CONTENTS

SECTION 01

What Is This Tool?

The **Value at Risk Analyzer** is a parametric, single-asset risk measurement tool that uses the **normal distribution** to answer the most fundamental question in financial risk management: what is the maximum loss I should expect on this position over a given time period, at a given level of confidence?

Every parameter — mean return, volatility, time period, asset value, and confidence level — updates the chart and all six result metrics **live** as you move a slider or type a number. There is no Calculate button. The distribution redraws instantly, letting you explore the entire risk landscape of an asset in seconds.

Who should use this tool?

Risk managers computing position-level VaR, portfolio managers setting stop-loss limits, treasury teams reporting market risk exposure, finance students mastering the parametric VaR methodology, and any professional who needs a rigorous, visual answer to "how much can I lose and how likely is it?"

What the tool does NOT do

This tool models a **single asset** using the **parametric (normal distribution) method** only. It assumes returns are normally distributed, which underestimates tail risk for assets with fat-tailed or skewed return distributions (equities during crises, options, cryptocurrencies). It does not compute historical simulation VaR or Monte Carlo VaR, and it does not model multi-asset portfolio VaR. These limitations are standard to the parametric approach and are discussed further in Section 8.

SECTION 02

Quick Start

The tool computes VaR and Expected Shortfall live — no button click required after the initial load. You can run a complete parametric VaR analysis in under a minute.

1 ENTER AN ASSET NAME (OPTIONAL)

Type a name for the position in the Asset Name field at the top of the parameters panel — for example, `Reliance Industries` or `Nifty 50 ETF`. This name appears as a gold italic watermark on the chart and in the PDF report. It is optional; leaving it blank has no effect on the calculations.

2 SET YOUR PARAMETERS

Enter the asset's annualised mean return and annualised standard deviation (volatility), select the time period from the dropdown, enter the current market value of your position, and set your desired confidence level. Every field has a paired slider for ±50% adjustment. The results update instantly with each change.

3 READ THE RESULTS BOX AND CHART

The dark results box below the parameters panel shows six metrics: Period Mean, Period Volatility, VaR (Absolute), Asset Value at VaR, VaR as % of Asset, and Expected Shortfall. The chart visualises the full return distribution, with the loss tail shaded green, the profit region shaded teal, and the VaR threshold marked by a gold dashed line.

> **No recalculation needed.** Every slider move and every keystroke updates all six result metrics and redraws the distribution chart in real time. Use the sliders to perform quick sensitivity analysis — for example, drag the confidence level from 90% to 99% to see how VaR escalates as you demand greater certainty.

SECTION 03

Understanding the Parameters

Five parameters drive all calculations. Each has a numeric input box for precise entry and a paired slider that adjusts the value ±50% of its current level for rapid sensitivity analysis. The Time Period parameter uses a dropdown instead of a numeric input.

MEAN RETURN (ANNUAL)

The annualised expected return of the asset, expressed as a percentage. This is the centre of the return distribution — the average annual gain (or loss) you expect the asset to deliver.

For equities, a common approach is to use the historical arithmetic average annual return (e.g., from Bloomberg, NSE/BSE data, or Yahoo Finance). For stress-testing purposes, you may wish to set this to zero or a negative value to model a bearish scenario.

`Period μ = Annual μ × t`

Default: 15%. Can be negative. Slider adjusts ±50% of the entered value.

STD DEVIATION (ANNUAL)

The annualised standard deviation of returns — commonly called volatility. This is the single most important input for VaR: it determines how spread out the return distribution is and therefore how large the tail losses can be.

Typical ranges: large-cap Indian equities 18-28%, mid-caps 25-40%, crypto 60-120%, government bonds 4-8%, gold 14-18%. Use the 252-day or 12-month historical standard deviation of daily or monthly returns, annualised.

$$\text{Period } \sigma = \text{Annual } \sigma \times \sqrt{t}$$

Default: 12%. Minimum: 0.01%. Slider adjusts ±50% of the entered value.

TIME PERIOD

The holding period over which VaR is measured. Regulatory and industry convention typically uses one-day VaR (Basel III) or ten-day VaR (market risk capital), but this tool offers four practitioner-friendly horizons:

Annual — t = 1.0 (one full year)
Quarterly — t = 0.25 (three months)
Monthly — t = 0.0833 (one month = 1/12)
Weekly — t = 0.0192 (one week ≈ 1/52)

Both mean and volatility are scaled from annual to the chosen period using the square-root-of-time rule. Longer horizons produce larger absolute VaR because uncertainty compounds with time.

A dropdown selector — no slider. Change to see how risk grows with horizon.

ASSET VALUE

The current market value of your position in rupees (or any currency — the tool uses ₹ as the display currency). This converts the percentage-based VaR into an absolute currency amount you can report to risk committees, regulators, or clients.

Enter the full position size — for example, ₹10,00,000 for a ten-lakh rupee holding. The Asset Value at VaR result will show the floor value of the position at the confidence threshold.

Default: ₹10,00,000. Slider adjusts ±50%. Minimum meaningful value: ₹1.

VAR CONFIDENCE LEVEL

The probability threshold for VaR. A 95% confidence level means: "We are 95% confident that losses over the period will not exceed the VaR figure." Equivalently, there is a 5% probability that losses will exceed VaR — the left tail of the distribution.

Common industry standards:
90% — used in less conservative internal risk models
95% — most widely used in practice and in academic literature
99% — Basel III regulatory capital requirement
99.9% — extreme stress-testing and operational risk

$$z_{VaR} = \Phi^{-1}(1 - \text{confidence})$$

Default: 90%. Range: 0.1% to 99.9%. Slider adjusts the confidence level directly. Note: z_VaR is always negative — it measures how many standard deviations into the left tail the threshold falls.

SECTION 04

The Results Box

Directly below the parameters panel sits the dark-background results box, showing six metrics arranged in a 2×3 grid. All six update live as you change any parameter. The teal-coloured metrics are descriptive statistics; the red-coloured metrics are risk measures. The green/red Asset Value at VaR changes colour based on whether the floor value is above or below the original position value.

PERIOD M

The **mean return scaled to the chosen time period**. This is the centre of the distribution plotted on the chart's x-axis. A monthly period with a 15% annual mean gives a period mean of 15% ÷ 12 = 1.25%.

```
Period μ = Annual μ × t
```

Displayed in teal. This is the expected return of the asset over the holding period, not the risk measure. A positive Period μ shifts the distribution rightward (away from losses).

PERIOD Σ

The **standard deviation scaled to the chosen time period**. This is the spread (width) of the normal distribution on the chart. A monthly period with 12% annual volatility gives a period volatility of 12% ÷ √12 ≈ 3.46%.

```
Period σ = Annual σ × √t
```

Displayed in teal. This is the single largest driver of VaR — doubling Period σ roughly doubles VaR (for a given confidence level and mean).

VAR (ABSOLUTE)

The **maximum expected loss in currency terms** at the chosen confidence level over the chosen period. This is the headline risk metric: the loss amount you should not expect to exceed with the stated probability.

```
Return at VaR = Period μ + z_VaR × Period σ
VaR = -Return at VaR × Asset Value
```

where z_VaR = $\Phi^{-1}(1 - \text{confidence})$ is a negative number (the left-tail z-score). The negative sign converts the negative return into a positive loss amount.

Displayed in red. Always a positive number — it represents a loss. Example: VaR of ₹82,000 means "we do not expect to lose more than ₹82,000 over this period with the stated confidence."

ASSET VALUE AT VAR

The **floor value of the position** at the VaR threshold — the asset value at the worst expected outcome with the stated confidence.

```
Asset Value at VaR = Asset Value × (1 + Return at VaR)
```

Displayed in green if the floor value is above the original position value (unusual — only when Return at VaR is positive, i.e., a very high mean or very low volatility). Displayed in red in the normal case where the floor value is below the original. Useful for reporting the minimum expected portfolio value in risk reports.

VAR AS % OF ASSET

The **VaR expressed as a percentage of the current position value**. This is the relative loss measure — independent of the absolute size of the position and therefore directly comparable

across positions of different sizes.

```
VaR % = (-Return at VaR) × 100%
```

Displayed in red. Example: VaR% of 8.2% means "the asset could lose up to 8.2% of its value over this period at the stated confidence." This is the figure typically used when comparing risk across a portfolio of positions of different sizes.

EXPECTED SHORTFALL

The average loss in the worst (1 − confidence)% of scenarios — also called CVaR (Conditional Value at Risk) or ES. Where VaR only tells you the loss threshold, ES tells you the average loss beyond that threshold. ES is always larger than VaR and is considered a superior risk measure because it is coherent and captures tail severity.

```
ES = -[Period μ - Period σ × φ(z_VaR) / (1 - cl)] × Value
```

where φ(z_VaR) is the standard normal PDF evaluated at the VaR z-score, and cl is the confidence level as a decimal.

Displayed in red. At 95% confidence, ES is typically 15-25% larger than VaR for a normal distribution. The Basel IV framework (FRTB) has replaced VaR with ES at 97.5% confidence as the primary regulatory risk measure.

SECTION 05

Reading the VaR Chart

The chart plots the **normal probability distribution of asset values** at the end of the chosen holding period. The x-axis shows the asset value at each possible return outcome; the y-axis shows the probability density at that value. The entire area under the curve sums to 1 (100% of all possible outcomes).

The Teal Gradient Curve — The Normal Distribution

The smooth bell-shaped curve is the **normal probability density function** of the asset's return over the period, scaled by the asset value to display on the x-axis in currency terms. The peak of the curve is at the Period Mean (Asset Value × (1 + Period μ)). The width of the curve is determined by Period σ — a higher volatility asset has a flatter, wider bell, meaning outcomes are spread over a much larger range of asset values. The entire curve is rendered with a teal-to-dark gradient fill for readability.

The Green Gradient Fill — The Loss Tail

The left portion of the distribution — from the minimum x-axis value up to the VaR threshold — is filled with a green gradient. This region represents the **(1 − confidence)% worst outcomes**: the scenarios where losses exceed the VaR level. For a 95% confidence level, this green tail contains exactly 5% of the total probability mass. The label inside this region displays the exact tail probability (e.g., "10.0% prob below VaR" for a 90% confidence level). The average loss within this tail is the Expected Shortfall.

The Teal Gradient Fill — The Profit Region

The right portion of the distribution — from the VaR threshold to the right side of the chart — is filled with a lighter teal gradient. This region represents the **confidence% of outcomes** where losses are less than VaR or the position ends in profit. The larger the

confidence level you set, the more of the distribution sits to the right of the VaR line, and the more extreme the VaR figure becomes to capture the remaining left tail.

The Gold Dashed Vertical Line — VaR Threshold

The gold dashed vertical line marks the **VaR level** — the asset value at the confidence threshold. Everything to the left of this line is the loss tail; everything to the right is within the confidence range. The label above the line shows "VaR XX%" where XX is the current confidence level. As you raise the confidence level, this line moves further left (deeper into the tail), and the VaR amount increases.

The Muted Dashed Vertical Line — Mean Return

A muted grey dashed line marks the **expected (mean) asset value** at end of period — the peak of the bell curve. This is Asset Value × (1 + Period μ). When the mean return is positive, this line sits to the right of the original asset value, indicating expected growth. The gap between the mean line and the VaR threshold line is the expected-return buffer that reduces VaR: a higher expected return shifts the distribution rightward, moving the left tail further from the loss zone.

Gold italic

The Gold Italic Watermark — Asset Name

If you enter a name in the Asset Name field, it appears as a **gold italic watermark** in the top-right area of the chart. This name also carries through to the PDF export header. It has no effect on calculations — it is a labelling convenience for reports and presentations.

Hover tooltip

Interactive Hover Tooltip

Hovering anywhere over the chart displays a **tooltip** showing four values at the cursor's x-position: the **Asset Value** at that point, the **Return %** relative to the original position, the **Probability Density** (height of the curve at that point), and the **Prob ≤ this** (the cumulative probability — the fraction of outcomes that are at or below this asset value). The Prob ≤ this value is most useful: when the cursor is at the VaR line, it will read exactly (1 − confidence).

Reading the chart at a glance

The visual size of the green tail relative to the total curve area equals (1 − confidence level). At 90% confidence the green tail is visibly wider than at 99%. The VaR gold line moves left as confidence rises. The deeper into the tail the line moves, the larger the absolute VaR number in the results box — this is the fundamental VaR-confidence trade-off made visible.

SECTION 06

Core Concepts Explained

The following concepts are the building blocks of the parametric VaR framework. Understanding them converts the tool's outputs from numbers into a coherent risk management language.

Value at Risk (VaR)

VaR is a statistical measure of the maximum expected loss over a defined period at a given confidence level. "Maximum expected loss" is a slightly misleading phrase — VaR is really a threshold: losses exceed it with probability (1 − confidence). A 95% one-month VaR of ₹80,000

means there is a 5% chance the loss will exceed ₹80,000 in any given month. VaR says nothing about how large the loss is when it does exceed the threshold — that is the role of Expected Shortfall.

The Normal Distribution

The parametric method assumes asset returns follow a normal (Gaussian) distribution, characterised by mean μ and standard deviation σ. This assumption makes VaR computation analytically tractable — a single formula yields the exact answer. The limitation is that real asset returns have fat tails (extreme losses are more common than the normal model predicts) and negative skew (large losses are more likely than large gains of the same magnitude). The normal model underestimates tail risk during market stress.

Expected Shortfall (ES / CVaR)

ES is the average loss conditional on the loss exceeding VaR — the mean of the left tail. It is mathematically defined as: E[Loss | Loss > VaR]. ES is always larger than VaR and provides a more complete picture of tail risk. For a normal distribution, ES at confidence c is: ES = −[μ − σ × φ(z) / (1 − c)] × Asset Value, where φ is the normal PDF. ES is a coherent risk measure (VaR is not) and has been adopted by Basel IV as the primary regulatory capital metric.

The Confidence Level

The confidence level determines how far into the left tail VaR is measured. Raising confidence from 90% to 99% does not mean the risk is "ten times safer" — it means you are capturing a more extreme tail event. The z-score for 90% is -1.28; for 95% it is -1.645; for 99% it is -2.326. The VaR roughly doubles from 90% to 99.9% confidence for a normal distribution. The choice of confidence level should match its intended use: 95% for internal management, 99% for Basel regulatory capital, 99.9% for operational risk and stress testing.

The Square-Root-of-Time Rule

Under the assumption that returns are independent and identically distributed (i.i.d.) over time, volatility scales with the square root of the time horizon: σ_T = σ_annual × √t. Returns (means) scale linearly: μ_T = μ_annual × t. This is the basis for converting annual parameters to any sub-annual holding period. The rule breaks down in practice when returns are autocorrelated (trending or mean-reverting), which is common over short horizons for many assets.

Parametric vs. Historical VaR

Parametric VaR (this tool) fits a distribution to return data and computes VaR analytically. It is fast, elegant, and produces a smooth risk surface — but inherits all assumptions of the chosen distribution. **Historical simulation VaR** uses actual past returns directly — no distribution assumption — and therefore captures fat tails and skewness automatically. Most large institutions use historical simulation or Monte Carlo as their primary method, and parametric VaR as a fast cross-check or for instruments with limited history.

Why VaR is not a "worst-case" loss

A common misconception is that VaR represents the worst possible loss. It does not. A 99% one-month VaR of ₹2,00,000 means losses exceed ₹2,00,000 in roughly 1% of months — that is approximately one month every eight years. In those tail months, the loss could be ₹2,10,000 or ₹10,00,000 — VaR provides no information on magnitude beyond the threshold. This is why regulators and risk managers use **Expected Shortfall** (which this tool computes) to quantify the average severity of tail losses, not just their frequency.

SECTION 07

Business Applications

Parametric VaR and Expected Shortfall are foundational tools across risk management, treasury, trading, and regulatory compliance:

- **Portfolio risk reporting.** Risk managers compute VaR for each position in a portfolio and aggregate them (with correlations) to report the total portfolio VaR to senior management and the board on a daily basis. This tool provides the single-asset building block for position-level risk.
- **Position sizing and stop-loss setting.** Traders use VaR to determine the maximum position size consistent with their risk limit. If the desk's VaR limit is ₹50 lakhs and the position-level VaR is ₹8 lakhs, the trader can hold at most 50 / 8 ≈ 6 positions of this size without breaching the limit. Sliders make this "what-if" sizing analysis instant.
- **Regulatory capital computation.** Under Basel III, market risk capital is computed as a multiple (typically 3×) of the 10-day, 99% VaR. Under Basel IV (FRTB), this has shifted to a 97.5% Expected Shortfall. This tool lets treasury teams estimate the capital requirement for individual positions and understand the sensitivity to the confidence level and holding period.
- **Client risk reporting.** Wealth managers and advisors use VaR to communicate risk to clients in plain language: "Your ₹10 lakh equity position has a one-month 95% VaR of ₹82,000 — meaning we do not expect to lose more than ₹82,000 in any given month with 95% confidence." The PDF export produces a client-ready report with this language embedded.
- **Stress testing and scenario analysis.** By adjusting volatility upward (e.g., to crisis-level values observed during COVID-19 or the 2008 GFC) and mean return downward, users can stress-test their positions to see how VaR escalates under adverse conditions. The live sliders make this exploration intuitive and fast.
- **Risk budgeting.** Fund managers allocate a total risk budget (e.g., portfolio VaR of ₹1 crore) across multiple strategies or positions. Each position's standalone VaR is computed using this tool, then aggregated (allowing for diversification benefits from correlation) to ensure the total stays within budget.
- **Academic and examination submissions.** The tool produces publication-quality charts with all parametric VaR formulas embedded in the results, suitable for FRM/CFA exam preparation, MBA assignments on market risk management, and academic case studies on the VaR methodology.
- **Comparing VaR across confidence levels.** The drop from VaR to ES across multiple confidence levels reveals the shape of the tail and helps risk managers decide which metric best reflects the tail risk of their specific asset. Equities at 95% typically show ES/VaR ratios of 1.2–1.4; ratios above 1.6 indicate significant tail risk beyond the normal assumption.

SECTION 08

Interpreting Your Results

How to read the VaR figure

The VaR (Absolute) result is a **currency loss amount, not a forecast.** A monthly 95% VaR of ₹75,000 on a ₹10 lakh position means: in the worst 5% of months (approximately one month every twenty months, or roughly once every 1.7 years), losses are expected to exceed ₹75,000. In the remaining 95% of months, losses are less than ₹75,000 — and in most of those months, the position gains value.

When to use 90% vs. 95% vs. 99% confidence

The choice of confidence level should be driven by the **purpose** of the VaR calculation:

- **90%** — internal desk-level monitoring; early-warning signals; less severe capital implications.
- **95%** — standard for most internal risk reports, academic literature, and client communications.
- **99%** — Basel III regulatory capital calculation; conservative institutional risk limits.
- **99.9%** — operational risk capital (Basel II Advanced Measurement Approach); extreme stress scenarios.

Using 99% when 95% is more appropriate will produce a VaR that looks much larger and may unnecessarily constrain trading positions. Using 90% for regulatory purposes will understate the required capital. Match the confidence level to the context.

VaR vs. Expected Shortfall — which to use?

Always look at both. VaR tells you the probability of a loss exceeding the threshold; ES tells you the average severity of losses when they do exceed it. If VaR is ₹80,000 and ES is ₹82,000, the loss tail is thin — losses rarely extend far beyond VaR. If ES is ₹1,40,000 against a VaR of ₹80,000, the tail is fat — when losses exceed VaR, they tend to be dramatically larger. The ES/VaR ratio is a practical measure of tail severity. For a standard normal distribution, ES/VaR at 95% confidence is approximately 1.25.

The role of mean return in VaR

The mean return acts as a buffer. An asset with a 15% annual mean return and 20% annual volatility has a meaningfully lower monthly VaR than one with a -5% mean and 20% volatility — even though the volatility is identical. This is because the positive expected return shifts the entire distribution rightward, moving the left tail further from the loss zone. For short holding periods (weekly, monthly), the mean effect is small relative to volatility because mean scales with t while volatility scales with √t. For annual horizons, mean return becomes significant.

Limitations of the parametric model

The normal distribution assumption is the central limitation of this tool. Real equity returns exhibit: (1) **Fat tails** — extreme losses occur more frequently than the normal distribution predicts. The true 99% loss is often 30-50% larger than the parametric estimate. (2) **Negative skew** — large down moves are more common than large up moves of the same magnitude. (3) **Volatility clustering** — high-volatility periods cluster together (GARCH effects), meaning a single static σ underestimates risk during stress. For a more accurate tail estimate, supplement this tool with historical simulation or a fat-tailed distribution model (Student-t, GEV) when making regulatory capital or large position decisions.

SECTION 09

Saving Models & Exporting PDFs

Two productivity features — **Save Model** and **Export PDF** — are available to Trial and Premium users. Free users can run the full parametric VaR analysis and view all six results on screen; saving and PDF export require an account.

FEATURE	FREE	TRIAL	PREMIUM
VaR Calculation	✓ Unlimited	✓ Unlimited	✓ Unlimited
Live Chart & Results	✓ Unlimited	✓ Unlimited	✓ Unlimited
Export PDF	✗ Not available	✓ Unlimited	✓ Unlimited
Save Model	✗ Not available	Up to 3 models	✓ Unlimited
Load Saved Model	✗ Not available	✓ All saved models	✓ All saved models

Exporting a PDF Report

Click the ⬇ **Export PDF** button after reviewing the on-screen results. The tool generates a formatted A4 landscape PDF entirely in your browser using jsPDF — no data is sent to a server — and downloads it immediately.

PARAMETERS BAR
A dark header bar at the top of the PDF lists every input used: Asset Name, Annual Mean Return, Annual Standard Deviation, Time Period, Asset Value, and VaR Confidence Level. This makes the PDF fully self-contained — a reader can reproduce the analysis without access to the tool.

CHART IMAGE
The full distribution chart — normal curve, green loss tail, teal profit region, gold VaR threshold line, mean return line, asset name watermark — is embedded in the PDF at high resolution. The chart occupies the left column of the A4 landscape layout.

METRICS COLUMN
The right column of the PDF shows all six result metrics: Period μ, Period σ, VaR (Absolute), Asset Value at VaR, VaR as % of Asset, and Expected Shortfall. Each metric is presented with its label and computed value, formatted for inclusion in risk reports or course submissions.

> Enter the Asset Name before exporting — it appears as a prominent label in the PDF chart watermark and in the parameters bar, turning a generic VaR report into an asset-specific document suitable for a risk management committee presentation or MBA assignment submission.

Saving and Loading Models

A **model** is a named snapshot of all your current inputs: Asset Name, Annual Mean Return, Annual Standard Deviation, Time Period, Asset Value, and Confidence Level. Saving a model lets you return to a previous risk analysis in any future session with a single click.

1 CLICK "SAVE MODEL"
The Save Model button appears once you are logged in with a Trial or Premium account. Click it after setting up the VaR analysis you want to preserve — for example, after entering a specific asset's risk parameters for a recurring monthly risk report.

2 ENTER A MODEL NAME (UP TO 10 CHARACTERS)

Choose a short, descriptive name — for example REL-MONTHLY, NIFTY-95, or GOLD-ANN. The name identifies the model in the dropdown. Saving under an existing name prompts you to confirm overwrite.

3 LOAD IT ANY TIME FROM THE DROPDOWN

Select any saved model from the Load saved model dropdown. All inputs are restored instantly — Asset Name, mean, volatility, time period, asset value, and confidence level — and all six results update immediately. No recalculation step required.

Trial — up to 3 saved models

During a Trial, you can save up to three distinct risk scenarios. Overwriting an existing model does not consume an additional slot. When the three-model limit is reached, overwrite an existing model or upgrade to Premium. A practical approach: save the three most commonly used positions (e.g., your largest equity holding, your bond position, and a benchmark index).

Premium — unlimited saves

Premium users can save as many models as needed — one per position in a portfolio, one per asset class, or a series of stress scenarios for the same position with different volatility assumptions. Premium is ideal for risk managers maintaining a library of position-level VaR scenarios.

What exactly is stored in a saved model?

Asset Name · Annual Mean Return (%) · Annual Standard Deviation (%) · Time Period (Annual / Quarterly / Monthly / Weekly) · Asset Value (₹) · VaR Confidence Level (%). All values are stored as entered at the time of saving. When you load a model, all six result metrics and the distribution chart update instantly to reflect the restored parameters.

SECTION 10

Glossary of Terms

A quick-reference table of every technical term used in the tool and this guide.

TERM	DEFINITION	IN THIS TOOL
Value at Risk (VaR)	The maximum expected loss over a defined holding period at a specified confidence level. Not a worst-case loss — losses exceed VaR with probability (1 − confidence). Introduced by J.P. Morgan's RiskMetrics in 1994.	Primary output. Shown as VaR (Absolute) in the results box (red).
Expected Shortfall (ES)	The average loss conditional on the loss exceeding VaR — the mean of the left tail. Also called CVaR or Conditional VaR. Always	Sixth result metric. Shown in red. Computed as −[μ_T − σ_T ×

	larger than VaR. A coherent risk measure. Adopted by Basel IV (FRTB) at 97.5% confidence.	φ(z) / (1−cl)] × Value.
CVaR	Conditional Value at Risk — another name for Expected Shortfall. The conditional expectation of loss given that the loss exceeds VaR.	Same as Expected Shortfall in the results box.
Normal Distribution	A symmetric, bell-shaped probability distribution characterised by mean μ and standard deviation σ. The parametric VaR method assumes asset returns follow this distribution, enabling analytical computation of VaR.	The teal bell-shaped curve on the chart. Computed using the standard normal PDF and CDF.
Confidence Level (cl)	The probability that the actual loss will not exceed the VaR figure. Common values: 90%, 95%, 99%, 99.9%. The complement (1 − cl) is the probability the loss exceeds VaR — i.e., the size of the left tail.	VaR Confidence Level parameter. Determines the z-score and the size of the green loss tail.
Parametric VaR	A VaR method that assumes a parametric distribution (usually normal) for returns and computes VaR analytically using the distribution's parameters. Fast and elegant, but sensitive to the distributional assumption.	The method used throughout this tool.
σ (Sigma)	Standard deviation of returns — the primary measure of volatility and risk in financial models. Annualised σ is the input; period σ is derived using the square-root-of-time rule.	Std Deviation (Annual) input. Period σ shown in results box (teal).
μ (Mu)	Mean (expected) return of the asset. Annualised μ is the input; period μ is derived by multiplying by t. A positive μ reduces VaR by shifting the return distribution rightward.	Mean Return (Annual) input. Period μ shown in results box (teal).
z-score (z_VaR)	The number of standard deviations from the mean corresponding to a given confidence level. For VaR, $z\_VaR = \Phi^{-1}(1 - cl)$. Always negative for loss-side VaR. Examples: z at 90% = −1.282, at 95% = −1.645, at 99% = −2.326.	Computed internally via the inverse normal CDF (rational approximation). Used in the VaR and ES formulas.
Φ^{-1} (Inverse CDF)	The inverse of the standard normal cumulative distribution function. Given a probability p, $\Phi^{-1}(p)$ returns the z-score such that $P(Z \leq z) = p$. Used to convert a confidence level into a VaR z-score.	Computed using a rational polynomial approximation (Beasley-Springer-Moro algorithm). Used for z_VaR.
φ (Normal PDF)	The standard normal probability density	Computed in the ES formula:

	function. φ(z) = (1/√2π) × exp(−z²/2). Appears in the Expected Shortfall formula as the height of the normal curve at the VaR z-score.	φ(z_VaR) / (1 − cl) scales the ES calculation.
Square-Root-of-Time Rule	The scaling rule for converting volatility across time horizons under i.i.d. return assumptions: σ_T = σ_annual × √t; μ_T = μ_annual × t. Volatility grows with √t because variance is additive but standard deviation is not.	Applied when computing Period σ and Period μ from annual inputs for any chosen time period.
Holding Period (t)	The time horizon over which VaR is measured, expressed as a fraction of a year. Annual = 1.0, Quarterly = 0.25, Monthly = 0.0833, Weekly = 0.0192. Longer horizons produce larger absolute VaR because uncertainty compounds.	Time Period dropdown. Determines t used in all scaling formulas.
Probability Density	The height of the probability distribution at a given outcome. Does not directly represent probability — the probability of an outcome in a range is the area under the density curve over that range.	Y-axis of the distribution chart. Shown in the hover tooltip as "Density".
Cumulative Distribution (CDF)	The probability that the random variable takes a value less than or equal to a given point. For a normal distribution, the CDF at the VaR point equals (1 − confidence level). Φ(x) denotes the standard normal CDF.	Shown in the hover tooltip as "Prob ≤ this". At the VaR threshold, reads (1 − confidence).
Fat Tails	A property of return distributions where extreme outcomes occur more frequently than predicted by the normal distribution. Measured by excess kurtosis (kurtosis > 3). Equity returns typically exhibit fat tails, meaning the normal model underestimates the probability and magnitude of large losses.	A key limitation of this tool's parametric model. Discussed in Section 8.
Basel III / IV	International banking regulatory frameworks (Bank for International Settlements). Basel III required 10-day, 99% VaR for market risk capital. Basel IV (Fundamental Review of the Trading Book, FRTB) replaced this with 97.5% Expected Shortfall.	Context for the 99% confidence level and the regulatory importance of ES alongside VaR.

← BACK TO VALUE AT RISK