

NPV Monte Carlo *Simulator*

CONTENTS

- | | |
|------------------------------|---------------------------------------|
| 1. What Is This Tool? | 2. Quick Start |
| 3. Deal Parameters | 4. Simulation Drivers & Distributions |
| 5. Reading the Results Bar | 6. Reading the Charts |
| 7. Core Concepts | 8. Business Applications |
| 9. Interpreting Your Results | 10. Saving Scenarios & Exporting PDFs |
| 11. Glossary of Terms | |

SECTION 01

What Is This Tool?

The **NPV Monte Carlo Simulator** answers a question that a single-point NPV calculation cannot: *not "is this project's NPV positive?" but "how likely is it to be positive, and how much could it swing either way?"* Instead of plugging one guess for sales growth, margins, and the discount rate into a discounted cash flow model, the tool draws thousands of random combinations of your key assumptions – each one internally consistent with the ranges and distribution shapes you specify – and recomputes the project's NPV for every single draw. The result is a full probability distribution of outcomes, not a single number.

Under the hood it runs the same free-cash-flow-to-firm engine used in a standard capital budgeting DCF: revenue is grown forward from a base year, cost and margin assumptions are applied, capex and working-capital investment are deducted, a terminal value captures cash flows beyond the explicit forecast horizon, and everything is discounted back at the WACC. What makes it a *Monte Carlo* simulator is that eight of those inputs – sales growth, COGS/Sales, SG&A/Sales, depreciation, capex/sales, terminal growth, WACC, and the working-capital drag – are treated as random variables rather than fixed numbers, each with its own uncertainty distribution.

Who should use this tool?

MBA students and analysts learning risk-adjusted capital budgeting, corporate finance teams evaluating a new plant, product line, or capex project under genuine uncertainty, private equity and VC analysts stress-testing a deal thesis, and anyone who has ever been asked "how confident are we in that NPV number?" and wanted a rigorous, quantitative answer.

Why simulate instead of just running one scenario?

A single deterministic NPV calculation hides the fact that every input is an estimate. Two projects can have the identical "base case" NPV of ₹50 lakh, yet one might have a 95% chance of a positive outcome while the other has only a 55% chance – because the second project's assumptions carry far more uncertainty. Monte Carlo simulation makes that difference visible. It also identifies, through sensitivity analysis, exactly *which* assumption is doing the most damage to (or the most for) the project's value – information a single-point estimate simply cannot provide.

SECTION 02

Quick Start

You can run your first risk-adjusted NPV simulation in under two minutes.

1 SET YOUR DEAL PARAMETERS

Enter Base Year Sales, Initial Capital Outlay (ICO), Tax Rate, and the NPV Threshold you're testing against (often ₹0, i.e. "what's the probability this project is value-accretive at all?"). Choose your Forecast Years (3-5) and the Number of Simulations (100-5,000 – more simulations give a smoother, more reliable distribution).

2 CONFIGURE THE EIGHT SIMULATION DRIVERS

For each driver (Sales Growth, COGS/Sales, SG&A/Sales, Depreciation, Capex/Sales, Terminal Growth, WACC, and $\Delta(\text{CA-CL})/\Delta\text{Sales}$), choose a distribution type – Triangular, Uniform, Normal, or Constant – and set its Base Case plus the distribution-specific parameters (Min/Max/Mode, or Mean/Std Dev).

3 READ THE LIVE RESULTS

The tool re-runs the full simulation automatically every time you change an input. Watch the stats box, the NPV Distribution histogram, and the Tornado sensitivity chart update in real time. The Prob NPV > Threshold cell is usually the single most important number in the whole model.

4 SAVE AND EXPORT

Name your scenario and save it (Trial and Premium) to return to it later, or export a full PDF report (Premium) for a board deck, investment committee memo, or coursework submission.

SECTION 03

Deal Parameters

These four inputs anchor the simulation. They are entered as single point values (with a base input and a $\pm 50\%$ sensitivity slider), not as random distributions – they define the scale and structure of the deal rather than its uncertainty.

BASE YEAR SALES

Year-0 (starting) revenue. All forecast-year revenue is grown forward from this figure using the simulated Sales Growth rate, compounding year over year.

Default: ₹1,00,00,000. Slider range: $\pm 50\%$ of base value.

INITIAL CAPITAL OUTLAY (ICO)

The upfront investment required at time zero – plant, equipment, working capital seed, or acquisition cost.

ICO is subtracted from the cumulative discounted cash flows in the terminal year to arrive at NPV, and it also anchors the Depreciation driver (depreciation is expressed as a percentage of ICO).

Default: ₹50,00,000. Slider range: ±50% of base value.

TAX RATE

Corporate income tax rate applied to operating profit (Sales - COGS - SG&A - Depreciation) in every forecast year, before depreciation is added back to arrive at free cash flow.

Default: 30%. Slider range: ±50% of base value.

NPV THRESHOLD

The bar the project must clear. Most commonly set to ₹0 (breakeven - "does this project create any value?"), but can be set to any hurdle, e.g. the NPV of the next-best alternative use of capital. The

Prob NPV > Threshold statistic answers: across all simulated draws, what fraction cleared this bar?

Default: ₹0. Slider range: ±50% of base value.



Forecast Years (3-5) and Number of Simulations (100-5,000) are set via their own sliders in the Parameters panel. More simulations reduce sampling noise in the output distribution - 1,000 is a reasonable default; use 5,000 for a final report where precision matters.

SECTION 04

Simulation Drivers & Distributions

Eight assumptions are drawn randomly on every simulated trial - these are the drivers of uncertainty in the model. Each one can be assigned any of four distribution shapes, independently of the others.

SALES GROWTH

Annual revenue growth rate applied compoundingly to Base Year Sales in each forecast year. Default base: 10% (range 5-20%).

COGS / SALES

Cost of goods sold as a percentage of that year's revenue. Drives gross margin. Default base: 40% (range 30-50%).

SG&A / SALES

Selling, general & administrative expense as a percentage of revenue. Default base: 15% (range 10-20%).

DEPRECIATION (% OF ICO)

Annual depreciation expressed as a percentage of the Initial Capital Outlay – a straight-line-style convention that keeps the depreciation schedule tied to the size of the original investment. Default base: 10% (range 5-15%).

CAPEX / SALES

Incremental capital expenditure as a percentage of the change in sales each year – growth-driven reinvestment, separate from the fixed ICO. Default base: 8% (range 5-12%).

TERMINAL GROWTH

The perpetual growth rate applied to the final forecast year's free cash flow when computing the terminal value (Gordon Growth). Must stay meaningfully below WACC. Default base: 3% (range 1-5%).

WACC

Weighted Average Cost of Capital – the discount rate applied to every year's free cash flow and to the terminal value. Default base: 12% (range 9-15%).

$\Delta(\text{CA-CL}) / \Delta\text{SALES}$

Incremental net working capital investment as a percentage of the change in sales – captures the cash tied up in receivables and inventory net of payables as the business grows. Default base: 10% (range 5-15%).

Choosing a distribution type

Triangular (T)

Defined by Min, Mode (most likely), and Max. The most common choice for business forecasting when you have a "most likely" estimate plus a plausible best and worst case but no hard statistical data – which describes most capital budgeting assumptions.

Uniform (U)

Defined by Min and Max only, with every value in between equally likely. Use this when you genuinely have no reason to believe any value in a range is more likely than another – a deliberately conservative, low-information assumption.

Normal (N)

Defined by Mean and Standard Deviation, bell-shaped and symmetric. Appropriate when a driver is the sum of many small independent effects (e.g. a well-diversified sales growth forecast) or when you have historical data suggesting a symmetric spread around an average.

Constant (C)

No randomness – always returns the Base Case value. Use this to "turn off" uncertainty on a driver you're highly confident in (e.g. a contractually fixed tax rate), isolating the risk analysis to the drivers that actually matter.

Validation hints

The tool flags inconsistent inputs in real time – e.g. a Triangular distribution where Mode falls outside the Min/Max range, or a Base Case value outside its own Min/Max. These are non-blocking warnings (the simulation still runs), but they signal the input needs a second look.

Twelve statistics summarize the full distribution of simulated NPVs, recomputed instantly on every input change.

STATISTIC	WHAT IT TELLS YOU
Base NPV	The deterministic NPV using only the Base Case value of every driver – the traditional single-point DCF answer, shown alongside the simulated distribution for comparison.
Mean NPV	The average NPV across all simulated trials – the expected value of the project under uncertainty.
Median NPV	The middle simulated value. When Mean and Median diverge noticeably, the distribution is skewed – check the Skewness statistic.
Std Dev	The spread of simulated outcomes around the mean – the primary measure of the project's total risk.
Variance	Std Dev squared – used less for direct reading, more as an input to further risk calculations.
Skewness	Distribution asymmetry. Positive skew means a long right tail (a small chance of a very large upside); negative skew means a long left tail (tail risk to the downside).
Kurtosis	Tail weight relative to a Normal distribution. High kurtosis means extreme outcomes (very good or very bad) are more likely than a bell curve would suggest.
1st / 3rd Quartile	The 25th and 75th percentile simulated NPVs – the middle 50% of outcomes falls between these two values.
Min / Max	The worst and best single simulated draw – useful for understanding tail scenarios, though any individual draw is just one random outcome, not a forecast.
Prob NPV > Threshold	The single most decision-relevant number: the percentage of simulated trials in which NPV exceeded your chosen threshold. This is computed analytically from the simulated mean and standard deviation via the Normal CDF, not by counting draws.

Reading the Charts



NPV Distribution — Histogram

Teal bars show how many of the simulated trials fell into each NPV bucket – the shape of the project's risk. The gold line and dots trace the cumulative percentage of trials at or below each bucket, letting you read off any percentile directly. The dashed grey vertical line marks the deterministic Base NPV, so you can see at a glance whether the "safe" single-point estimate sits in a thick or thin part of the distribution. Hover anywhere on the chart for exact bucket frequency and cumulative %.



Sensitivity — Tornado Chart

Each horizontal bar isolates one driver at a time: it re-runs the NPV calculation twice, holding every other driver at its Base Case but setting the one driver in question first to the lowest value it drew across the whole simulation, then to the highest. The wider the bar, the more that single driver moves the project's NPV – the bars are sorted so the most influential driver appears at the top. This is the fastest way to identify which assumption deserves the most diligence, negotiation effort, or hedging attention.

Why the tornado bars aren't symmetric

A driver's low-draw and high-draw NPV impacts are not mirror images of each other, because the DCF engine is non-linear (compounding growth, terminal value with a WACC-minus-growth denominator, tax effects). A bar that extends much further in one direction than the other is telling you something real about the shape of that risk – not a charting artifact.

Core Concepts Explained

Monte Carlo Simulation

A technique that estimates the distribution of an uncertain outcome by generating a large number of random scenarios – each drawing every uncertain input from its own probability distribution – and computing the outcome for each one. Named after the Monte Carlo casino for its reliance on repeated random sampling.

Net Present Value (NPV)

The sum of all future free cash flows from a project, each discounted back to today's value at the WACC, minus the initial investment. A positive NPV means the project is expected to create value above the cost of the capital used to fund it.

$$NPV = \sum FCF_t / (1+WACC)^t - ICO$$

Terminal Value

The present value of all cash flows beyond the explicit forecast horizon, computed with the Gordon Growth (perpetuity growth) formula and discounted back to today alongside the explicit-period cash flows. Terminal value is highly sensitive to the spread between WACC and Terminal Growth – as that spread narrows, terminal value grows explosively.

$$TV = FCF_{final} \times (1+g) / (WACC - g)$$

Sensitivity vs. Uncertainty

Sensitivity (the Tornado chart) tells you how much NPV *moves* when a driver changes. Uncertainty (the driver's chosen distribution) tells you how *likely* that driver is to actually move. A high-sensitivity, low-uncertainty driver may matter less in practice than a moderate-sensitivity, high-uncertainty one – the histogram captures the combination of both effects; the tornado isolates sensitivity alone.

Why Probability Beats a Single Number

"The NPV is ₹50 lakh" implies false precision. "There is a 78% chance the NPV exceeds zero, with a 25th percentile outcome of ₹5 lakh and a 75th percentile of ₹95 lakh" gives decision-makers the actual risk-return tradeoff – which is what capital allocation decisions really require.

The Central Limit Theorem, in practice

Even when individual drivers follow Triangular or Uniform distributions, the resulting NPV distribution across thousands of trials often looks approximately Normal – because NPV is a sum-like combination of many random effects. This is why Prob NPV > Threshold can be computed analytically from the simulated mean and standard deviation via the Normal CDF rather

than requiring a manual tally of draws.

SECTION 08

Business Applications

- **Capital expenditure approval.** Present a probability of success alongside the base-case NPV in a capex request – far more persuasive to a capital committee than a single number that hides the underlying uncertainty.
- **New product / plant launch risk analysis.** Quantify how much of the project's risk comes from demand uncertainty (Sales Growth) versus cost uncertainty (COGS, SG&A) versus the discount rate itself (WACC) – and prioritize risk mitigation accordingly.
- **Private equity and venture diligence.** Stress-test a target's projections by widening the Min/Max ranges on aggressive growth or margin assumptions and observing how much the probability of a positive NPV degrades.
- **Negotiating deal terms.** The Tornado chart tells you exactly which contractual terms (e.g. a capex cap, a minimum volume guarantee that stabilizes Sales Growth) would most reduce the deal's downside risk.
- **Comparing mutually exclusive projects.** Two projects with similar Base NPV can have very different risk profiles; save a scenario for each and compare their Prob NPV > Threshold and Std Dev side by side.
- **MBA coursework** Bring a textbook DCF case to life with a full Monte Carlo risk analysis, export a polished PDF, and demonstrate probabilistic thinking beyond a static spreadsheet model.

and case
studies.

SECTION 09

Interpreting Your Results

A high-confidence, low-risk project

Prob NPV > Threshold above ~80%, a narrow gap between the 1st and 3rd quartiles relative to the mean, and low kurtosis. Base NPV and Mean NPV are close together (low skew). This project's value creation is robust to reasonable variation in your assumptions.

A coin-flip project

Prob NPV > Threshold near 50%, a wide quartile spread, and a Base NPV that sits close to the threshold itself. Check the Tornado chart – usually one or two drivers dominate the spread. Tightening the range on those specific drivers (through better data, contracts, or hedges) is the highest-leverage next step before approving the project.

A project with fat-tail risk

High Kurtosis and/or strongly negative Skewness even with a decent Mean NPV. This signals that while the "typical" outcome looks fine, there's a meaningfully elevated chance of a severe downside outcome relative to what a Normal distribution would suggest – worth a closer look at whichever driver's distribution is generating that tail.

Saving Scenarios & Exporting PDFs

Access to saving and exporting depends on your account tier.

FEATURE	FREE	TRIAL	PREMIUM
Run Simulation	✓ Unlimited	✓ Unlimited	✓ Unlimited
Save Scenario	✗ Not available	Up to 3 scenarios	✓ Unlimited
Load / Delete Saved Scenario	✗ Not available	✓ All saved scenarios	✓ All saved scenarios
Export PDF	✗ Not available	✗ Not available (Premium only)	✓ Unlimited

Exporting a PDF Report

Premium users can click **Export PDF** to generate a formatted A4 landscape report entirely in-browser (via jsPDF – no data leaves your device). It includes a header band with the project title and metadata, the full results stats bar (Base/Mean/Median NPV, Std Dev, Prob NPV > Threshold, quartiles, skewness, and kurtosis), snapshots of both the NPV Distribution histogram and the Tornado sensitivity chart, and a summary table of every simulation driver's distribution assumptions – a fully self-contained record of the analysis.

Saving and Loading Scenarios

A saved scenario captures every input: the four deal parameters, forecast years, number of simulations, and all eight drivers' distribution type and parameters. Trial accounts can keep up to three scenarios at a time – delete one to free a slot, or upgrade to Premium for unlimited scenarios (useful for comparing many project variants side by side).

Glossary of Terms

TERM	DEFINITION
NPV	Net Present Value – the sum of discounted future free cash flows minus the initial investment.
WACC	Weighted Average Cost of Capital – the blended required return used to discount project cash flows.
Free Cash Flow (FCF)	Cash generated by the project after operating costs, taxes, and reinvestment (capex and working capital), available to all capital providers.
Terminal Value	The present value of all cash flows beyond the explicit forecast horizon, via the Gordon Growth perpetuity formula.
Monte Carlo Simulation	Estimating an outcome's distribution by running many trials, each drawing random values for the uncertain inputs.
Triangular Distribution	A distribution defined by a minimum, maximum, and most-likely (mode) value – the standard choice when only expert judgment, not statistical data, is available.
Uniform Distribution	Every value between a minimum and maximum is equally likely.
Normal Distribution	A symmetric, bell-shaped distribution defined by a mean and standard deviation.
Standard Deviation	A measure of dispersion around the mean – the primary summary statistic for a distribution's risk.
Skewness	A measure of a distribution's asymmetry; positive skew means a longer right (upside) tail.
Kurtosis	A measure of how much probability sits in the tails relative to a Normal distribution – high kurtosis means more extreme outcomes than a bell curve implies.

TERM	DEFINITION
Quartile	Values that divide the distribution into four equal parts; the 1st (25th percentile) and 3rd (75th percentile) quartiles bound the middle half of outcomes.
Tornado Chart	A one-at-a-time sensitivity chart that isolates each input's impact on the output, sorted from most to least influential.
ICO	Initial Capital Outlay – the upfront cash investment required at the start of the project.