers are omitted from this public copy.

Part A - Data Sets and Data Cleaning:

For the analysis one full year of data was used which spanned from the 13th of February 2025 to the 13th of February 2026. The dataset includes the following:

- Daily closing prices of the S&P 500 index
- Daily 6-month US Treasury bill discount rate (DTB6)
- Daily 3-month US Treasury bill discount rate (DTB3)

All the information was obtained from the Federal Reserve Economic Data (FRED) database.

Over the period, the S&P 500 contains 252 trading days. The daily log returns were calculated.

$$r_t = \ln \frac{S_t}{S_{t-1}} \tag{1}$$

As the first price does not generate a return due to the lag structure, to ensure the first observation also had an associated return, the closing price from the 12th of February 2025 was included when calculating returns. Therefore, the dataset contains 252 daily log observations which is fully aligned with the trading days in the sample period.

The calendar dates on which both the stock market as well as the bond market were closed were removed from the dataset to ensure that there was alignment across all the series. This ensured that each observation corresponded to an active trading day. However, there were three days, in which the bond market was open but there was no data for the stocks (Appendix A), due to short term interest rates typically exhibit strong persistence and adjust gradually, these missing values were forward filled using the most recent available observation for each maturity. This preserves a continuous daily risk-free rate series required for consistent binomial tree calculation, while being able to remain economically reasonable. This method is standard empirical practice in financial time series analysis when dealing with separately missing short rate data, as it avoids introducing artificial jumps.

Discount basis rates were converted into bond prices assuming a face value of 1. Bond Prices were calculated as follows:

$$P = 1 - \left(\frac{discount\ basis}{100}\right) \cdot T \tag{2}$$

The corresponding yields were then calculated and expressed in annualized % terms to then be further used in the binomial pricing models.

Part B - Data Analysis:

Part A

The estimated daily mean log-return over the one-year sample is seen to be 0.048%, which is annualized to 12.10%. The daily standard deviation (volatility) of the log-returns is 1.17% annualizing to 18.57%.

$$\overline{r}_{annual} = \overline{r}_{daily} \times 252 \tag{3}$$

$$\sigma_{annual} = \sigma_{daily} \times \sqrt{252} \tag{4}$$

The annualized volatility of 18.57% falls within the typical long-run S&P 500 range, indicating a period of moderate market risk. The positive annualized mean return of 12.10% reflects overall index growth during the year, indicating a broadly upward market trend over the sample period. These estimates further assist in the binomial model, as the annualized volatility determines how large the potential up (u) and down price (d) movements are in each period.

$$u = e^{\sigma\sqrt{T}} \quad \text{and} \quad d = \frac{1}{u} \tag{5}$$

Part B

Figure 1 : 30-day rolling volatility

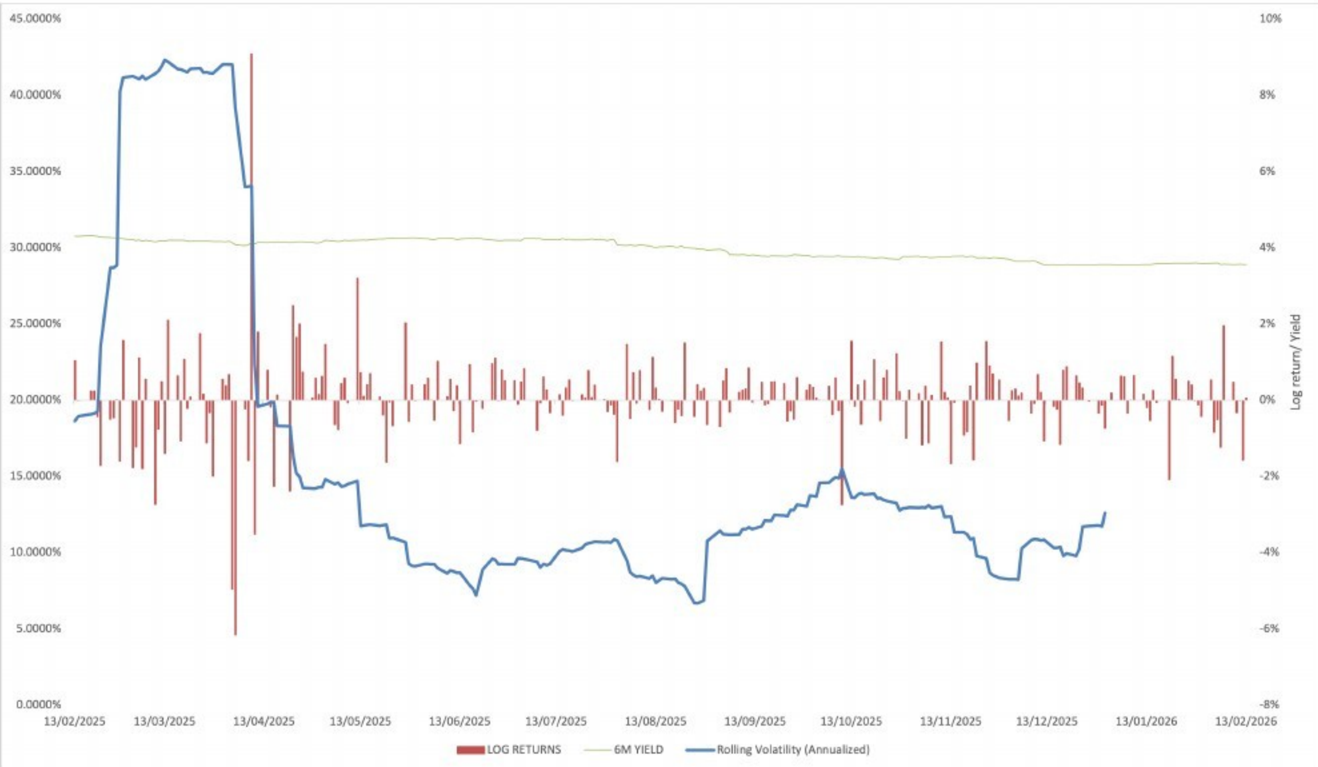

The 30-day rolling volatility was calculated using overlapping 30-day periods. The rolling volatility ranges from 6.6780% and 42.3192% over the sample.

The rolling volatility clearly indicates that market risk changes over time rather than remaining constant. A sharp increase in volatility is visible in late February and early March 2025, where it rises above 40%. This spike occurs at the same time as several negative daily log returns, including a -6.16% in early April. This indicates that large price movements, particularly declines tend to coincide with higher market uncertainty.

Following the spike, volatility gradually falls and stabilizes between roughly 8% and 15% for most of the remaining sample. This pattern represents well-known behaviour of financial markets where volatility clusters in certain periods rather than moving randomly.

Part C

Over the one-year sample of 252 trading days, there are clear patterns that can be observed in returns, volatility and interest rates. 109 days gave negative returns meaning the market moved down almost as often as it moved up. Yet, the S&P 500 produced a total return of 11.79%, corresponding to an annualized mean log-return of 12.18%. This contrast highlights that the S&P 500 grew through long-term compounding rather than steady daily gains. The average daily return was 0.048%, while daily volatility was 1.17%, showing that short-run price movements were dominated by fluctuations rather than trend.

Extreme returns were also notable. The maximum daily log-return reached 9.089%, while the minimum was 6.161%. Relative to a daily standard deviation of 1.17%, these observations were many standard deviations away from the mean and were extremely unlikely under a normal distribution. This indicates fat tails in the return distribution, meaning that large shocks occur more frequently than standard models would predict. Economically, markets do not evolve smoothly but are occasionally hit by substantial news or uncertainty shocks.

Next, there is a clear sign of volatility clustering. This means that when uncertainty rises, it tends to remain elevated, and when things are calm, they tend to remain calm. Annualized 30-day rolling volatility ranged from 6.678% to 42.319%, rising above 40% early in the sample before declining to 8–12%. Risk therefore moves in phases rather than staying constant, contradicting basic models that assume stable volatility.

There is also evidence of the leverage effect, which refers to the tendency for volatility to increase when equity prices fall. The largest negative return (6.161%) occurred during the high-volatility period. When prices fall, perceived risk rises, either because firms become more financially leveraged or because investor fear increases, risk increases exactly when investors experience loss.

By contrast, the 6-month Treasury bill yield declined smoothly from 4.3% to 3.6% (a 0.7 percentage point drop) without sharp regime shifts. Overall, instability during this period was driven primarily by changing equity risk rather than movements in interest rates, as there is no clear sign of volatility clustering or sudden shifts in yields in the same way we see for equity risks.

Part C - One Period binomial tree option Pricing:

Part A

We constructed a one period binomial tree to price and at the money (ATM) European call option written on the S&P 500 index. The issuance date $t = 0$ is august 13th 2025, on which the index level is:

$$S_0 = 6466.58$$

Since the option is priced at the money the strike price is set equal to the spot price:

$$K = S_0 = 6466.58$$

The option maturity is 6 months therefore:

$$T = 0.5$$

The annualized volatility estimated from historical log returns prior to the issuance is:

$$\sigma = 9.6278\%$$

The annual risk-free rate is taken from the 6-month Treasury bill yield on the issuance date:

$$r = 3.9984\%$$

The up and down factors for the tree are calculated in Eq 2 Resulting in:

$$u = 1.070449571$$

$$d = 0.934186931$$

Therefore, the possible index values at maturity are therefore:

$$S_u = S_0 u = 6922.15$$

$$S_d = S_0 d = 6040.99$$

The risk neutral probability q is determined by ensuring that the expected growth of the index equals the risk-free rate which therefore yields:

$$q = \frac{(1+r)^T - d}{u - d} = 61.3\%$$

Figure 2 : One period Binomial tree

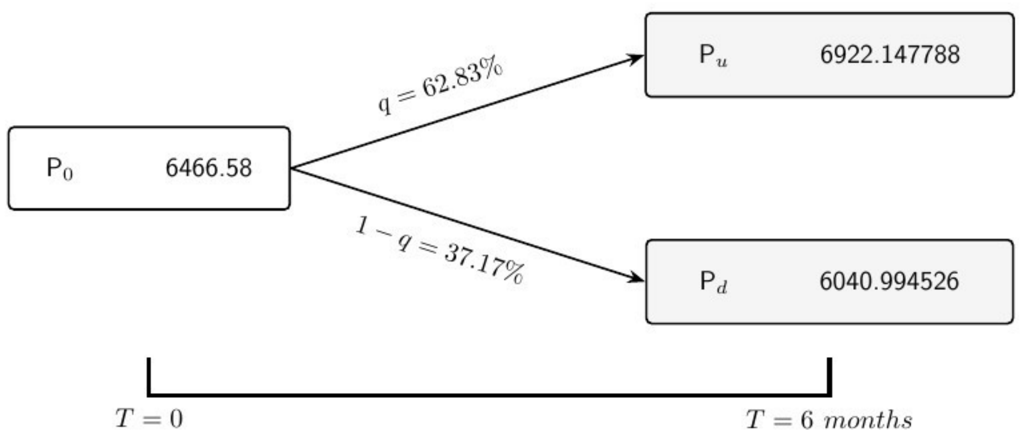

Part B

The options price at initiation, realized Payoff and Return. The calls payoffs at maturity are:

$$C_u = \max(S_u - k, 0) = 455.57$$

$$C_d = \max(S_d - k, 0) = 0$$

The price at $t = 0$ would be calculated as shown below:

$$C_0 = e^{-rT} r[qC_u + (1-q)C_d] \tag{6}$$

$$C_0 = 281.85$$

Figure 3 : One period European call Binomial tree

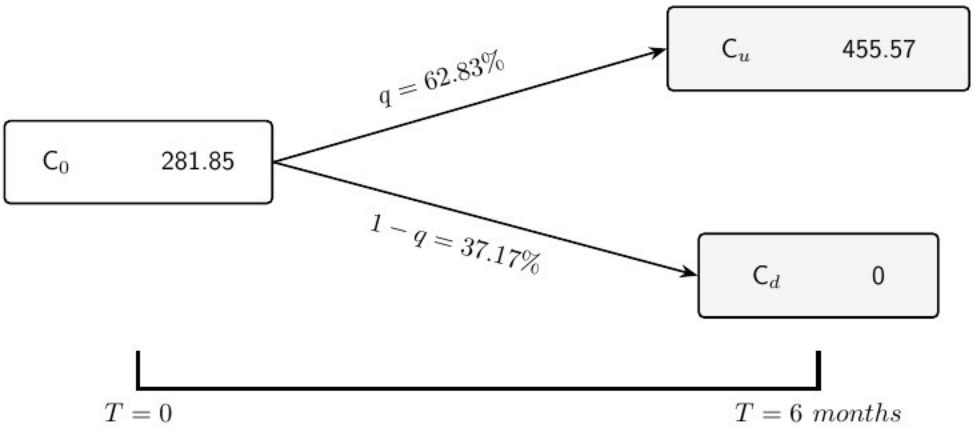

At maturity (13th February 2026), the index level is:

$$S_T = 6836.17$$

The realized payoff:

$$\max(S_T - k, 0) = 369.59$$

The simple holding-period return on the option is

$$\frac{369.59}{281.85} - 1 \approx 31.1\%$$

Part C

The realized return exceeds the return implied by the risk neutral pricing, which assumes that the asset grows at the risk-free rate in expectation. if the volatility increases by 10% then:

$$\sigma' = 1.1$$

$$\sigma = 10.906\%$$

Therefore, once recalibrating the tree like steps above:

$$u' = 1.0778$$

$$d' = 0.92780$$

Therefore, the new option price becomes:

$$C_0' = 303.31$$

We can see that the option prices increase from 281.85 to 303.62.

Figure 4 : One period Binomial tree with 10% increase in Volatility

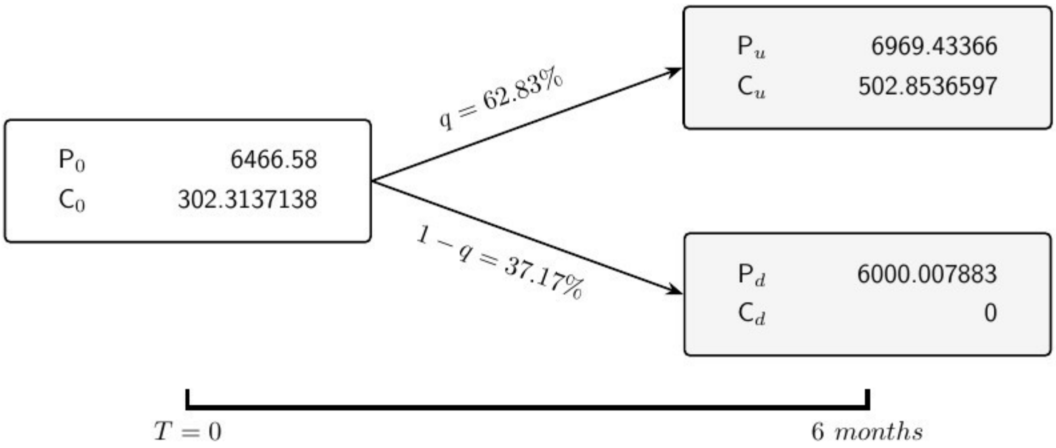

When volatility increases by 10%, the option price rises from 281.85 to 303.31. Economically, this occurs because a European call option has a convex payoff structure: its downside is limited to zero, while its upside potential is unlimited. An increase in volatility widens the distribution of possible future index values, increasing the probability of large upward movements. Since losses are capped but gains are not, greater uncertainty raises the expected payoff under the risk-neutral measure. This highlights that option values are highly sensitive to volatility, even relatively moderate increases in volatility can materially increase the value of the options.

Part D

If the risk-free rate increases by 1% then:

$$r' = 4.0384\%$$

Recalculating the risk neutral probability will give:

$$q' = 0.6327$$

Therefore, the new option price becomes:

$$C_0' = 282.47$$

Figure 5 : One period Binomial tree with 1% increase in Risk free rate

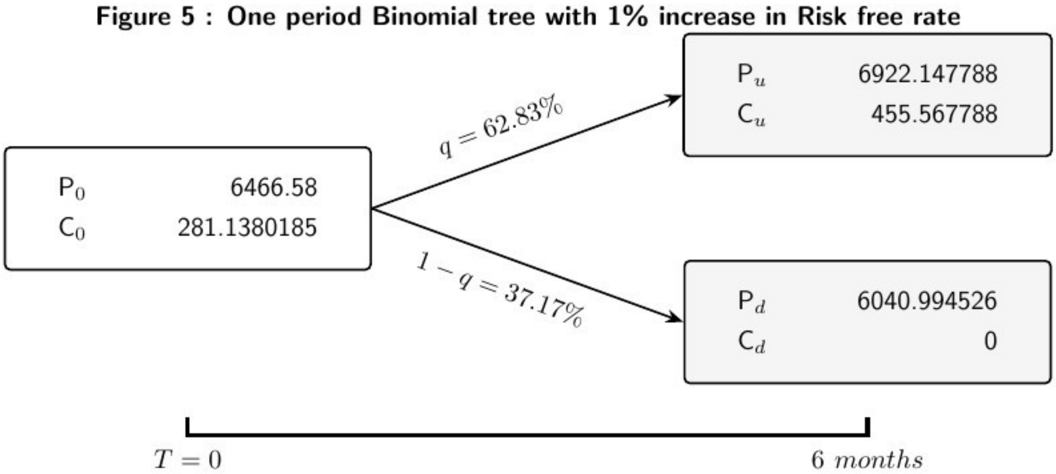

When the risk-free rate increases by 1%, the option price increases slightly. This happens because the stock is assumed to grow faster when interest rates are higher. For a call option this is seen to be beneficial because it increases the chance that the option finishes in the money.

At the same time a higher interest rate also means future payoffs are discounted more heavily. However, in this case the positive effect of a higher expected future stock price outweighs the stronger discounting effect, so therefore the price only rises slightly.

Part D - Two Period binomial tree option Pricing

Part A

This section extends the one period tree, by splitting the six-month window of the ATM European call (issued 13/08/2025) into two 3-month sub-periods. The aim is to compare a constant-parameter two period tree to a time varying two period tree and evaluate how the parameter instability affects the option price. With the following,

$$S_0 = k = 6466.58$$

$$\Delta t = 0.25$$

The following estimated statistics were used:

$$Mean\ return\ t_0 \to t_1 = 16.90\%$$

$$Mean\ return\ t_1 \to t_2 = 5.91\%$$

$$\sigma_1 = 0.1128$$

$$\sigma_2 = 0.1172$$

The first sub period is seen to have experienced stronger average growth, but slightly higher volatility and the second sub period was seen to generate lower returns but marginally higher volatility.

This pattern can suggest a shift in the market environment indicating that the earlier period shows stronger upward momentum and relatively stable conditions. Whereas the later period shows weaker growth but increased uncertainty.

This is consistent with the idea that risk and return are time-varying and not constant.

Part B

The constant model uses one volatility and one interest rate for both steps, after calculations

$$\sigma = 0.1146$$

$$u = 1.05896$$

$$d = 0.94432$$

$$q = 0.5724$$

This approach assumes that market risk and financing conditions remain unchanged over the full six months.

The time-varying model allows the following:

$$\sigma_1 \neq \sigma_2$$

$$R_1 \neq R_2$$

with:

$$\sigma_1 = 0.1128$$

$$\sigma_2 = 0.1172$$

$$R_1 = 1.01038$$

$$R_2 = 1.00959$$

Figure 6: 2 period Constant model binomial tree

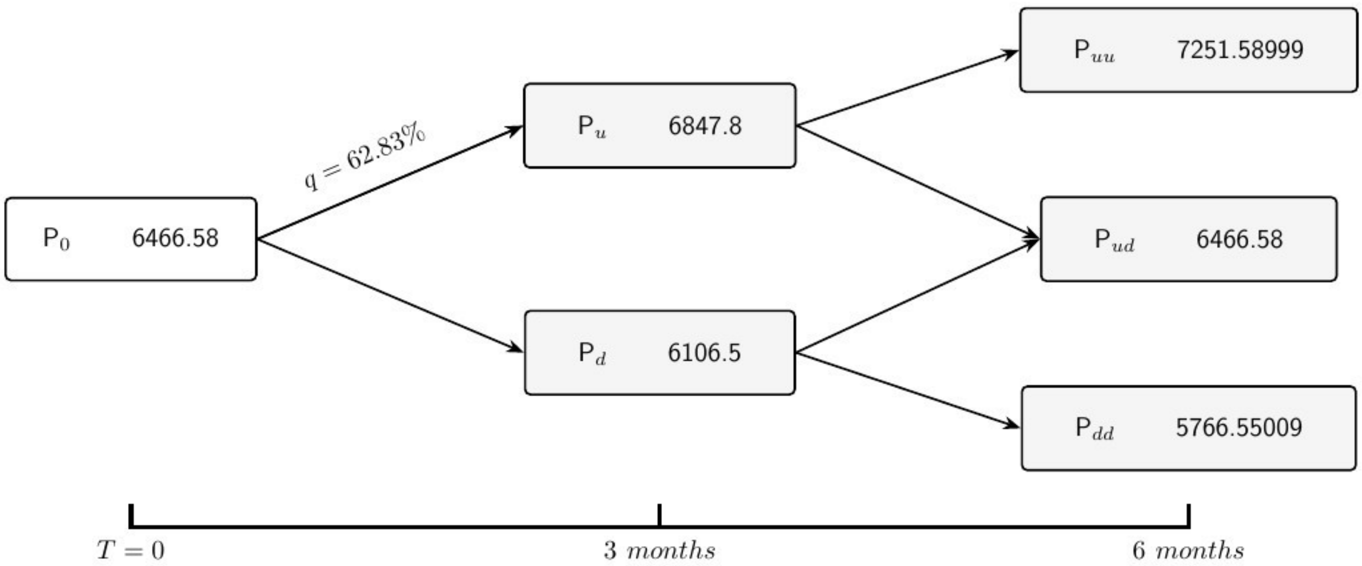

Figure 7: 2 period time-varying model binomial tree

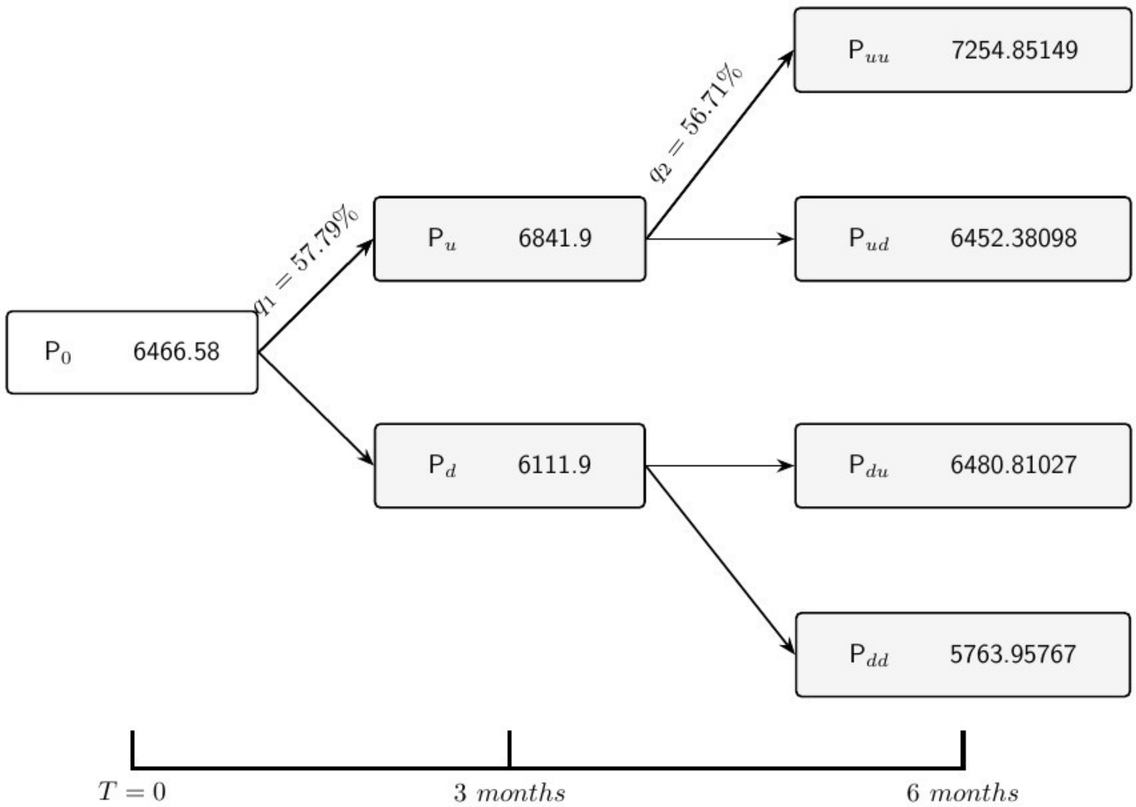

This model however shows what has happened, the volatility was slighter higher in the second-month period. However, at the same time the risk-free rate was slightly lower. The small increase in the volatility is seen as an important indication in financial markets as markets often show volatility clustering, when uncertainty rises it usually stays high for a while instead of dropping back down quickly. The time vary tree captures the real-world behaviour more accurately than the contestant model, as the constant model smooths out everything and misses out on these important changes over time. Due to the volatility increase in the second half, the time-varying model should give a slightly higher option price than the constant model.

Part C

$$C_{0\ (constant\ model)} = 252.20$$

$$C_{0\ (Time\ varying\ model)} = 256.59$$

This is seen to be intuitive as a European call option has a convex payoff which is where the downside is limited to zero while the upside is unlimited. The higher volatility increases the dispersion of possible future index levels at maturity. Because the losses are capped but gains are not greater dispersion raises the expected payoff under risk neutral measure. The convexity of the option increases the effect of even small volatility increases; this further enhances why the time-varying yields a higher price despite the small changes in variance. Due to the price difference being relatively small, it indicates that market conditions remained stable over the 6-month life of the option.

Figure 8: 2 period Constant model European call binomial tree

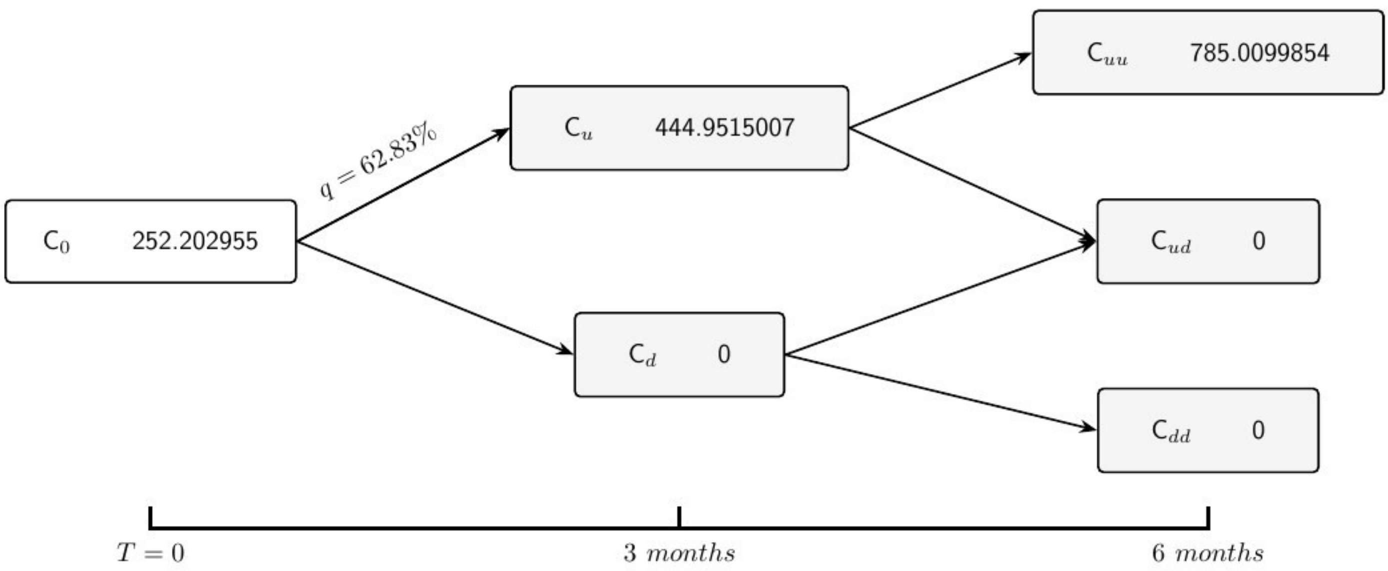

Figure 9: 2 period Time-varying model European call binomial tree

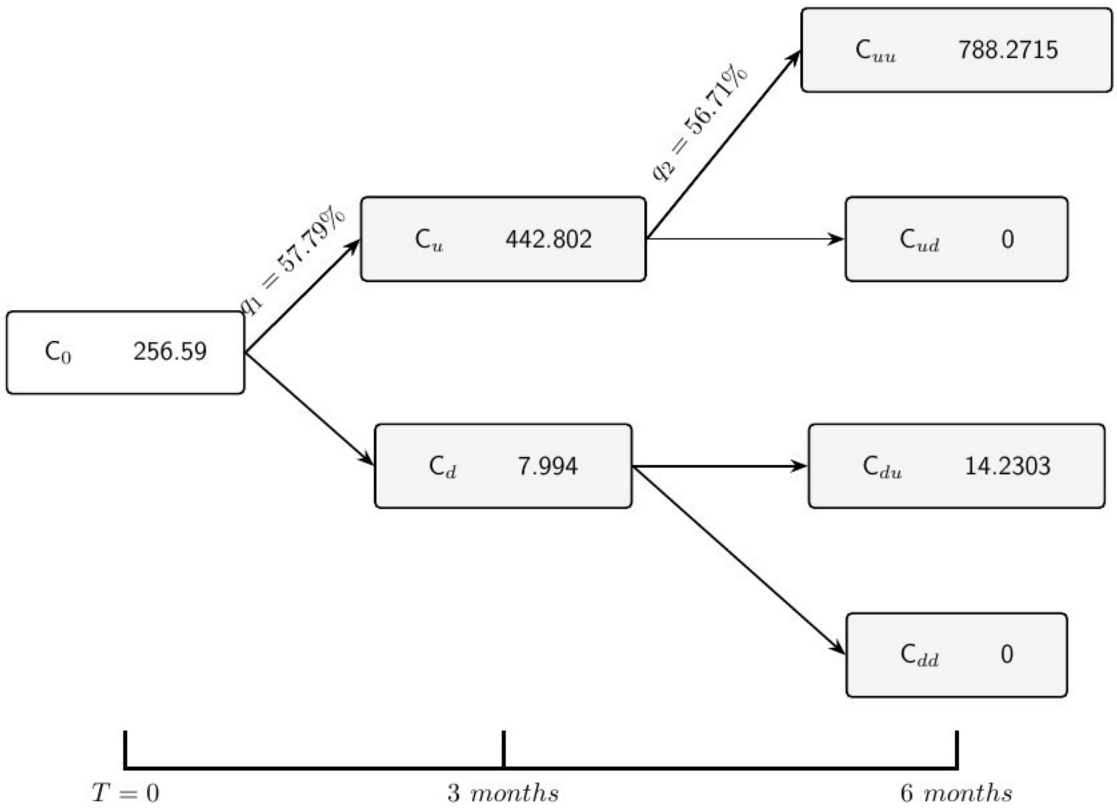

Part D

Modelling multi-period binomial trees with time-varying parameters presents challenges because future volatility and interest rates are unknown at the issuance date. In Part C, parameters were assumed constant over the option's life, offering a clear modelling approach. In Part D, parameters were allowed to vary across the different sub-periods, better reflecting changing market conditions. This approach relied on realised (ex-post) data that would not have been observable at issuance. In real market applications, time-varying trees must therefore be constructed using forward-looking information rather than realised outcomes. This means using current option prices to estimate how the market expects volatility to change over time and using the yield curve to estimate how interest rates are expected to evolve in the future. The parameters u_t, d_t, q_t, R_t would then be calculated separately for each period using these market-implied inputs. While time-varying models better capture volatility clustering and shifts in interest rates, they introduce estimation risk and greater model complexity. Forecasting mistakes can significantly affect the option price.

Appendix

Appendix A

The three days of missing stock data that the bond market was open are:

- 13th October 2025
- 11th November 2025
- 13th February 2026