

CallCulator Handbook

Advanced version

Aaron Gonsior

February 18, 2026

Contents

| | | |
|------------|--|-----------|
| I | Foundations | 4 |
| 1 | Welcome to CallCulator | 4 |
| 2 | Basics of Options Trading | 4 |
| 2.1 | What is an Option? | 4 |
| 2.2 | Visualizing Option Performance | 4 |
| 2.3 | American vs. European Options | 5 |
| 2.4 | What is a Spread? | 6 |
| II | Core Philosophy: Why Probability Distributions? | 6 |
| 3 | The Problem with Price Targets | 6 |
| 4 | Our Approach: Probability Density Functions | 6 |
| 4.1 | From PDF to Risk Profile: The Mathematical Bridge | 7 |
| 5 | Probability Distributions Over Time | |
| | (In Development) | 8 |
| III | Using CallCulator | 9 |
| 6 | Basic Calculation Mode | 9 |
| 6.1 | Workflow | 9 |
| 6.2 | Results Displayed | 9 |
| 7 | Probability & Risk Mode | 9 |
| 7.1 | Workflow | 10 |
| 7.2 | Search Algorithms | 10 |
| 7.2.1 | Brute Force | 10 |
| 7.2.2 | Hybrid Search (GSGD) | 11 |
| 7.2.3 | Performance & Complexity | 11 |
| 7.3 | The Optimization Engine | 11 |
| 7.4 | Results Displayed | 12 |
| 8 | Max Min Mode | 14 |
| 8.1 | Single Stock Mode | 14 |
| 8.1.1 | Workflow | 14 |
| 8.1.2 | Mathematical Formulation | 14 |
| 8.2 | Multi Stock Mode | 15 |
| 8.2.1 | Workflow | 15 |
| 8.2.2 | Cross-Asset Diversification | 15 |
| 8.3 | Soft Minimum Optimization (Percentile-Based Targets) | 15 |
| 8.3.1 | When to Use Soft Minimum | 15 |
| 8.3.2 | How It Works | 15 |
| 8.3.3 | Hybrid Approach: VaR Floor with Growth | 16 |
| 9 | Market Data Overview | 16 |

| | |
|---|-----------|
| 10 Implied Probability Distributions | 16 |
| 10.1 What the Market “Thinks” | 16 |
| 10.2 Dynamic Implied Volatility | 17 |
| 11 Data Pipeline | 18 |
| | |
| IV Risk Theory | 18 |
| | |
| 12 Risk Profiles | 18 |
| 12.1 Personal Risk Tolerance Profile | 18 |
| 12.2 Investment Risk Profile | 18 |
| 12.3 Matching Profiles | 18 |
| | |
| 13 On the Merit of Spreads | 19 |
| 13.1 The Single-Step Argument Against Spreads | 19 |
| 13.2 The Multi-Step Counter-Argument: Geometric Growth | 19 |
| 13.2.1 Alice and Bob: A Story of Ruin | 19 |
| 13.2.2 Position Sizing Saves the Day | 19 |
| 13.3 Beyond the Coin Flip: The Ergodicity Problem | 19 |
| 13.4 Monte-Carlo Simulation: Quantifying Hidden Risk | 20 |
| 13.5 FFT Convolution: Deterministic CAGR Distributions | 20 |
| 13.5.1 Core Insight: Quantile Wealth Multipliers | 20 |
| 13.5.2 Protocol | 21 |
| 13.5.3 Advantages over Monte-Carlo | 21 |
| 13.6 Kelly Criterion: Optimal Reinvestment Fraction | 22 |
| 13.7 Log-Optimal Strategy: Maximizing Median CAGR | 22 |
| 13.7.1 The Log-Optimal Payoff | 22 |
| 13.7.2 Adding the Risk Tolerance Floor | 23 |
| 13.7.3 Computing λ : Newton-Raphson on the Budget Equation | 23 |
| 13.8 Implications for Options | 23 |
| | |
| 14 Mathematical Constructions: The Six Optimization Regimes | 23 |
| 14.1 Regime 0: The Scavenger (Micro-Arbitrage) | 24 |
| 14.1.1 Residual Mispricing | 24 |
| 14.1.2 Butterfly Spreads and Neighbor Selection | 24 |
| 14.1.3 Boosting Macro Strategies | 24 |
| 14.2 Regime 1: The Sniper (Unconstrained Maximization) | 24 |
| 14.3 Regime 2: The Insurer + Sniper (Constrained Optimization) | 25 |
| 14.4 Regime 3: The Farmer (Log-Optimal / Median CAGR Maximization) | 25 |
| 14.4.1 Risk-Tolerance Constrained Farmer | 25 |
| 14.4.2 Solving for λ : Newton-Raphson | 26 |
| 14.5 Regime 4: The Bunker (Max Min Optimization) | 26 |
| 14.6 Regime 5: The Sentinel (Soft Minimum / Percentile-Based Targets) | 26 |
| 14.7 Summary: Comparison of All Regimes | 27 |
| | |
| V Reference | 27 |
| | |
| 15 Glossary | 27 |

Part I

Foundations

1 Welcome to CallCulator

CallCulator is an interactive web-based tool for quantitative options analysis. It enables users to move beyond zero-dimensional metrics (price targets) and instead express their market beliefs through probability distributions, which are then combined with risk tolerance profiles to find mathematically optimal investment strategies.

The platform offers three main modes, accessible via the tabs at the top of the page:

1. **Basic Calculation** – Manually define a custom option spread and compute its full performance profile, Black-Scholes time-decay surface, and Monte-Carlo CAGR statistics.
2. **Probability & Risk** – Supply a probability distribution and risk tolerance to let the backend search engine find the globally optimal strategy across hundreds of millions of option combinations.
3. **Max Min** – Find strategies that maximize the worst-case outcome (minimax optimization), either for a single stock or across a multi-stock portfolio.

2 Basics of Options Trading

2.1 What is an Option?

An **option** is a financial derivative that grants the holder the right—but not the obligation—to buy or sell an underlying asset at a pre-agreed price (the **strike price** K) on or before a specified date (the **expiration date** T).


- **Call Option** $C(K, T)$: Right to *buy* at K . Payoff at expiration: $\max(S_T - K, 0)$ where S_T is the stock price at time T .
- **Put Option** $P(K, T)$: Right to *sell* at K . Payoff at expiration: $\max(K - S_T, 0)$.

The **premium** p is the market price of the option. In percentage terms, the profit/loss of a call option is:

$$g_c(S_T) = \max\left(-1, \frac{S_T - K - p}{p}\right) \quad (1)$$

2.2 Visualizing Option Performance

The graph below shows the percentage gain or loss of different call options (red lines) compared to simply owning the stock (blue line):



Red: percentage gain of various call options. Blue: the underlying stock.

Notice how the red lines are much steeper than the blue line. A 10% rise in the stock price might translate to a 50% or even 200% rise in a call option's value. But if the stock doesn't reach the strike price, the option goes to zero—a 100% loss. This is the leverage effect in action.



A single call option's performance at expiration.

2.3 American vs. European Options

While you can trade (buy/sell) any option on the market at any time, there is a distinction in when you can **exercise** the right to buy or sell the underlying stock:

- **American-Style:** You can exercise the option at any point up to and including the expiration date. Most standard equity options are American-style.
- **European-Style:** You can only exercise the option on the expiration date itself. Index options are often European-style.

For most traders, this distinction is technical because options are rarely exercised early; they are usually sold to close the position.

2.4 What is a Spread?

A **spread** is a convex combination of financial products:

$$\sum_i w_i g_i \quad \text{where} \quad \sum_i w_i = 1, \quad w_i \in [0, 1] \quad (2)$$

Spreads allow you to engineer a specific risk/reward profile. CallCulator supports multi-leg strategies including **short (write) positions**, enabling bull spreads, bear spreads, iron condors, and custom combinations. The search algorithms don't know any difference between those traditionally grouped spreads into names, it is completely arbitrary.

Part II

Core Philosophy: Why Probability Distributions?

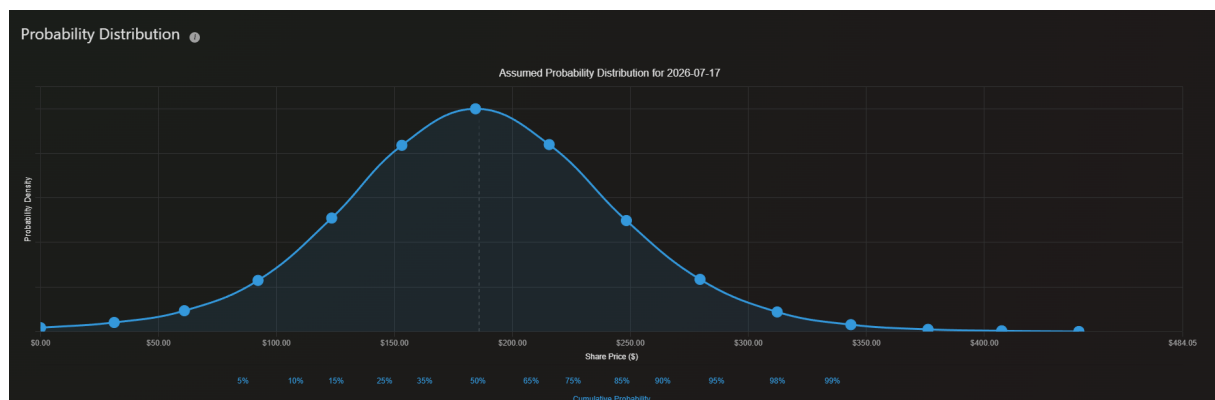
3 The Problem with Price Targets

Conventional analysis relies on a “price target”—a single number X representing where you think a stock will be by some date. This is a zero-dimensional metric that collapses all uncertainty into a point estimate.

The problem: two investors may both have a target of \$150, yet one might be very confident (narrow bell curve around \$150) while the other sees wild uncertainty (flat distribution from \$80 to \$220). A price target cannot distinguish these beliefs, leading to fundamentally different appropriate strategies.

4 Our Approach: Probability Density Functions

CallCulator replaces the price target with a **Probability Density Function (PDF)** $\mu(x)$ for the stock price at a future date. This captures the full shape of uncertainty:

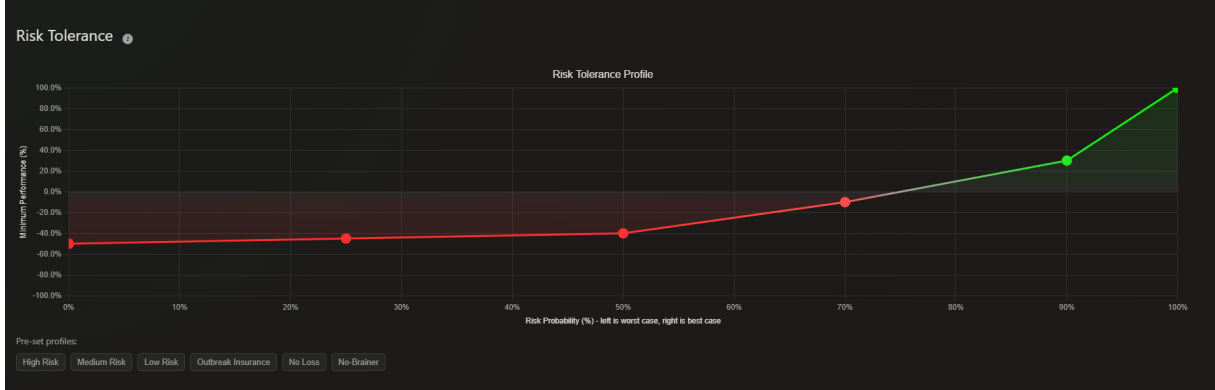


A probability distribution supplied by the user: the peak at \sim \$180 indicates the most likely price.


Given $\mu(x)$ and a set of chosen options $\{g_i\}$, the **expected return** of a strategy is:

$$\mathbb{E}_\mu \left[\sum_i w_i g_i \right] = \sum_i w_i \int_{\mathbb{R}} g_i(x) \mu(x) dx \quad (3)$$

Combined with a **Risk Tolerance Profile**, we can find the strategy that maximizes expected return subject to the constraint that its risk profile stays within the user’s tolerance at every cumulative probability level.



The user’s risk tolerance profile $\hat{r}_{tol}(\lambda)$: a continuous boundary from worst case (left) to best case (right).



The found strategy’s risk profile $\hat{r}_c(\lambda)$ (blue) satisfying $\hat{r}_c(\lambda) \geq \hat{r}_{tol}(\lambda) \forall \lambda \in [0, 1]$

4.1 From PDF to Risk Profile: The Mathematical Bridge

The probability distribution $\mu(x)$ and a strategy’s performance function $p_c(x)$ together determine a **risk profile** $\hat{r}_c : [0, 1] \rightarrow \mathbb{R}$ —a non-decreasing curve that maps each cumulative probability level λ to the return achieved at that quantile. This is the object plotted in the Risk Profile Comparison charts.

Monotone strategies (single legs). For a single option whose performance $p_c(x)$ is monotone in the share price, the construction is direct. Define the CDF $\Psi(x) = \int_0^x \mu(y) dy$, which maps prices to cumulative probabilities. The risk profile is:

$$\hat{r}_c(\lambda) = p_c(\Psi^{-1}(\lambda)) \quad (4)$$

The CDF inverse $\Psi^{-1}(\lambda)$ converts the probability level λ to the corresponding share price (“at probability level 30%, the stock is at \$X”), and p_c converts that price to the strategy’s return. Since both Ψ^{-1} and p_c are monotone, \hat{r}_c is itself non-decreasing.

Non-monotone strategies (spreads). For multi-leg spreads, $p_c(x)$ is generally not monotone—it may rise, peak, and fall. In this case, different share prices can produce the same return. The risk profile is then the **quantile function of the induced return distribution**:

$$\hat{r}_c(\lambda) = \inf \left\{ z \in \mathbb{R} : \int_{\{x : p_c(x) \leq z\}} \mu(x) dx \geq \lambda \right\} \quad (5)$$

The integral computes, for each return level z , the total probability of achieving a return *at most* z —integrating over all (possibly disjoint) price regions where the strategy underperforms z . Inverting this CDF produces the non-decreasing risk profile.

In practice, the computation proceeds as follows: First, the multi-leg spread is converted into a performance curve $p_c(x)$ by combining all option payoffs with their respective weights. A series of return levels z is sampled uniformly across the range from the minimum to maximum possible performance. For each return level, we identify all price intervals where the strategy’s payoff does not exceed z —these are the sub-level sets $\{x : p_c(x) \leq z\}$, which may consist of multiple disjoint regions when the performance curve is non-monotone. We then integrate the probability density $\mu(x)$ over all these intervals using polynomial spline arithmetic, accumulating the cumulative probability for this return level. Finally, these (return, probability) pairs are inverted and fit to a spline to construct the risk profile.

The optimization constraint. The risk tolerance profile $\hat{r}_{tol}(\lambda)$ is the user’s minimum acceptable return at each quantile. The optimization engine enforces:


$$\hat{r}_c(\lambda) \geq \hat{r}_{tol}(\lambda) \quad \forall \lambda \in [0, 1]$$

This ensures that at *every* probability level—from the worst tail to the best—the strategy meets or exceeds the user’s requirements.

5 Probability Distributions Over Time (In Development)

For dynamic strategies, we model beliefs as a family of distributions $\{\mu_i(x)\}_{i \in [n]}$ at times $\{d_i\}$, then interpolate continuously using linear transport:

$$\nu(x, d) = \frac{d_{i+1} - d}{d_{i+1} - d_i} \mu_i(x) + \frac{d - d_i}{d_{i+1} - d_i} \mu_{i+1}(x), \quad d \in [d_i, d_{i+1}] \quad (6)$$



Grayscale heatmap of a probability distribution evolving over time.
(This is my personal belief of \$TSLA share price development as per 2024)

Part III

Using CallCulator

6 Basic Calculation Mode

6.1 Workflow

1. **Enter a Ticker:** Type a stock symbol (e.g., TSLA) and press Enter. The system fetches the current price, company name, FIGI code, and full options chain via the backend Data Manager.
2. **Select an Expiration Date:** Choose from the list of available expirations.
3. **Build Your Strategy:** Add legs (Buy Call, Buy Put, Sell Call, Sell Put) with strike prices and weights (percentages representing capital allocation, e.g., 60% + 40%).
4. **Compute:** The backend `BasicCalcs` Lambda calculates the full analysis.

6.2 Results Displayed

- **Strategy Performance Chart:** A graph showing your percentage gain/loss (y-axis) at every possible stock price (x-axis) at expiration.
- **Black-Scholes Table:** A grid of estimated option values across time (columns) and stock price (rows), using the Black-Scholes PDE solution. Color-coded: green = profit, red = loss.

The table axes are:

- **X-axis (Time):** Time progresses from left (today) to right (expiration date).
- **Y-axis (Price):** Possible share prices of the underlying stock.

The cell at coordinate (x, y) shows the estimated P/L of your strategy at that specific future time and stock price. The rightmost column (at expiration) corresponds exactly to the values in the Strategy Performance Chart above it (just rotated 90 degrees). Values at expiration are certain (intrinsic value), while values earlier in time are theoretical estimates based on the Black-Scholes model.

- **Monte-Carlo CAGR Statistics:** Simulated compound annual growth rate assuming repeated investment with the given probability distribution.
- **Action Buttons:** Share (generates public link), Track, Invest (in development).

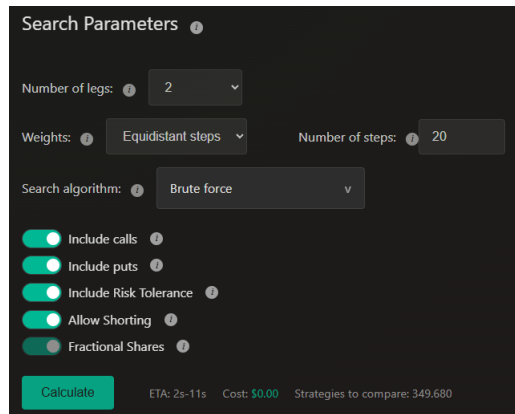
Note on Black-Scholes realism: The standard Black-Scholes model assumes constant implied volatility (IV) per option leg. In reality, IV changes as the stock price moves. CallCulator offers a **Dynamic IV** mode (enabled via a toggle in Search Parameters) that models this effect. See Section 10.2 for details.

7 Probability & Risk Mode

This is CallCulator's flagship mode. Instead of manually defining a strategy, you define your *beliefs* and let the optimization engine find the best strategy.

7.1 Workflow

1. **Enter a Ticker:** Same as Basic mode.
2. **Select an Expiration Date:** The system pre-fetches all implied PDFs in a single backend call (/impliedPDFAll).
3. **Define Your Probability Distribution:** Adjust the probability curve $\mu(x)$ to match your belief about where the stock will be at expiration (see figure below).
4. **Set Risk Tolerance:** Define a continuous tolerance function $\hat{r}_{tol}(\lambda)$ mapping each cumulative probability λ to the minimum acceptable return at that level (e.g., “in the worst 40% of outcomes, I require at least -5% ”).
5. **Configure Search Parameters:** Choose the number of **legs** (option positions per strategy), the number of **steps** (granularity of the strike/weight grid), and the search algorithm. Currently a **Brute Force** scan is used; smarter algorithms (Quasi-Monte Carlo / Sobol seeding, adaptive hill-climbing) are in development. Complexity grows rapidly with legs and steps, affecting computation time and cost.



Search parameter configuration panel.

6. **Launch Search:** The backend Search Manager distributes the optimization across multiple Lambda worker nodes, each evaluating hundreds of millions of option combinations in parallel (up to 625 threads per node, ~ 80 calculations per thread).

7.2 Search Algorithms

CallCulator employs two distinct algorithmic approaches to solve the non-convex optimization problem of finding the best spread.

7.2.1 Brute Force

The Brute Force algorithm discretizes the search space (strike prices and weights) into a uniform grid and evaluates every single valid combination.

- **Pros:** Guaranteed to find the global optimum within the resolution of the grid.
- **Cons:** Exponential time complexity $\mathcal{O}(N^k)$ where k is the number of legs. Infeasible for $k > 3$.

7.2.2 Hybrid Search (GSGD)

The **Grid Search Gradient Descent (GSGD)** algorithm is a hybrid global-local optimizer designed for high-dimensional spaces (4+ legs). It operates in two phases:

1. **Global Exploration (Sobol Seeding):** Instead of a uniform grid, the algorithm generates initial starting points (“seeds”) using **Sobol sequences**. These are quasi-random low-discrepancy sequences that cover the multi-dimensional search space more evenly than pseudo-random numbers, avoiding clustering and gaps.
2. **Local Exploitation (Gradient Descent):** From each seed, the algorithm performs a gradient descent optimization to “climb the hill” of expected return. It uses **Powell-Wolfe step rules** to dynamically adjust the step size, ensuring that each update provides a sufficient decrease in the objective function (or increase in return) while satisfying curvature conditions.

This approach allows CallCulator to efficiently navigate the complex, multi-modal landscape of option strategies, finding high-quality solutions in spaces that are too vast for brute force.

7.2.3 Performance & Complexity

- **Brute Force:** Computational cost explodes exponentially with complexity ($\mathcal{O}(S^N)$). Great for 2-3 legs, impossible for 4+.
- **Hybrid Search (GSGD):** Computational cost scales linearly with complexity ($\mathcal{O}(K \cdot N)$). Feasible for 4-5+ legs. The runtime depends primarily on the number of seeds (K) chosen by the user, not the number of legs.
- **Mathematical Construction:** The theoretical “Insurer + Sniper” solution (Regime 2) is constructed analytically in $\mathcal{O}(M \log M)$ time by sorting the likelihood ratio. This is effectively instantaneous (milliseconds) regardless of strategy complexity, but produces an ideal continuous payoff that must be approximated by discrete options.

7.3 The Optimization Engine

The backend considers all 2- and 3-leg option spreads from the available chain. The optimization problem is formally defined as:

$$\max_{s \in \mathcal{S}} \mathbb{E}_P[g_s] = \max_{s \in \mathcal{S}} \int_0^\infty g_s(x) P(x) dx \quad (7)$$

Subject to the risk constraint:

$$\forall q \in [0, 1] : R_s^{-1}(q) \geq T(q) \quad (8)$$

where $T(q)$ is the user’s risk tolerance profile (minimum return at quantile q), and $R_s(y)$ is the cumulative distribution function of the strategy’s return:

$$R_s(y) = \int_{\{x | g_s(x) \leq y\}} P(x) dx \quad (9)$$

Expanding the strategy s as a weighted sum of options $o_i \in \mathcal{O}$:

$$\max_{\substack{o_i \in \mathcal{O} \\ w_i \in \mathbb{R} \\ \sum w_i = 1}} \int_0^\infty \left(\sum_i w_i \cdot g_{o_i}(x) \right) P(x) dx \quad (10)$$

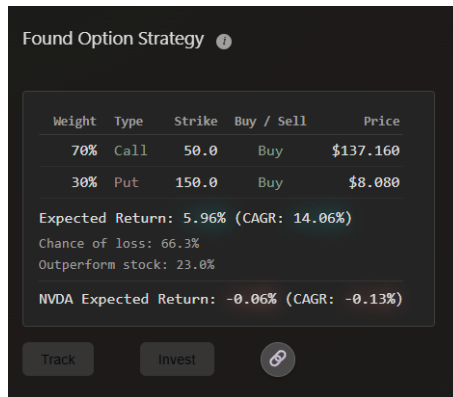
For each candidate:

1. Compute the performance function $g(S_T)$ using spline interpolation.
2. Compute the expected return: $\mathbb{E}_\mu[g] = \int g(x) \mu(x) dx$.
3. Compute the risk profile and verify it satisfies the risk tolerance constraint.
4. If valid, compare against the current best strategy.


Weights represent **capital allocation** (percentage of portfolio). The backend derives option quantities as weight/option_price.

7.4 Results Displayed

All results from Basic mode, plus:



Found Option Strategy card with legs, weights, expected return, and key statistics.




Strategy Performance Analysis: percentage return across all possible stock prices.

| Time | 2026 | | | | | | | | | | | | | | | | | | | | |
|-------|-------|--------|--------|-------|--------|--------|--------|-------|--------|--------|--------|-------|--------|--------|--------|-------|--------|--------|-------|--------|--------|
| Price | Feb 7 | Feb 15 | Feb 23 | Mar 3 | Mar 11 | Mar 19 | Mar 27 | Apr 4 | Apr 12 | Apr 20 | Apr 28 | May 6 | May 14 | May 22 | May 31 | Jun 8 | Jun 16 | Jun 24 | Jul 2 | Jul 10 | Jul 18 |
| 301 | 251 | 258 | 295 | 293 | 291 | 289 | 287 | 286 | 285 | 284 | 283 | 283 | 282 | 282 | 281 | 281 | 281 | 281 | 280 | 280 | 280 |
| 291 | 268 | 261 | 260 | 257 | 255 | 252 | 250 | 249 | 247 | 246 | 245 | 244 | 244 | 243 | 243 | 243 | 242 | 242 | 242 | 242 | 242 |
| 286 | 232 | 223 | 225 | 223 | 219 | 216 | 214 | 212 | 210 | 208 | 207 | 206 | 206 | 205 | 205 | 204 | 204 | 204 | 204 | 204 | 204 |
| 278 | 193 | 194 | 190 | 187 | 183 | 180 | 177 | 175 | 173 | 170 | 168 | 167 | 167 | 166 | 166 | 166 | 166 | 166 | 165 | 165 | 165 |
| 271 | 166 | 162 | 157 | 153 | 149 | 145 | 142 | 139 | 136 | 134 | 132 | 131 | 130 | 129 | 128 | 128 | 128 | 127 | 127 | 127 | 127 |
| 263 | 136 | 130 | 125 | 120 | 115 | 111 | 107 | 104 | 100 | 97 | 95 | 93 | 91 | 90 | 89 | 89 | 89 | 88 | 88 | 88 | 88 |
| 256 | 106 | 100 | 94 | 88 | 83 | 78 | 73 | 69 | 65 | 61 | 58 | 56 | 54 | 52 | 51 | 51 | 50 | 50 | 50 | 50 | 49 |
| 248 | 79 | 71 | 64 | 58 | 52 | 47 | 41 | 35 | 31 | 26 | 23 | 20 | 17 | 15 | 13 | 12 | 12 | 12 | 11 | 11 | 11 |
| 241 | 54 | 46 | 38 | 32 | 24 | 17 | 10 | 4 | -1 | -6 | -11 | -15 | -19 | -21 | -23 | -25 | -26 | -26 | -26 | -26 | -26 |
| 233 | 30 | 22 | 14 | 6 | -2 | -10 | -17 | -22 | -26 | -31 | -35 | -40 | -44 | -47 | -49 | -50 | -50 | -50 | -50 | -50 | -50 |
| 226 | 13 | 3 | -4 | -10 | -16 | -22 | -28 | -33 | -37 | -41 | -45 | -49 | -52 | -54 | -55 | -55 | -55 | -55 | -55 | -55 | -55 |
| 218 | -17 | -26 | -34 | -42 | -48 | -53 | -56 | -58 | -59 | -60 | -60 | -60 | -60 | -60 | -60 | -60 | -60 | -60 | -60 | -60 | -60 |
| 211 | -38 | -48 | -56 | -63 | -68 | -71 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 | -72 |
| 203 | -63 | -74 | -82 | -88 | -92 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 | -94 |
| 196 | -92 | -104 | -112 | -117 | -120 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 | -121 |
| 188 | -124 | -137 | -145 | -149 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 | -150 |
| 180 | -160 | -174 | -178 | -178 | -177 | -175 | -172 | -168 | -163 | -157 | -150 | -142 | -134 | -125 | -116 | -106 | -95 | -83 | -71 | -59 | -47 |
| 173 | -202 | -218 | -222 | -221 | -218 | -213 | -207 | -200 | -192 | -183 | -173 | -162 | -150 | -138 | -125 | -111 | -97 | -83 | -69 | -55 | -41 |
| 165 | -250 | -268 | -271 | -268 | -262 | -255 | -247 | -238 | -228 | -217 | -205 | -192 | -178 | -164 | -149 | -134 | -118 | -102 | -86 | -70 | -54 |
| 158 | -304 | -324 | -326 | -320 | -312 | -303 | -293 | -282 | -270 | -257 | -243 | -228 | -212 | -195 | -178 | -160 | -141 | -122 | -103 | -84 | -65 |
| 150 | -366 | -388 | -389 | -380 | -369 | -357 | -344 | -330 | -315 | -299 | -282 | -264 | -245 | -225 | -204 | -182 | -158 | -133 | -107 | -80 | -52 |
| 143 | -436 | -460 | -460 | -450 | -438 | -425 | -411 | -395 | -378 | -359 | -339 | -318 | -295 | -270 | -243 | -215 | -185 | -153 | -120 | -86 | -51 |
| 135 | -514 | -540 | -540 | -530 | -517 | -503 | -487 | -469 | -449 | -427 | -404 | -379 | -352 | -323 | -293 | -261 | -228 | -194 | -159 | -124 | -88 |
| 128 | -600 | -628 | -628 | -618 | -604 | -588 | -570 | -551 | -530 | -507 | -482 | -455 | -426 | -395 | -362 | -327 | -291 | -254 | -217 | -180 | -143 |
| 120 | -694 | -724 | -724 | -714 | -700 | -683 | -664 | -644 | -622 | -598 | -572 | -544 | -514 | -482 | -448 | -412 | -374 | -335 | -296 | -257 | -218 |
| 113 | -796 | -828 | -828 | -818 | -804 | -786 | -766 | -744 | -720 | -694 | -666 | -636 | -604 | -570 | -534 | -496 | -457 | -417 | -376 | -334 | -292 |
| 105 | -906 | -940 | -940 | -930 | -916 | -897 | -876 | -853 | -828 | -801 | -772 | -741 | -708 | -674 | -638 | -599 | -559 | -518 | -476 | -433 | -390 |
| 98 | -1024 | -1060 | -1060 | -1050 | -1036 | -1017 | -996 | -972 | -946 | -918 | -888 | -856 | -822 | -786 | -748 | -709 | -669 | -628 | -586 | -543 | -500 |
| 90 | -1150 | -1190 | -1190 | -1180 | -1166 | -1147 | -1126 | -1102 | -1076 | -1048 | -1018 | -986 | -952 | -916 | -878 | -839 | -799 | -758 | -716 | -673 | -630 |
| 83 | -1284 | -1328 | -1328 | -1318 | -1304 | -1285 | -1264 | -1240 | -1214 | -1186 | -1156 | -1124 | -1090 | -1054 | -1016 | -977 | -937 | -896 | -854 | -811 | -768 |
| 75 | -1426 | -1474 | -1474 | -1464 | -1450 | -1431 | -1409 | -1384 | -1358 | -1330 | -1299 | -1266 | -1231 | -1194 | -1156 | -1117 | -1077 | -1036 | -994 | -951 | -908 |
| 68 | -1576 | -1628 | -1628 | -1618 | -1604 | -1585 | -1563 | -1538 | -1512 | -1484 | -1454 | -1421 | -1386 | -1349 | -1311 | -1272 | -1232 | -1191 | -1149 | -1106 | -1063 |
| 60 | -1734 | -1790 | -1790 | -1780 | -1766 | -1747 | -1724 | -1698 | -1672 | -1644 | -1614 | -1581 | -1546 | -1509 | -1471 | -1432 | -1392 | -1351 | -1309 | -1266 | -1223 |
| 53 | -1900 | -1960 | -1960 | -1950 | -1936 | -1917 | -1894 | -1868 | -1841 | -1812 | -1781 | -1748 | -1713 | -1676 | -1638 | -1599 | -1559 | -1517 | -1474 | -1431 | -1388 |
| 45 | -2074 | -2138 | -2138 | -2128 | -2114 | -2095 | -2072 | -2046 | -2019 | -1990 | -1959 | -1926 | -1891 | -1854 | -1816 | -1777 | -1737 | -1695 | -1652 | -1609 | -1566 |
| 38 | -2256 | -2324 | -2324 | -2314 | -2300 | -2281 | -2258 | -2232 | -2205 | -2176 | -2145 | -2112 | -2077 | -2040 | -2002 | -1963 | -1923 | -1881 | -1838 | -1794 | -1751 |
| 30 | -2446 | -2518 | -2518 | -2508 | -2494 | -2475 | -2452 | -2426 | -2399 | -2370 | -2339 | -2306 | -2271 | -2234 | -2196 | -2157 | -2117 | -2075 | -2032 | -1988 | -1945 |
| 23 | -2644 | -2720 | -2720 | -2710 | -2696 | -2677 | -2654 | -2628 | -2599 | -2568 | -2535 | -2500 | -2463 | -2425 | -2386 | -2346 | -2305 | -2262 | -2218 | -2173 | -2128 |
| 15 | -2850 | -2930 | -2930 | -2920 | -2906 | -2887 | -2864 | -2838 | -2809 | -2778 | -2745 | -2710 | -2673 | -2635 | -2596 | -2556 | -2515 | -2472 | -2427 | -2381 | -2335 |
| 8 | -3064 | -3148 | -3148 | -3138 | -3124 | -3105 | -3082 | -3056 | -3027 | -2996 | -2963 | -2928 | -2891 | -2853 | -2814 | -2774 | -2732 | -2689 | -2644 | -2597 | -2550 |
| 0 | -3286 | -3374 | -3374 | -3364 | -3350 | -3331 | -3308 | -3282 | -3253 | -3222 | -3189 | -3154 | -3117 | -3079 | -3040 | -3000 | -2958 | -2914 | -2868 | -2821 | -2773 |

Black-Scholes table. The blue glow on the right edge represents the user's Assumed Probability Density: brighter glow indicates higher probability of the stock ending at that price.

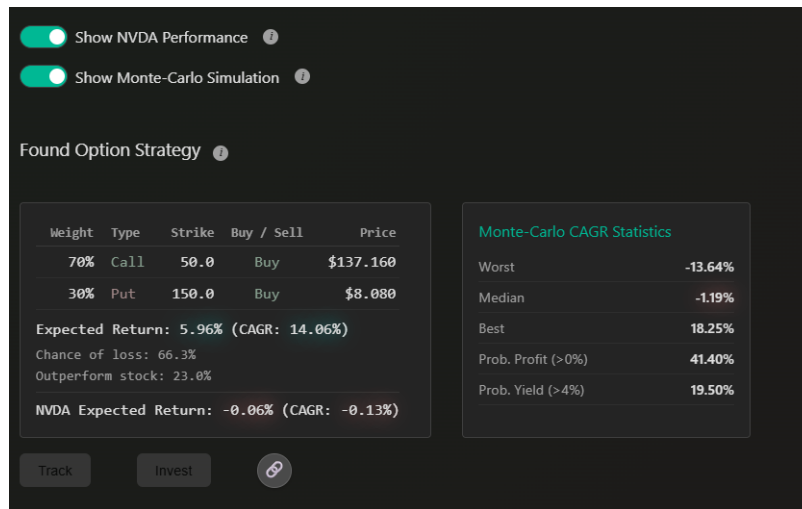
- **Found Option Strategy:** The optimal spread with its legs, percentage weights, and expected return.
- **Risk Profile Comparison:** Overlay of the user's risk tolerance \hat{r}_{tol} (red) and the strategy's actual risk profile \hat{r}_c (blue).



The blue line (strategy) strictly dominates the red line (tolerance).

- **Underlying Performance Toggle:** “Show {ticker} Performance” reveals the stock's expected return and risk profile for comparison.
- **Outperform Probability:** The probability that the found strategy outperforms simply holding the stock, computed via spline intersections with the PDF.

- **Chance of Loss:** The probability that the strategy results in a net loss.
- **Monte-Carlo CAGR Statistics:** Simulated compound annual growth rate—see Risk Theory (Part IV) for why this matters beyond expected return.



Found strategy card with Monte-Carlo CAGR statistics.

8 Max Min Mode

Max Min takes a fundamentally different approach to portfolio construction. Instead of maximizing expected return (which depends on the accuracy of a user-supplied probability distribution), it maximizes the **worst-case outcome** (minimax optimization):

$$\max_{\mathbf{w}} \min_{S_T \in [0, \infty)} \sum_i w_i g_i(S_T)$$

where \mathbf{w} is the weight vector over available option legs.

This is a **distribution-free** criterion: it makes no assumption about where the stock will end up. It answers the question: “If I had to prepare for the absolute worst, what is the strategy whose floor is the highest?”

8.1 Single Stock Mode

8.1.1 Workflow

1. **Enter a Ticker:** Same data pipeline as the other modes (price, options chain, implied PDFs).
2. **Select an Expiration Date:** Choose from the available expirations list.
3. **Launch Search:** The optimizer evaluates all spread combinations and finds the one whose minimum across all possible stock prices S_T is maximized.

8.1.2 Mathematical Formulation

For a single stock with available option payoff functions $\{g_i\}_{i \in [n]}$, the optimization problem is:

$$\max_{s \in \mathcal{S}} \min_{x \in \mathbb{R}} g_s(x) = \max_{\substack{o_i \in \mathcal{O} \\ w_i \in \mathbb{R} \\ \sum w_i = 1}} \min_{x \in \mathbb{R}} \sum_i w_i \cdot g_{o_i}(x) \quad (11)$$

Because each g_i is piecewise linear (for vanilla options), the minimum over S_T occurs at a vertex of the piecewise-linear surface. The problem reduces to a finite-dimensional linear program.

8.2 Multi Stock Mode

Multi mode extends the minimax framework across multiple tickers. The user builds a portfolio by adding stocks one at a time, each with its own expiration and option chain.

8.2.1 Workflow

1. **Add Stocks:** Enter a ticker and click “Add.” Repeat for each stock.
2. **Select Expirations:** Choose an expiration date per stock.
3. **Launch Search:** The optimizer finds both the cross-asset allocation *and* the per-stock option spread that maximizes the portfolio’s worst case.

8.2.2 Cross-Asset Diversification

With M stocks, each having payoff functions $\{g_{m,i}\}$, the problem becomes:

$$\max_{\mathbf{w}_1, \dots, \mathbf{w}_M, \alpha \in \Delta^M} \min_{S_{T,1}, \dots, S_{T,M}} \sum_{m=1}^M \alpha_m \sum_i w_{m,i} g_{m,i}(S_{T,m})$$

where α is the inter-stock capital allocation and \mathbf{w}_m is the intra-stock option spread. Because the stocks are independent, the global minimum decomposes into per-stock minima, and the outer allocation α optimally diversifies across the per-stock floors.

Feature Isolation: Actions in Single mode (selecting tickers, dates) never affect the Multi mode state, and vice versa. This is enforced at the UI level.

Note: The Max Min optimization algorithm is currently in development. The UI for both sub-modes is complete and functional.

8.3 Soft Minimum Optimization (Percentile-Based Targets)

While pure Max Min optimizes the absolute worst-case outcome (the 0th percentile), users may prefer to relax this extreme conservatism and instead optimize a higher percentile—such as the 5th or 10th percentile. This approach is known as **Soft Minimum maximization**, **Max Min Soft** or **Value-at-Risk (VaR) optimization**.

8.3.1 When to Use Soft Minimum

- **Acceptable Risk Tolerance:** You are willing to accept a 5% chance of outcomes below a threshold to improve average performance for the other 95% of scenarios. - **Realistic Worst-Case Budgets:** The absolute worst-case floor may be so expensive (or impossible) to guarantee that you prefer to focus on a more attainable probabilistic floor. - **Long-Term Investing:** For multi-period strategies, the 5th percentile often provides more meaningful protection than the theoretical minimum (which might represent a scenario with 0.001% probability).

8.3.2 How It Works

Instead of optimizing $\max(\min g)$, soft minimum optimization solves:

$$\max_{\mathbf{w}} (\text{Value-at-Risk}_{p^*}[g])$$

where p^* is the chosen percentile level (e.g., $p^* = 0.05$ for the 5th percentile). This means: “Maximize the return level such that 5% of scenarios breach below it.”

Under the hood:

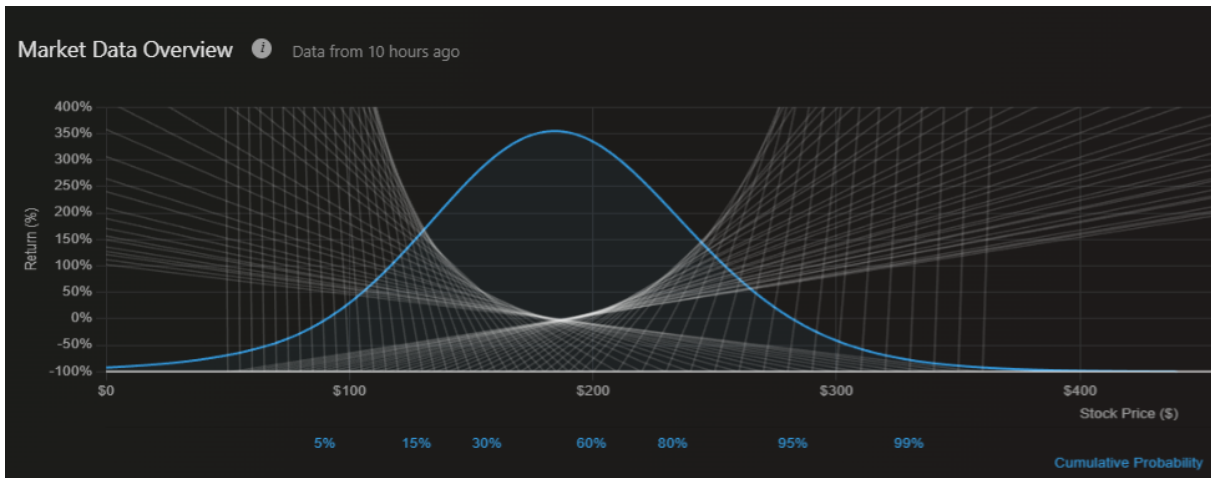
1. The optimizer identifies the worst $p^* \times 100\%$ of stock price outcomes (the “sacrifice zone”).
2. Within this zone, the strategy can yield any payoff (potentially zero or negative, even bankruptcy if not protected against it).
3. For the remaining $100(1 - p^*)\%$ of scenarios, the optimizer maximizes the floor.
4. The freed-up budget from sacrificing the worst outcomes can significantly improve protection for the preserved percentiles.

8.3.3 Hybrid Approach: VaR Floor with Growth

Users can also combine a soft-minimum floor with upside potential. Instead of capping the protected percentiles at a flat floor, the optimizer can use the Farmer approach (from Probability & Risk mode) to create a curved payoff that still respects the VaR floor. This results in: a guaranteed 5th percentile return, with probability-weighted growth opportunities above that floor.

9 Market Data Overview

When a ticker and expiration are selected, the Market Data Overview chart is displayed:



Market Data Overview: thin white lines show individual option performance curves; the colored curve is the Implied PDF (peak \sim \$180); the CDF is overlaid.

10 Implied Probability Distributions

10.1 What the Market “Thinks”

The Implied PDF is derived from current option prices using the following pipeline:

1. Compute implied volatility (IV) for each out-of-the-money option using the Black-Scholes formula.
2. Fit a **Quadratic Curve** to IV as a function of strike: $IV(K) = aK^2 + bK + c$.
3. Use the smoothed IV to generate dense Black-Scholes prices $C(K)$ across a fine strike grid.

4. Derive the PDF via the **Breeden-Litzenberger** finite-difference formula:

$$\mu(K) \approx e^{rT} \frac{C(K + \Delta K) - 2C(K) + C(K - \Delta K)}{(\Delta K)^2} \quad (12)$$

This **IV Smoothing** method avoids the instabilities (Runge’s phenomenon) of fitting splines directly to noisy option prices.

10.2 Dynamic Implied Volatility

The standard Black-Scholes model uses a *constant* IV per option leg, calibrated from current market prices. This “static IV” assumption means the model predicts the same volatility regardless of where the stock moves. In reality, implied volatility varies with the stock price — a phenomenon known as the **volatility smile** or **volatility skew**.

CallCulator’s **Dynamic IV** feature models this by adjusting each leg’s IV based on the stock’s displacement from its current price:

$$\sigma^{\text{dyn}}(S) = \sigma_{\text{base}} \times (1 + \alpha \cdot \Delta + \beta \cdot \Delta^2) \quad (13)$$

where $\Delta = (S - S_0)/S_0$ is the relative price change, α is the **skew** parameter, and β is the **curvature** parameter.

Why does IV change with price? Three effects drive the relationship between stock price and IV:

1. **Crash fear and put demand (Skew):** When stocks fall, investors rush to buy protective puts, driving up option prices and thus IV. When stocks rise, this panic subsides and IV compresses. This directional asymmetry is captured by the skew parameter ($\alpha < 0$ for equities).
2. **Leverage effect (Skew):** A falling stock price increases the company’s debt-to-equity ratio, making it objectively riskier. Higher risk translates to higher expected future volatility.
3. **Regime shifts (Curvature):** Any large move — up *or* down — signals that the market’s volatility regime may have changed. A stock that jumps 30% in either direction is clearly in a more volatile state than one that barely moves. This symmetric effect is captured by the curvature parameter ($\beta > 0$).

Practical examples:

- With $\alpha = -1, \beta = 0$: A 10% stock drop increases IV by 10% of its base value (e.g., 30% → 33%). A 10% stock rise *decreases* IV by 10% (30% → 27%).
- With $\alpha = -1, \beta = 0.5$: A 10% drop increases IV by 10.5% (skew + curvature). A 10% rise decreases IV by only 9.5% (skew partially offset by curvature).
- With $\alpha = 0, \beta = 1$: Pure smile — IV increases symmetrically for any large move.

Initial condition guarantee: At the current stock price ($S = S_0$), $\Delta = 0$, so $\sigma^{\text{dyn}} = \sigma_{\text{base}}$. This means the Black-Scholes table always shows exactly 0% performance at today’s price and time, regardless of the skew and curvature settings.

Usage in CallCulator: Enable the **Dynamic IV** toggle in Search Parameters and set the **Skew** value (typical equity values: -2 to -0.5). The Black-Scholes table will then reflect how your strategy performs under a more realistic volatility surface. Use the **Show IV X-ray** button (hold to reveal) above the table to see the effective IV at each cell.

11 Data Pipeline

- **Data Manager:** Serves market data to the frontend. Checks S3 data freshness and triggers updates if stale.
- **Data Updater:** Fetches live data from the Alpaca API when triggered.
- **Caching:** Three-layer architecture: S3 (backend), CloudFront (edge), Browser (30-second window).
- **Implied PDFs stay in sync:** Any options refresh automatically triggers a recalculation of all implied PDFs for that ticker.
- **Thundering Herd Protection:** Uses `isBeingUpdated` flag + Probabilistic Early Expiration (PEE) to prevent cache stampedes.

Part IV

Risk Theory

12 Risk Profiles

Every investment carries risk. CallCulator quantifies risk through two mathematical objects:

12.1 Personal Risk Tolerance Profile

A function $\hat{r}_{tol} : [0, 1] \rightarrow \mathbb{R}$ mapping cumulative probability λ to the minimum acceptable return at that probability level.

Example: “I never want to lose more than 25% (at $\lambda = 0$), in 50% of cases I want to profit (at $\lambda = 0.5$, return ≥ 0), and my best case should be at least +100% (at $\lambda = 1$).”

12.2 Investment Risk Profile

For a strategy with performance function $p_c(x)$ and assumed PDF $\mu(x)$, the risk profile is:

$$\hat{r}_c(\lambda) = p_c(\Psi^{-1}(\lambda)), \quad \Psi(x) := \int_0^x \mu(y) dy \quad (14)$$

where Ψ is the cumulative distribution function. This maps each cumulative probability level $\lambda \in [0, 1]$ to the corresponding return of the strategy.

12.3 Matching Profiles

A strategy is **valid** if its risk profile dominates the tolerance at every point:

$$\hat{r}_c(\lambda) \geq \hat{r}_{tol}(\lambda) \quad \forall \lambda \in [0, 1]$$

The optimization engine searches for the valid strategy with the highest expected return.

13 On the Merit of Spreads

13.1 The Single-Step Argument Against Spreads

For a single, isolated investment, the expected return of a spread is:

$$\mathbb{E}\left[\sum_i w_i g_i\right] = \sum_i w_i \mathbb{E}[g_i] \quad (15)$$

Since this is a convex combination of real numbers, its maximum occurs at the boundary ($w_j = 1$ for some j). Therefore, in a single-step game, a single option always has the highest expected return. Spreads appear pointless.

13.2 The Multi-Step Counter-Argument: Geometric Growth

Investing is not a single event—it is a **repeated game**. In repeated games, maximizing the arithmetic expectation can lead to guaranteed ruin.

13.2.1 Alice and Bob: A Story of Ruin

Alice offers Bob a coin (80% Heads, 20% Tails):

- Heads: +100% (investment doubles)
- Tails: −100% (investment lost)

Rule: Bob must reinvest his entire bankroll.

Expected return per flip: +60%. But if Bob plays long enough:

$$P(\text{at least one Tails in } n \text{ flips}) = 1 - 0.8^n \xrightarrow{n \rightarrow \infty} 1$$

Bob is bankrupt with probability 1. The arithmetic expectation (+60%) is misleading; the **geometric growth rate** is $\ln(0) = -\infty$ on a Tails outcome.

13.2.2 Position Sizing Saves the Day

By betting only a fraction f of his capital, Bob's geometric growth rate per flip becomes:

$$G(f) = 0.8 \ln(1 + f) + 0.2 \ln(1 - f)$$


Maximizing $G(f)$ yields the **Kelly fraction**—the optimal bet size.

13.3 Beyond the Coin Flip: The Ergodicity Problem

This coin-flip example might seem extreme, but the core issue persists in almost all repeated investing. The mathematical term for this is **ergodicity breaking**.

In financial markets, the **Expected Return** (average of all possible outcomes) is often very different from the **Median Outcome** (what actually happens to most people). A strategy might have a positive Expected Return but still lead to a loss for 90% of investors over time because a few massive winners skew the average.

When you invest repeatedly (re-investing your gains), you don't experience the average of parallel universes; you experience one single path through time. If that path hits zero, you are out of the game. This is why maximizing Expected Return is not enough—you must manage the risk of the "time-average" diverging from the "ensemble-average."



Simulation of average return per flip vs. fraction of capital wagered.

13.4 Monte-Carlo Simulation: Quantifying Hidden Risk

Expected return alone is an insufficient metric for investment decisions in a repeated-game context. Two strategies with identical $\mathbb{E}[g]$ can have vastly different geometric growth rates depending on the variance and tail behavior of their return distributions.

A **Monte-Carlo CAGR simulation** would expose this by random sampling:

1. Draw N samples $\{x_j\}_{j=1}^N$ from the user's PDF $\mu(x)$.
2. For each sample, compute the strategy return $g(x_j)$.
3. Simulate T rounds of compounding: $W_T = W_0 \prod_{t=1}^T (1 + g(x_{\sigma(t)}))$ where σ is a random draw.
4. Report the annualized geometric mean: $\text{CAGR} = (W_T/W_0)^{1/T} - 1$, averaged over many runs.

However, Monte-Carlo is stochastic: each run gives slightly different results, tail events are under-sampled, and convergence is slow ($\mathcal{O}(1/\sqrt{K})$ for K paths). CallCulator replaces this with a **deterministic** alternative.

13.5 FFT Convolution: Deterministic CAGR Distributions

Instead of sampling random paths, CallCulator computes the **exact** compound return distribution by exploiting the Convolution Theorem on log-return probability densities.

13.5.1 Core Insight: Quantile Wealth Multipliers

The strategy's risk profile is a **return quantile function** $Q : [0, 1] \rightarrow \mathbb{R}$, mapping each percentile p to the return at that probability level. From this, we define the **wealth multiplier quantile function**:

$$W(p) = 1 + Q(p), \quad p \in [0, 1] \tag{16}$$

Each percentile gives a *different* wealth multiplier: $W(0)$ is the worst-case multiplier (e.g. 0.6 if the worst loss is -40%), while $W(1)$ is the best-case (e.g. 3.0 for $+200\%$).

When we compound over N independent periods, the total wealth is the product of the single-period multipliers. Taking the logarithm converts this product into a sum:

$$\mathcal{W}_N = \prod_{t=1}^N W_t \quad \implies \quad \ln \mathcal{W}_N = \ln \left(\prod_{t=1}^N W_t \right) = \sum_{t=1}^N \ln W_t \quad (17)$$

This transforms the problem of multiplying random variables into the simpler problem of adding them.

The probability distribution of a sum of independent random variables is the ****convolution**** of their individual distributions. (Intuition: to get a sum $Z = X + Y$, for every value x that X can take, Y must take exactly $Z - x$; summing these mutually exclusive possibilities weighted by their probabilities is the definition of convolution).

In the frequency domain (Fourier space), this complex convolution operation becomes simple pointwise multiplication:

$$\widehat{p_L^{*N}} = (\widehat{p_L})^N \quad (18)$$

where p_L is the single-period log-return density. To compute this efficiently, we approximate the continuous density with a discrete histogram and use the Fast Fourier Transform (FFT).

13.5.2 Protocol

Given the strategy's return quantile function Q , the number of compounding periods N , and a reinvestment fraction f :

1. **Discretize** the quantile function at M bin-center percentiles to obtain discrete quantile returns: $r_k = Q\left(\frac{k+0.5}{M}\right)$.
2. **Form quantile wealth multipliers**: $W_k = 1 + r_k/100$. Each W_k corresponds to a different percentile of the outcome distribution.
3. **Apply partial reinvestment**: $W_k^{(f)} = (1 - f) + f \cdot W_k$ and clip $W_k^{(f)} > 0$. Only a fraction f of capital is exposed each period.
4. **Log-transform** to discrete quantile log-returns: $\ell_k = \ln W_k^{(f)}$.
5. **Build a normalized histogram PDF** of the log-returns over power-of-2 bins. Since the quantile samples are uniformly spaced in probability, each contributes equal mass to the histogram.
6. **Zero-pad** to length $\geq N \cdot N_b$ to prevent circular wraparound.
7. **Forward FFT** (Cooley-Tukey radix-2), raise spectrum to power N , **inverse FFT**.
8. **Map convolved bins to CAGR**: $\text{CAGR}_m = \exp(S_m/N) - 1$ where S_m is the total log-return at bin m .
9. **Build CDF** to produce the final CAGR quantile function—a mapping from percentile to compound growth rate, describing N -period compounded outcomes.

13.5.3 Advantages over Monte-Carlo

- **Deterministic**: Identical inputs always produce identical outputs.
- **Exact tails**: Unlike Monte-Carlo which under-samples rare events, the FFT method resolves the full distribution uniformly—critical for quantifying ruin risk.
- **Speed**: $\mathcal{O}(M \log M)$ regardless of N , compared to $\mathcal{O}(K \cdot N)$ for K Monte-Carlo paths.
- **Multi- N animation**: CallCulator computes profiles for $N = 1, 2, \dots, 20$, visualizing how the CAGR quantile function concentrates over time.

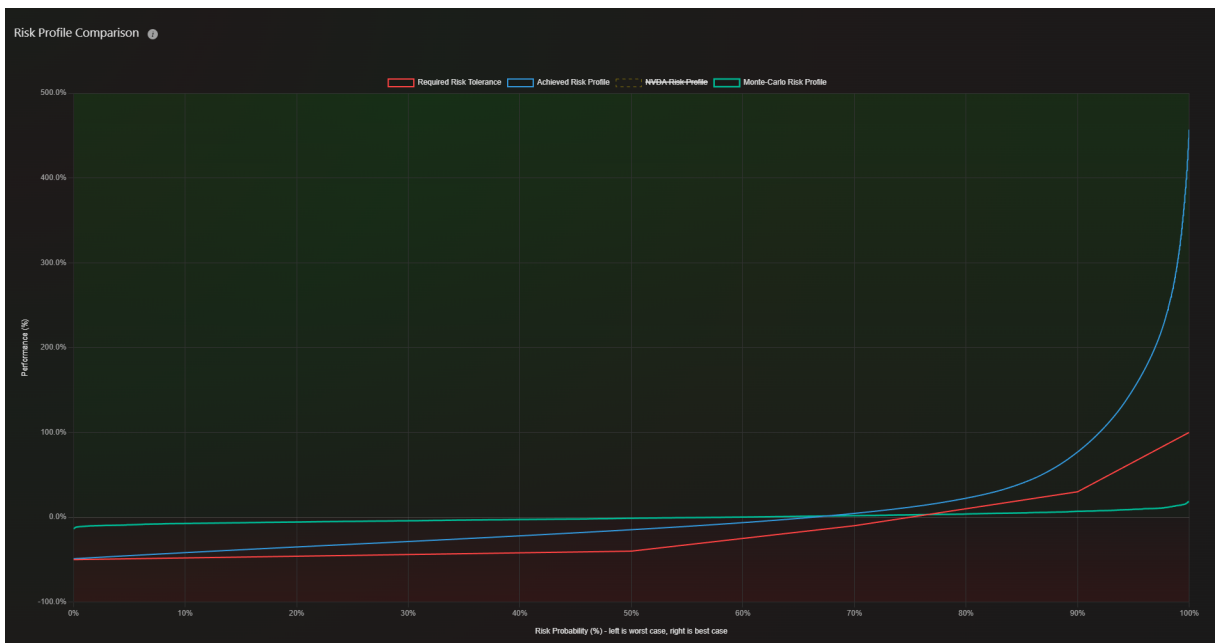
13.6 Kelly Criterion: Optimal Reinvestment Fraction

The FFT convolution requires a reinvestment fraction f . The **Kelly Criterion** finds the growth-optimal fraction f^* by maximizing the expected log-growth rate over the quantile returns:

$$f^* = \operatorname{argmax}_{f \in [0, f_{\max}]} G(f) = \operatorname{argmax}_f \frac{1}{M} \sum_{k=0}^{M-1} \ln(1 + f \cdot r_k) \quad (19)$$

where r_k are the single-period fractional returns sampled from the return quantile function. The upper bound $f_{\max} < -1/r_{\min}$ ensures solvency ($1 + f \cdot r_{\min} > 0$). Since $G(f)$ is concave, f^* is unique and is found via golden-section search.

CallCulator feeds f^* into the FFT convolution as the reinvestment parameter. The Kelly fraction is also displayed as an informational metric alongside the CAGR statistics, giving users insight into how much of their capital should be allocated to the strategy for optimal long-term growth.



Risk Profile Comparison with FFT-computed CAGR profiles.

13.7 Log-Optimal Strategy: Maximizing Median CAGR

The Probability & Risk optimization engine (Section 3.2) maximizes the **expected return** $\mathbb{E}_P[g]$ —a linear objective. As the Alice & Bob story and the ergodicity discussion demonstrate, expected return can be a misleading metric for repeated investment. A strategy with the highest expected return may have a median compound growth rate of zero.

An alternative objective is to maximize the **Median Monte-Carlo CAGR** directly. Mathematically, this is equivalent to maximizing the expected logarithm of wealth:

$$\max_g \int_0^\infty \ln(g(x)) P(x) dx \quad (20)$$

subject to the budget constraint $\int_0^\infty g(x) Q(x) dx = K$.

13.7.1 The Log-Optimal Payoff

Solving via the Lagrangian method yields the optimal payoff:

$$g^*(x) = \frac{1}{\lambda} \cdot \frac{P(x)}{Q(x)} \quad (21)$$

where λ is the Lagrange multiplier (budget normalizer) and $P(x)/Q(x)$ is the likelihood ratio $\Lambda(x)$.

This is fundamentally different from the expected-return-maximizing “Sniper” strategy:

- **Sniper (Max Expected Return):** Concentrates all capital at a single spike at $\operatorname{argmax} \Lambda(x)$. Highest average return, but median wealth of \$0 over time (most spikes miss).
- **Farmer (Max Median CAGR):** Spreads capital across the *entire* state space proportional to $\Lambda(x)$. Lower average return, but highest median wealth over time. The payoff never reaches zero, avoiding the ruin that destroys geometric growth.

13.7.2 Adding the Risk Tolerance Floor

When a risk tolerance $T(x)$ is imposed (minimum required payoff at each state x), the constrained log-optimal solution becomes:

$$g^*(x) = \max\left(T(x), \frac{1}{\lambda} \cdot \frac{P(x)}{Q(x)}\right) \quad (22)$$

This is a **threshold strategy**: below the threshold, you buy “insurance” (paying the minimum cost to satisfy T); above the threshold, you allocate according to the Kelly-optimal Λ -shaped profile. In financial terms, this resembles a CPPI (Constant Proportion Portfolio Insurance) strategy—a protective floor plus a growth allocation.

13.7.3 Computing λ : Newton-Raphson on the Budget Equation

The multiplier λ is found by solving the budget equation $f(\lambda) = K$, where:

$$f(\lambda) = \int_0^\infty \max\left(T(x), \frac{1}{\lambda} \cdot \frac{P(x)}{Q(x)}\right) Q(x) dx \quad (23)$$

The derivative is:

$$f'(\lambda) = -\frac{1}{\lambda^2} \int_{\Omega_\lambda} P(x) dx \quad (24)$$

where Ω_λ is the “active set” of states where the Kelly bet exceeds the floor. Since f is monotonically decreasing and convex, Newton-Raphson converges rapidly (typically 3–4 iterations). CallCulator’s polynomial spline infrastructure enables exact evaluation of both $f(\lambda)$ and $f'(\lambda)$ via coefficient-shifting integration, avoiding discretization noise entirely.

13.8 Implications for Options

A spread is precisely this: instead of betting everything on a single call option (maximizing single-step expected return), a spread allocates capital across multiple payoff profiles, avoiding absorbing barriers (total loss) and maximizing long-term geometric growth. This is why CallCulator’s risk-constrained optimization produces spreads—they are the geometrically optimal strategy for repeated investing.

14 Mathematical Constructions: The Six Optimization Regimes

The various optimization objectives available in CallCulator (and planned for future implementation) can be understood as distinct mathematical “regimes,” each with a characteristic optimal payoff structure. These regimes form a hierarchy from micro-arbitrage to macro-strategic allocation.

14.1 Regime 0: The Scavenger (Micro-Arbitrage)

Before constructing a global strategy based on your belief $P(x)$, the Scavenger exploits static mispricings within the market data itself—deviations between observed option prices and the fitted theoretical curve $Q(x)$.

14.1.1 Residual Mispricing

Let $\{(K_i, p_{mkt}^i)\}$ be the observed market prices and $C_{fit}(K)$ be the theoretical price curve from the IV Smoothing pipeline. The **residual** for each option is:

$$\epsilon_i = p_{mkt}^i - C_{fit}(K_i) \quad (25)$$

Strategy:

- **Buy** options with $\epsilon_i < -\theta$ (“cheap”)
- **Sell** options with $\epsilon_i > \theta$ (“expensive”)

where θ is a transaction cost threshold.

14.1.2 Butterfly Spreads and Neighbor Selection

To isolate a mispricing without taking directional risk, construct a butterfly spread. If option K_i is cheap, buy it and sell neighbors K_{i-m} and K_{i+n} to hedge.

Neighbor Selection Rule: Expand the window $[K_{i-m}, K_{i+n}]$ outward until you reach strikes where:

$$\text{sgn}(\epsilon_{i\pm k}) \neq \text{sgn}(\epsilon_i) \quad \text{or} \quad |\epsilon_{i\pm k}| < \delta \quad (26)$$

This ensures you hedge a “cheap” option against “fair” or “expensive” ones, maximizing residual capture.

14.1.3 Boosting Macro Strategies

The Scavenger’s role is not standalone—it **boosts** the macro regimes (1–5). When discretizing a theoretical payoff $g^*(x)$ into actual options, preferentially choose “cheap” outliers. This lowers the actual cost below the theoretical cost, freeing up budget to scale up the position or improve the floor.

14.2 Regime 1: The Sniper (Unconstrained Maximization)

Objective: Maximize expected return with a fixed budget K .

$$\max_g \int_0^\infty g(x)P(x) dx \quad \text{s.t.} \quad \int_0^\infty g(x)Q(x) dx = K \quad (27)$$

Optimal Payoff: A Dirac delta (infinitely sharp spike) at the point maximizing the likelihood ratio:

$$x^* = \text{argmax}_x \frac{P(x)}{Q(x)} \quad (28)$$

Interpretation: The Sniper bets everything on the single most mispriced outcome. In practice, this is approximated by a very tight butterfly spread centered at x^* . This strategy has the highest possible expected return but a median outcome of \$0 over repeated trials (most spikes miss).

14.3 Regime 2: The Insurer + Sniper (Constrained Optimization)

Objective: Maximize expected return subject to a risk tolerance floor $T(q)$.

$$\max_g \mathbb{E}_P[g] \quad \text{s.t.} \quad \forall q \in [0, 1] : R_g^{-1}(q) \geq T(q) \quad (29)$$

Optimal Payoff: A two-part portfolio:

$$g^*(x) = g_{safe}(x) + g_{spec}(x) \quad (30)$$

where:

- **Safety Portfolio:** The cheapest portfolio satisfying $T(q)$, constructed via the Hardy-Littlewood rearrangement:

$$g_{safe}(x) = T(F_\Lambda(\Lambda(x))) \quad (31)$$

This matches the worst required outcomes (lowest T) to the most expensive states (lowest Λ).

- **Speculative Portfolio:** Any remaining budget is allocated to the Sniper bet (Dirac delta at $\operatorname{argmax} \Lambda$).

Interpretation: “Insure then speculate.” First, guarantee the risk floor at minimum cost. Then, use the surplus for maximum expected return.

14.4 Regime 3: The Farmer (Log-Optimal / Median CAGR Maximization)

Objective: Maximize the expected geometric growth rate (equivalently, the Median CAGR).

$$\max_g \int_0^\infty \ln(g(x)) P(x) dx \quad \text{s.t.} \quad \int_0^\infty g(x) Q(x) dx = K \quad (32)$$

Optimal Payoff: Proportional to the likelihood ratio:

$$g^*(x) = \frac{1}{\lambda} \cdot \frac{P(x)}{Q(x)} = \frac{1}{\lambda} \Lambda(x) \quad (33)$$

Interpretation: Unlike the Sniper (which concentrates at a single spike), the Farmer spreads capital across the *entire* state space in proportion to the mispricing signal $\Lambda(x)$. The payoff never reaches zero, avoiding the ruin that destroys geometric growth. This strategy has a lower expected return than the Sniper but the highest median wealth over repeated trials.

14.4.1 Risk-Tolerance Constrained Farmer

Adding a risk tolerance floor $T(x)$:

$$g^*(x) = \max\left(T(x), \frac{1}{\lambda} \Lambda(x)\right) \quad (34)$$

This is a **threshold strategy**: below the threshold, you buy insurance (paying the minimum cost to satisfy T); above the threshold, you allocate according to the Kelly-optimal Λ -shaped profile. This resembles a CPPI (Constant Proportion Portfolio Insurance) strategy.

14.4.2 Solving for λ : Newton-Raphson

The Lagrange multiplier λ is found by solving the budget equation:

$$f(\lambda) = \int_0^\infty \max\left(T(x), \frac{1}{\lambda}\Lambda(x)\right) Q(x) dx = K \quad (35)$$

The derivative is:

$$f'(\lambda) = -\frac{1}{\lambda^2} \int_{\Omega_\lambda} P(x) dx \quad (36)$$

where $\Omega_\lambda = \{x \mid \frac{1}{\lambda}\Lambda(x) > T(x)\}$ is the “active set.” Since f is monotonically decreasing and convex, Newton-Raphson converges rapidly (typically 3–4 iterations).

CallCulator’s polynomial spline infrastructure enables exact evaluation of both $f(\lambda)$ and $f'(\lambda)$ via coefficient-shifting integration, avoiding discretization noise entirely.

14.5 Regime 4: The Bunker (Max Min Optimization)

Objective: Maximize the worst-case outcome.

$$\max_g \left(\min_x g(x) \right) \quad \text{s.t.} \quad \text{Cost}[g] = K \quad (37)$$

Optimal Payoff: A constant (flat line):

$$g^*(x) = M \quad (\text{constant}) \quad (38)$$

Proof: Write $g(x) = M + \epsilon(x)$ where $\epsilon(x) \geq 0$ is any upside. The cost is:

$$\text{Cost} = M \int Q(x) dx + \int \epsilon(x)Q(x) dx = K \quad (39)$$

To maximize M , minimize the “Cost of Upside.” Since $\epsilon \geq 0$ and $Q > 0$, the integral is minimized only when $\epsilon(x) = 0$ everywhere.

Implementation: A Box Spread (Long Call K_1 , Short Call K_2 , Long Put K_2 , Short Put K_1), which creates a synthetic risk-free bond with constant payoff $K_2 - K_1$.

14.6 Regime 5: The Sentinel (Soft Minimum / Percentile-Based Targets)

The pure Bunker is extremely conservative—it protects against *every* possible outcome, including those with near-zero probability. The **Sentinel** relaxes this by maximizing a higher percentile (e.g., 5th or 10th) rather than the absolute minimum.

Objective: Maximize the Value-at-Risk at level p^* :

$$\max_g (q_g^*) \quad \text{where} \quad q_g^* = \inf \{z : P(g(x) \leq z) \geq p^*\} \quad (40)$$

Interpretation: The Sentinel guards the important percentiles (e.g., 95% of outcomes) while accepting a “sacrifice zone” (the worst 5%). This frees up budget to improve the protected floor.

Hybrid Approach: Combine a soft-minimum floor with the Farmer’s growth allocation:

$$g^*(x) = \max\left(M, \frac{1}{\lambda}\Lambda(x)\right) \quad (41)$$

where M is the soft-minimum floor (determined by the p^* percentile constraint) and the remaining budget is spent on the $\frac{1}{\lambda}\Lambda(x)$ curve above it. This creates a guaranteed floor for the protected percentiles with probability-weighted growth on top.

14.7 Summary: Comparison of All Regimes

| Regime | Optimization | Objective | Optimal Payoff $g^*(x)$ |
|---------------------|---|-------------------------------|--|
| 0: Scavenger | max residual $ \epsilon_i $ | max $ \epsilon_i $ | Outlier isolation |
| 1: Sniper | \max_g | $\mathbb{E}_P[g]$ | Dirac-delta at $\operatorname{argmax}(\Lambda(x))$ |
| 2: Insurer + Sniper | $\max_{\substack{s \in \mathcal{S} \\ \forall q \in [0,1]: R_s^{-1}(q) \geq T(q)}}$ | $\mathbb{E}_P[g]$ | Floor T + Dirac-delta at $\operatorname{argmax}(\Lambda(x))$ |
| 3: Farmer | \max_g | $\mathbb{E}_P[\ln(g)]$ | $\frac{1}{\lambda} \Lambda(x)$ |
| 4: Bunker | \max_g | $\min_x g(x)$ | Constant (Box Spread) |
| 5: Sentinel | \max_g | $\operatorname{VaR}_{p^*}[g]$ | Sacrifice zone floor M + growth |

Current Implementation: CallCulator’s Probability & Risk mode implements Regime 2 (Insurer + Sniper). The “Use Monte-Carlo Median as objective function” toggle is a placeholder for Regime 3 (Farmer). Max Min mode implements Regime 4 (Bunker), with Regime 5 (Sentinel) planned for future implementation.

Part V

Reference

15 Glossary

- **PDF** – Probability Density Function. Describes the likelihood of each possible stock price.
- **CDF** – Cumulative Distribution Function. $\Psi(x) = \int_{-\infty}^x \mu(y) dy$.
- **IV** – Implied Volatility. The market’s expectation of future volatility, back-solved from option prices via Black-Scholes.
- **Strike Price (K)** – The price at which an option can be exercised.
- **Premium (p)** – The market price of an option contract.
- **Spread** – A combination of two or more option legs.
- **CAGR** – Compound Annual Growth Rate. The annualized return assuming reinvestment.
- **Kelly Criterion** – The optimal fraction of capital to wager for maximum geometric growth.
- **PEE** – Probabilistic Early Expiration. A backend caching strategy that refreshes data before it strictly expires.