Explorations in Quantitative Finance and Financial Mathematics

By Jaychandran Padayasi

A set of four mini-projects that I completed during the Summer 2025 offering of Intro to Quantitative Methods in Finance bootcamp by The Erdos Institute.

Portfolio Optimization with Liquidity Constraints

Objective:

This project explores portfolio optimization using Modern Portfolio Theory. The objective was to create sample portfolios consisting of stocks with different risks. I aimed to create a highly customizable Jupyter notebook that can be modified easily to create many custom portfolios

Methodology:

I began by scraping the S&P 500 ticker list along with their sector classifications, then randomly selected 20 stocks from 3 different sectors. Using historical price data, I computed mean daily returns and the covariance matrix for these stocks. I used yfinance for stock data.

The efficient frontier of a set of stocks is given by the set of portfolios which are not dominated by any other portfolio. This means that one cannot construct a portfolio that has both a higher expected return and a lower risk. The project notebook includes a customizable function to generate the efficient frontier. I used the mean of simple returns over a year as proxy for expected return of the stocks.

Additionally, I incorporated the effects of liquidity, with the understanding that low liquidity implies that one cannot trade high volumes of the stock without altering the market price. The "liquidity aware objective" was defined with a Lagrange multiplier λ that punishes low liquidity as follows:

$$O[w] = w^T \Sigma w + \lambda \left(\frac{w^T w}{V} \right)$$

where w is the weight vector, $\Sigma_{\{ij\}} = Cov(S_i, S_j)$, and V_i are the average volumes traded over the period of one year. However, I found that the S&P 500 companies are all to big for this to make a noticeable difference in the efficient frontier. Nevertheless, the feature is included for use with more illiquid assets.

In the notebook, I have generated a sample portfolio for both the high-risk and low-risk scenarios by sampling the efficient frontier. Their stats are as follows:

- High-risk portfolio: 20.05% annualized volatility, 0.21% mean daily return,
- Low-risk portfolio: 13.23% volatility, 0.07% mean daily return

Key goals achieved:

- Developed a customizable module that can be used to compute the efficient frontier of any given set of stocks
- Implemented a liquidity aware optimizer that can be used to protect against illiquid assets.

Hypothesis Testing of Log-Normality in Asset Returns

Objective:

This project explored whether the common financial assumption of log-normal returns holds in empirical data. Using 5000 days of SPY data, I computed log returns and applied several statistical normality tests, all of which rejected the log-normal hypothesis over long horizons. I then broke the data into rolling windows to determine whether normality holds locally. The results consistently showed that shorter windows (45–50 days) passed normality tests, even for assets with high volatility, including cryptocurrencies.

Methodology:

I divided the historical stock returns from the last 5000 days into trading windows of size n_days. Then, I defined a simple measure called percent_normal which is the number of chunks for which normaltest returns a p-value > 0.05 (normality cannot be rejected).

Key Findings:

- Over long time periods, log returns clearly deviate from normality.
- Log-normality of stock returns confirmed: The notebook demonstrates clear evidence that for trading windows of size ~50 days or smaller, there is over a 60% chance that the behavior of the index SPY is log-normal.
- Generalizability: Log-normality is statistically defensible for most assets, including famously volatile ones like crypto. This validates its use in short-term derivative pricing and risk models.
- However, this does not mean that the volatilities are constant for all the windows. The volatility shows expected characteristics such as erratic spikes (volatility clustering) and reverting to its long-term mean soon after spiking (autoregressive).

Analytical Limits of Black-Scholes results for European options pricing

Objective:

The objective of this project was to investigate the behavior of European call and put option prices under the Black-Scholes model in various asymptotic regimes – as time to expiration $T \to 0$ or ∞, and as spot price $\frac{S_0}{K} \to 0$ or ∞.

Methodology:

I used the analytical expression for European options prices as derived using Black-Scholes assumptions. I explored key dependencies, especially on the time-to-expiration and the spot price, by taking carefully controlled limits of the expressions. The results obtained analytically were confirmed by visually plotting the full functions and the calculated asymptotes on the same plot.

Key Findings:

- Short-term options exhibit strong $\sqrt{T}$ dependence.
- Long-dated call options plateau to the spot price; long puts decay to zero in high-rate environments. In the real-world, this translates to the intuition that it does not make sense to sell long put options contracts in a market where the risk-free interest rate is high.

Volatility Modeling and its effects on Hedged Profits

Objective:

In this project, I explored whether modeling stochastic volatility improves hedging performance. As witnessed clearly in Mini-Project 2, stock returns can be safely assumed to be log-normally distributed given one allows for the volatility to change with time. However, the Black-Scholes model assumes a constant volatility, σ. This begs the question, how important is volatility modeling to protect the delta-hedged profits of selling European call options? I perform an in-depth analysis to investigate it using Monte Carlo simulations.

Methodology:

I simulated 5000 stock price paths using two models widely used in the industry for volatility modeling, namely:

- Generalized Autoregressive Conditional Heteroskedasticity GARCH(1,1), and
- Exponential Generalized Autoregressive Conditional Heteroskedasticity EGARCH(1,1),

both of which model time-varying volatility and fat-tailed behavior. I found that EGARCH(1,1) reproduces qualitatively the skewness and the fat tail observed in the SPY index data from Mini-Project 2 more accurately. Using Black-Scholes deltas, I performed discrete delta hedging across these paths and compared hedging profits. Surprisingly, the volatility model had limited impact on outcomes. The explanation lies in the mean-reverting nature of volatility and the fact that hedging is performed at discrete, mismatched intervals relative to volatility spikes.

Key Findings:

- GARCH(1,1) and EGARCH(1,1) capture realistic volatility dynamics, but their impact on hedging profits is minimal.
- Mean reversion and hedging frequency dominate volatility modeling in determining hedging accuracy.
- For European calls, delta-hedging matters far more than detailed volatility modeling.