Adaptive Optimization for Options Trading

Anthony Baez
MIT
acbaez@mit.edu

Jose Ramos
MIT
jricramc@mit.edu

Michael Wong
MIT
mwong21@mit.edu

Abstract

Options trading is a lucrative yet highly complex domain in finance due to its inherent risks and lack of closed-form pricing solutions. Traditional trading strategies fail to adequately address the uncertainties and non-linearities of options trading. In this report, we present an adaptive optimization framework for options trading under uncertainty. By incorporating optimization with both first-degree constraints (e.g., options' greeks) and historical market data, we aim to maximize portfolio profitability while accounting for uncertainty. We utilize benchmarks of traditional trading strategies, such as mean-reversion and long holding, and of adaptive optimization strategies, deterministic and wait-and-see to evaluate the effectiveness of our approach. Our findings demonstrate very significant improvements in risk-adjusted returns compared to traditional methods with some limitations. This paves the way for more resilient and informed decision making in options trading.

1 Introduction

Options are powerful and versatile financial instruments that grant the holder the right, but not the obligation, to buy (call option) or sell (put option) an underlying asset at a predetermined price, known as the strike price. This contract lasts within a specified time frame determined by the expiration date of the option. This flexibility allows options to serve multiple roles in financial markets, including hedging against adverse price movements, speculating on market directions with leverage, and generating additional income through premium collection. Their broad utility has made them a cornerstone of modern financial markets, widely adopted by institutional investors, retail traders, and corporations alike.

In Figure 1, the diagram shows the payoff of a call option with a strike price of $K = 25$ and a premium cost of 120. The option begins to gain value only when the underlying stock price exceeds the strike price, with the breakeven point at $K + \text{Premium} = 29$. Below this level, the buyer incurs a maximum loss equal to the premium paid. Above the breakeven point, the payoff increases linearly as the stock price rises, offering unlimited upside potential while limiting the buyer's downside to the initial premium. This illustrates the asymmetric payoff of call options, making them an effective tool for leveraged exposure to rising stock prices.

However, the inherent complexity of options trading poses significant challenges. Unlike simpler financial assets like stocks or bonds, the pricing and management of options are influenced by a wide array of factors, including the passage of time (time decay), changes in market expectations (implied volatility), and the non-linear relationship between the price of the option and the price of the underlying asset. These factors result in a dynamic and multifaceted risk profile that requires care when trading.

Moreover, the absence of a universally applicable closed-form solution for pricing options, beyond basic models like the Black-Scholes formula, underscores the complexity of these instruments. Real-world trading scenarios often deviate from the assumptions underlying traditional pricing models, such as constant volatility, efficient markets, and the absence of transaction costs. This gap between

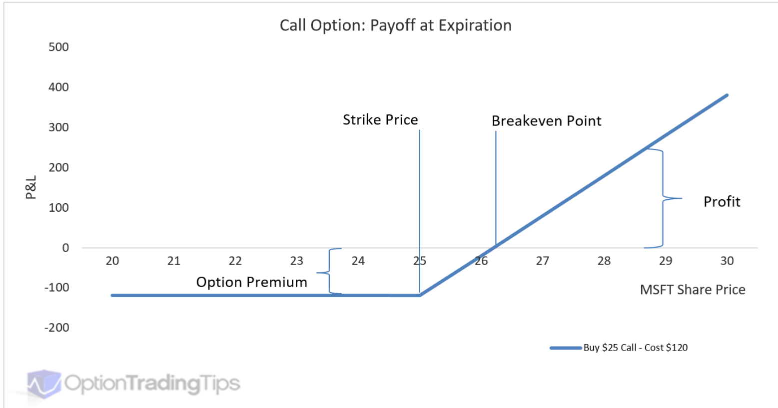

Figure 1: Call option payoff diagram

theoretical models and practical realities complicates decision-making for both those who use and make the models.

The challenges extend further into the realm of portfolio optimization. Constructing and managing an option portfolio involves balancing competing objectives, such as maximizing returns, minimizing risk, and adhering to constraints such as capital requirements, regulatory limits, and liquidity concerns. Static optimization methods can struggle to address the dynamic nature of markets and the nuanced factors influencing options prices. As a result, many strategies either fail to capture the full potential of options or expose traders to excessive risk.

In this report, we address these challenges by proposing an adaptive optimization framework for managing options portfolios. Our approach differs from traditional methods by dynamically incorporating real-world constraints and market uncertainties, thereby improving both risk management and potential returns. This framework also considers important factors in options trading such as implied volatility, theta, diversification, timing of options, and accounting for uncertainty.

Adaptive optimization differs from non-adaptive optimization techniques by taking into account uncertainty and probability by optimizing over different scenarios, s, and outputting an expected value over these scenarios.

$$
\begin{aligned}
\min \quad & \mathbf{c}^\top \mathbf{x} + \mathbb{E}(\mathbf{f}_s^\top \mathbf{y}_s) \\
\text{s.t.} \quad & \mathbf{A}\mathbf{x} = \mathbf{b}, \\
& \mathbf{U}_s \mathbf{x} + \mathbf{V}\mathbf{y}_s = \mathbf{w}_s, \quad \forall s \in S.
\end{aligned}
$$

The above equation is the compact form of an adaptive optimization problem. The objective function contains both the conventional cost term $\mathbf{c}^\top \mathbf{x}$ and the adaptive cost term $\mathbb{E}(\mathbf{f}_s^\top \mathbf{y}_s)$. In this term, there are different costs and decision variable for different scenarios, and the total cost is the expectation over these scenarios. This also applies to constraints, with the first constraint being conventional and the second constraint containing different values across scenarios for the constraint matrix and decision variables. It is this ability to gain insight from scenarios that makes it uniquely advantageous in the options trading domain.

2 Motivation

The options market, known for its high volatility and potential for significant returns, can be both enticing and difficult to navigate for traders and portfolio managers. While various risk management and portfolio construction techniques—such as delta hedging for managing directional exposure, or mean-variance optimization for allocating capital—are commonly employed, they often fall short under real-world conditions. One core issue is market uncertainty: because future asset prices are inherently unpredictable, strategies based on static assumptions can quickly become ineffective when conditions change. Additionally, many approaches do not adequately consider portfolio diversification, neglecting correlations between positions or the potential for sector-specific shocks. This lack of holistic risk assessment leaves portfolios vulnerable to sudden losses.

Moreover, gathering reliable historical price data is both time-consuming and costly, as it must encompass a sufficiently broad and granular data set to capture the underlying asset's volatility patterns and market anomalies. Beyond the data itself, accurately computing first- and second-order derivatives introduces another layer of complexity. These measures—such as Delta, Gamma, Vega, and Theta—require advanced mathematical and statistical methods to account for the multitude of factors influencing option prices, including underlying asset fluctuations, interest rates, dividends, and varying implied volatility surfaces. Even small errors in these calculations can lead to flawed assessments of risk and opportunity. The combined difficulties of obtaining high-quality data and performing intricate derivative computations underscore the need for more efficient, adaptive, and scalable approaches that can navigate the market's complexity while controlling costs.

In this project, we leverage adaptive optimization to address the complexities and uncertainties of options trading. By allowing the model to make adjustments for different price scenarios, this approach supports the construction of portfolios that can dynamically respond to evolving market conditions. Importantly, it allows for the incorporation of option Greeks and portfolio-level metrics, ensuring that our strategies remain balanced and diversified. As a result, adaptive optimization provides a practical, responsive, and scalable framework that aligns closely with the project's goals of delivering maximum profit under uncertainty.

By addressing these gaps, our framework aims to deliver a practical and efficient solution for managing options portfolios under uncertainty. Importantly, over 99% of investors lack access to the sophisticated tools, real-time analytics, and vast resources that institutional players on Wall Street take for granted. This disparity places retail traders and smaller portfolio managers at a considerable disadvantage when competing in options markets, where split-second decisions and advanced predictive insights can determine success or failure. Traditional trading strategies often fail to account for this resource gap, leaving non-institutional investors to navigate highly complex instruments without the same computational power, proprietary models, or access to granular data. By developing an adaptive optimization framework that leverages advanced methods and real-time market adjustments, we aim to level the playing field, democratizing access to cutting-edge technology and predictive insights. This approach not only provides a competitive edge for retail investors but also enhances their ability to construct well-diversified, dynamically managed portfolios that can withstand market volatility and capitalize on opportunities previously reserved for institutional giants.

3 Methodology

3.1 Data Collection

We collected options data using Market Data, a market data provider that offers both an API and a Google Sheets add-on, making it seamless to retrieve historical options data at scale. For this project, we focused on five major stocks: AMZN, TSLA, AAPL, GOOG, and META. For each stock, we selected 10 different options and gathered a month's worth of price data for each option. The key data points used in our model included the mid-price of the option, the underlying stock price, and the strike price.

However, one notable limitation of the Market Data platform was the absence of first- and second-order derivatives, which are essential for understanding an option's sensitivity to changes in market variables. Without access to pre-computed Greeks, we would have been unable to directly incorporate them into our model as constraints, which are needed for more sophisticated risk management and optimization strategies.

To address this gap, we developed our own formulas to approximate the Greeks using standard option pricing models. By doing so, we were able to derive estimates of these values, which provided a more comprehensive view of the options' behavior and allowed us to enhance the robustness of our optimization framework. While these calculations introduced an additional layer of complexity, they were necessary to ensure our model could better capture the nuanced relationships between option prices and their underlying factors.

Theta Formula

Theta measures the time decay of an option's value as expiration approaches, quantifying how much the option price decreases per day, assuming other factors remain constant. This decay is most significant near expiration and at-the-money options, where θ grows in magnitude. Accurate modeling of θ enables better pricing and risk management.

$$\theta(K, S, D) = -\frac{2.1}{\sqrt{D}} \exp\left(-\frac{(K-S)^2}{2 \times 20^2}\right)$$

Where:

- θ: The option's theta
- K: The strike price
- S: The underlying stock price
- D: Days until expiration

Key Components:

- Time Decay Factor: $-\frac{2.1}{\sqrt{D}}$ models faster time decay as expiration ($D \to 0$) nears.
- Moneyness Adjustment: $\exp\left(-\frac{(K-S)^2}{2\times 20^2}\right)$ ensures θ peaks at-the-money ($K \approx S$) and decreases for deep in/out-of-the-money options.

Inverse Black-Scholes Algorithm

This algorithm estimates implied volatility σ, the value that equates the Black-Scholes call price to a given market price, using the Newton-Raphson method.

Steps:

1. Black-Scholes Price:

$$C(S, K, T, r, \sigma) = SN(d_1) - e^{-rT} KN(d_2), \quad d_1 = \frac{\ln(S/K) + (r + 0.5\sigma^2)T}{\sigma\sqrt{T}}, \; d_2 = d_1 - \sigma\sqrt{T}.$$

2. Vega:

$$\text{Vega} = SN'(d_1)\sqrt{T}, \quad N'(\cdot) \text{ is the standard normal PDF.}$$

3. Iterative Update:

$$\sigma_{n+1} = \sigma_n + \frac{C_{\text{target}} - C(S, K, T, r, \sigma_n)}{\text{Vega}(\sigma_n)}.$$

The algorithm starts with an initial guess for σ (e.g., 0.5) and iteratively refines it until the difference between the target and model price is within a small tolerance (e.g., 10^{-5}).

3.2 Data Processing

To introduce uncertainty to our model, we created 20 different scenarios for the price of the underlying stock for the adaptive optimization. This was done by first calculating the implied volatility of the stock using our inverse Black-Scholes algorithm. This gives us a stock price path with the same returns but a higher standard deviation of returns. Then, we create the different scenarios using a Monte Carlo simulation that adds Gaussian noise at each time step from a normal distribution of mean 0 and a standard deviation of the scaled returns. This results in 20 different price scenarios that both mimic the real price trend but also simulates uncertainty on the future price.

The final processing step was to line the matrix of options data by calendar day in the time axis, not by days until expiration. This converted a matrix of size (days, option stock) = (21, 10, 5) to (66, 10, 5) by padding each 21 day option by zeros to get 66 days. We also calculate the mean underlying stock price on each day and make all values across all options equal to this mean to ensure that all options have the same underlying stock price on the same day.

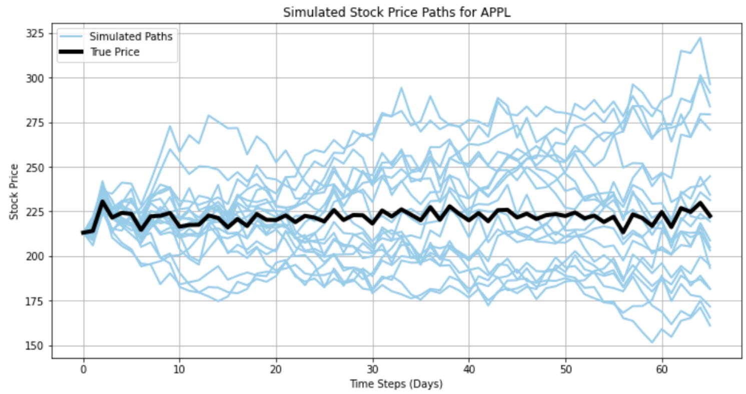

Figure 2: APPL stock with 20 generated scenarios

3.3 Optimization Model

We modeled options trading with an adaptive linear mixed-integer model in Julia using Gurobi as a solver. We formulate the process with the first-stage decision being which 4 out of the 5 stocks to trade with, which is controlled by decision variable c below. Then, we introduce uncertainty using the future price scenarios. The second-stage decisions afterward are purchase and sale of options, controlled by the b and s decision variables. We describe and explain the objective function and constraints below:

Parameters:
$T = 66,\ n_{scen} = 20,\ n_{opt} = 10,\ n_{stock} = 5,\ p = 100$

T was the total number of days covered by our options. n_{scen} was the number of different future underlying stock price scenarios we considered. n_{opt} was the total number of options per stock. n_{stock} was the total number of stocks included in our model. $p = 100$ was the starting number of dollars in our portfolio.

Decision Variables:
$b_{t,n,i,j} \in \mathbb{Z}$: Number of options to buy at time t = 1, ..., T in option i = 1, ..., n_{opt} in stock j = 1, ..., n_{stock} in scenario n = 1, ..., n_{scen}
$s_{t,n,i,j} \in \mathbb{Z}$: Number of options to sell at time t = 1, ..., T in option i = 1, ..., n_{opt} in stock j = 1, ..., n_{stock} in scenario n = 1, ..., n_{scen}
$c_j \in \{0, 1\}$ Whether or not stock j = 1, ..., n_{stock} can be bought

The buying and selling decisions were decided by decision variables, with the integer value entry being the number of options bought at that specific time t, scenario n, option i, and stock j. The stocks available to purchase was controlled by binary decision variable c.

Objective Function:

$$\max \quad \frac{1}{n_{\text{scen}}} \cdot \sum_{n=1}^{n_{\text{scen}}} \sum_{t=1}^{T} \sum_{i=1}^{n_{\text{opt}}} \sum_{j=1}^{n_{\text{stocks}}} \left(price_{s,n,i,j} \cdot s_{t,n,i,j} - (strike_{t,i,j} + mid_{t,i,j}) \cdot b_{t,n,i,j} - 0.1 \cdot b_{t,n,i,j} \right)$$

This objective function describes the maximization of total profit of our options trading strategy under adaptive optimization. The profit is the price of the underlying stock times the number of it we sell at a given time and scenario minus the value of the option, which is the sum of the strike price and the market value mid, times the number of it we buy at a given time and scenario. We also subtract by a transaction cost of 10 cents per option bought. We sum over all price scenarios and divide by the number of scenarios to get the weighted sum of the profit. This assumes that each price scenario is equally likely.

Non-Negativity Constraints:
$b_{t,n,i,j} \geq 0,\ \forall t = 1, \ldots, T;\ n = 1, \ldots, n_{\text{scen}};\ i = 1, \ldots, n_{\text{opt}};\ j = 1, \ldots, n_{\text{stocks}}$
$s_{t,n,i,j} \geq 0,\ \forall t = 1, \ldots, T;\ n = 1, \ldots, n_{\text{scen}};\ i = 1, \ldots, n_{\text{opt}};\ j = 1, \ldots, n_{\text{stocks}}$

Option Existence Constraints:

$b_{t,n,i,j} \leq 0,\ \forall i = 1, \ldots, 10;\ n, j;\ \forall t \notin \{\, 5(i-1)+1, \ldots, 5(i-1)+21 \,\}$

This constraint was used to prevent the option from being purchased outside of its time of existence. This is a byproduct of choosing to align our options data by time in market days as opposed to days until expiration for each option.

Cannot Sell Before Buying:
$$\sum_{\tau=1}^{t-1} b_{\tau,n,i,j} \geq \sum_{\tau=1}^{t} s_{\tau,n,i,j} \quad \forall t = 2, \ldots, T;\ n, i, j$$

Buy-Sell Balance:
$$\sum_{t=1}^{T} b_{t,n,i,j} = \sum_{t=1}^{T} s_{t,n,i,j} \quad \forall n, i, j$$

Each option contract must be bought before being sold and be bought and sold the same number of times.

Portfolio Budget Constraint:
$$\sum_{t=1}^{T} \sum_{i=1}^{n_{\text{opt}}} \sum_{j=1}^{n_{\text{stock}}} b_{t,n,i,j} \cdot mid_{t,i,j} \leq p \quad \forall n$$

The total value of the options contracts we buy must be less than the portfolio. This constraint also prevents rapidly buying and selling the same option to more realistically simulate options trading.

Enforcing Chosen Stock Constraint:
$$\sum_{t=1}^{T} \sum_{i=1}^{n_{\text{opt}}} b_{t,n,i,j} \leq M \cdot c_j \quad \forall n, j, \qquad \sum_{j=1}^{n_{\text{stock}}} c_j = 4$$

Diversification Constraint:
$$\sum_{t=1}^{T} \sum_{i=1}^{n_{\text{opt}}} b_{t,n,i,j} \ \leq\ 20 \quad \forall s = 1, \ldots, n_{\text{scen}},\ j = 1, \ldots, n_{\text{stocks}}$$

The model can only buy a total of 20 options contracts for each stock. This is used to enforce diversification over stocks of the portfolio.

Theta Constraint:
$$-\sum_{\tau=1}^{t} \sum_{i=1}^{n_{\text{opt}}} \sum_{j=1}^{n_{\text{stocks}}} (b_{\tau,n,i,j} - s_{\tau,n,i,j}) \cdot \theta_{t-1,i,j} \leq 0.1 \cdot \sum_{\tau=1}^{t} \sum_{i=1}^{n_{\text{opt}}} \sum_{j=1}^{n_{\text{stocks}}} (b_{\tau,n,i,j} - s_{\tau,n,i,j}) \cdot mid_{t,i,j}\ \forall n; i = 2, \ldots, T$$

The total theta of the options contracts in the portfolio must always be at most 10% of the total value of the options contracts held. This reflects real-world strategy to not hold onto contracts too close to their expiration date.

3.4 Benchmarks

To evaluate the performance of adaptive optimization when applied to options trading and our specific implementation, we compared our method to the traditional trading techniques of mean reversion and long calling. Mean reversion is based on the idea that an asset's price will revert to its historical average over time after deviating significantly. Therefore, the stock is bought when its price deviates below a certain chosen threshold, which we chose as 2 times the standard deviation of the standard deviation, and sold when the stock price deviates above a certain threshold; in our case, a 10-day running average. This strategy was implemented only with the underlying stock price and not with the options mechanics. Long holding is a simplified version of options trading where a call is bought and held until expiration, at which point it is either worthless if it finishes out of the money (meaning that the price of the stock is less than the strike price), or it is in the money and the call is exercised at a profit. This was implemented using the same model as the deterministic optimization model but with the added constraint that once bought, options can only be sold at their time of expiration.

We also compared our adaptive optimization strategy to deterministic and wait-and-see benchmarks. The wait-and-see solution provides us with the profit under perfect knowledge of the future. In our model, this means that there is a c_j for every scenario. The stochastic solution makes fixed decisions based on expected outcomes, offering a more practical approach. In our model, this means that decision variables b and s must be the same for all scenarios, and the average prices across the scenarios are used to decide options trades. This allows us to calculate the value of adaption to future scenarios (by comparing deterministic and adaptive) and the value of perfect future information (by comparing adaptive to wait-and-see).

4 Results

4.1 Traditional Benchmarks

The results of our adaptive optimization model, traditional trading benchmarks, and optimization benchmarks can be found in Table 1. Our adaptive optimization model had a end portfolio value of 4956.66 dollars or 4856% return. This was significantly higher, at 1 order of magnitude, than the worst performing traditional benchmark of mean reversion at 276 dollars or 176% return. The adaptive optimization model was also significantly better than the best performing traditional benchmark of long holding at 2265 dollars or 2165% return, a roughly 2x improvement. These results show that our adaptive optimization model significantly outperforms traditional trading strategies that trade both stocks and options.

4.2 Optimization Benchmarks

We can also compare our adaptive model results to the deterministic and wait-and-see benchmarks. Our adaptive model, at a final portfolio value of 4956.66 dollars and a return of 4956% performed, as expected, marginally better than the deterministic model at 4846.23 dollars and 4746% return and marginally worse than the wait-and-see model at 5178.24 dollars and 5078% return.

4.3 Performance Metrics

Using the values from our adaptive optimization benchmarks, we can calculate performance metrics that quantify the value of certain knowledge in our problem. The value of stochastic solution (VSS) is the difference between the objective value between the adaptive (stochastic) model and the deterministic model. VSS measures value of the adaptive solution taking into account uncertainty using the scenarios, which it does by making different buying and selling actions according to the different prices in each scenario, as opposed to no different price scenario. We calculate VSS to be 109.43 dollars. While this value is low compared to our final portfolio values, the assumed knowledge of a single perfect price is impossible to determine when actually trading on the market.

We can also calculate the expected value of perfect information (EVPI), which is calculated by taking the difference between the objective value of the wait-and-see and adaptive models. EVPI measures the value of that perfect information of the future has to the model to the first-stage decision. In the our options trading model, this is the value of the model buying a different 4 of the 5 stocks in each price scenario when compared to needing to decide on the same 4 of the 5 stocks for all scenarios. We find that the EVPI is 332.01 dollars. This number is also low compared to the final value in the portfolios in all optimization models, which makes sense as you are only leaving 1 stock out with the decision.

Strategy	Traditional		Optimization		
Type	Mean Reversion	Long Hold	Adaptive	Deterministic	Wait-and-See
Portfolio Value ($)	276.24	2265.08	4955.66	4846.23	5178.24
Total Returns (%)	176	2165	4856	4746	5078

Table 1: Number of dollars in portfolio and total returns at end of option trading with different strategies.

5 Impact

5.1 Trading Sector and Economy

An optimization-based options trading strategy that outperforms traditional approaches, such as mean reversion or long-only strategies, could have significant implications for financial markets and investment management. By systematically exploiting inefficiencies in the options market through quantitative methods, this approach has the potential to generate higher risk-adjusted returns while reducing exposure to downside risks. Such strategies enable traders and portfolio managers to better navigate market volatility, optimize capital allocation, and enhance overall portfolio resilience.

The adoption of optimization-based techniques could also contribute to greater market efficiency. By encouraging more frequent and precise trades, driven by real-time price rebalancing and risk adjustments, these methods may increase liquidity demand in options markets. This, in turn, could lead to narrower bid-ask spreads, reducing transaction costs and improving market accessibility for all participants.

However, these benefits come with trade-offs, particularly the increased computational costs and resource demands required to implement and execute complex optimization algorithms in real time. As financial institutions and advanced trading systems increasingly rely on such models, the need for high-performance infrastructure and low-latency execution systems will grow.

From a broader perspective, the integration of quantitative optimization methods into options trading could drive positive long-term economic impacts. More efficient options pricing contributes to better price discovery, which promotes stability and transparency in financial markets. This benefits not only institutional players but also retail investors, as it fosters an environment where risks are more accurately priced and opportunities are more fairly distributed. As opposed to the current state of market making, where quant firms capitalize on the fact that retail traders are sending inefficient orders to the market due to a lack of data-driven insights.

5.2 Limitations

Despite the significant out-performance of the adaptive optimization model, our current model is limited in its application to real world trading. This is mainly because it relies on the true future price of stocks that we have in historical data. Even though we simulate uncertainty by adding noise and with different scenarios, this is still significantly more and specific information than what real options traders and algorithms work with. Another drawback is that our options data was selected when there was sufficient market liquidity and efficiently priced options. In real options trading, there can be times with low liquidity and inefficiently priced options that can interfere with the precise calculations of the optimization model. Despite these limitations however, the considerable results demonstrate that adaptive optimization-based options trading could be a promising direction to explore and apply to more realistic market scenarios.

6 Future Work

If we had another week, we could increase the scale of our data by introducing more stocks, options for each stock, and longer time until expiration for each stock. This would bring the model closer to a real world application where the whole market is available. We would also implement more constraints based on the options greeks, such as rho and gamma, that describe valuable information about the position being taken that options traders use. Another constraint that was implemented successfully but became computationally infeasible was a constraint on the average correlation on stocks purchased. While this is something that would have been extremely helpful in minimizing the risk, it made the optimization problem a quadratic-constrained mixed-integer problem and caused the solver's runtime to increase exponentially. This could be implemented if we had access to more compute or a better quadratic program solver. The inverse Black-Scholes model used to calculate IV could also be improved as there were some assumptions that we made, such as a constant risk-free rate for the entire year (which is never the case), and the same precision level for all of the iterations, that we can challenge in order to get even more granular insights.