FinLaxmi

Personal finance planner · built for India

The FinLaxmi Methodology

How your plan is computed — and how we keep it honest.

250
simulated market futures

22
glidepath strategies compared

81
years modelled, year by year

1-in-20
the safety bar every plan must clear

Version 1.0 · July 2026 · Finominal AI Pvt Ltd

WHY WE PUBLISH THIS

1 · Advice you can audit

Most financial advice in India is either a product pitch wearing a plan, or a thumb rule (“save 30%”, “the 4% rule”) wearing authority. FinLaxmi is neither. Your plan is computed by a deterministic simulation engine — roughly 450,000 interlocking formulas — and this paper explains, openly, what that engine does with your answers. We publish it because trust should not require faith: the model, the assumptions, the safety standard and the honesty mechanisms are all described here. A small number of implementation details are our intellectual property and are deliberately not published; section 8 tells you exactly which ones, so nothing is quietly hidden.

Five commitments the engine is built around:

- **Every number is computed, never guessed.** The engine produces the figures; the website words them; the AI advisor may only explain figures that already exist in your plan.
- **The safety bar is the worst 1-in-20 market.** A plan is only called safe if it survives the worst 5% of the simulated futures — not the average one.
- **Everything is quoted in today’s rupees.** A crore forty years from now is not a crore. Deflating to today’s purchasing power is a unit conversion, not an opinion.
- **We sell no products and take no commission.** Nothing in the engine knows what any fund house pays anyone, because nobody pays us.
- **Every assumption is yours to change.** Each default in Appendix A can be edited inside your plan, and the whole simulation re-runs.

THE ARCHITECTURE

2 · Three layers, one direction of truth

FinLaxmi separates computing from communicating. The engine owns every number. The website owns every sentence — one voice, one vocabulary, composed from the engine’s numbers. The AI advisor sits on top for explanation and conversation, and is structurally prevented from inventing figures: every number in an AI answer is mechanically checked against your plan’s own data, and an answer that fails the check twice is discarded rather than shown.

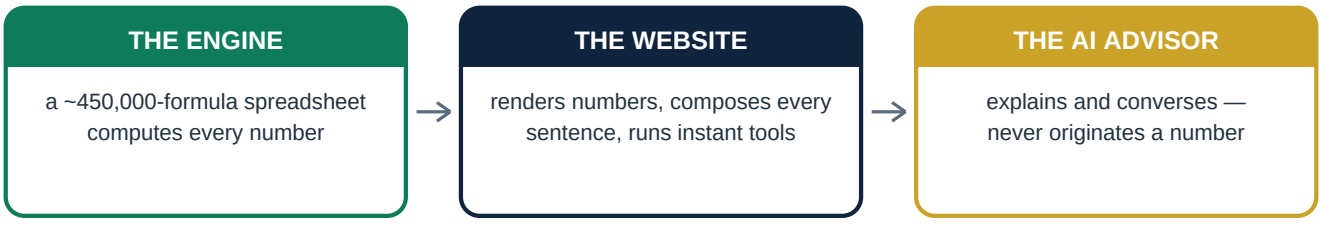

Numbers flow one way. Every figure the AI speaks is mechanically audited against the plan payload before it reaches you.

> **Why this matters**
>
> When a tool blends computation and narration, errors hide in the prose. Separating them means a wrong number can only come from one place — and that place is a spreadsheet that can be tested, audited and replayed, line by line.

YOUR LIFE, AS THE ENGINE SEES IT

3 · A year-by-year ledger, not a formula

The engine does not compress your life into one compounding equation. It builds an explicit ledger, year by year from your current age to your life expectancy (up to ~80 years ahead), and pushes every rupee through it:

- **Income** grows at your stated rate; your annual bonus rides along and both stop at retirement. EPF contributions are credited with the employer match (double your employee deduction).
- **Rental income** from property flows for life and rises with inflation.
- **Spending** starts from what you actually spend today and inflates every year. Yearly recurring costs (insurance, travel, festivals) are carried separately.
- **Children** follow a four-stage cost model — infant, school, college, independence — each stage adding a percentage bump to base spending, per child, ending at the independence age. The engine works with the change relative to today, since today's children are already inside today's spending.
- **Loans** pay their actual EMIs until each tenure ends — they are cash-flows, not an afterthought.
- **Goals** (education, a home, a wedding) are withdrawn in the year they fall due, at their inflation-adjusted cost — so a goal seven years out competes with retirement the way it really does.

This ledger is the engine's spine. The Monte Carlo below runs 250 market futures over exactly this ledger — which is why a loan, a second child or a rental flat changes every answer, not just one card.

THE MARKET MODEL

4 · 250 futures on common random numbers

Markets do not return their average every year, and plans that assume so fail quietly in bad decades. FinLaxmi models annual mutual-fund and FD returns as lognormal random variables — returns can never fall below –100%, large up-years are possible, and the arithmetic mean is preserved exactly:

$$\sigma^2_{log} = \ln\left(1 + \left(\sigma/(1+\mu)\right)^2\right) \qquad \mu_{log} = \ln(1+\mu) - \sigma^2_{log}/2$$

so that a 13% expected return simulates to a 13% mean — the workbook carries verification cells proving it. Gold is modelled as a per-lifetime regime: each simulated future draws one long-run gold return around your gold assumption, reflecting that gold's risk in a plan is decade-scale, not day-scale.

All 250 futures are generated once, from a frozen seed, and **every strategy is tested on the same 250 futures** — the Common Random Numbers method. Strategies are compared on identical market luck, so the ranking measures the strategy, not the dice. It also makes every plan exactly reproducible: same inputs, same answer, always.

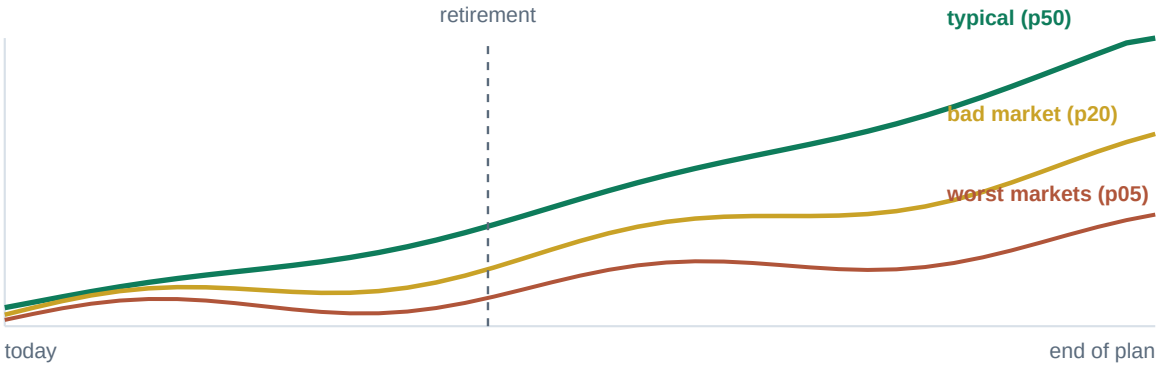

250 lifetimes, three honest summaries — always shown in today's rupees.

We summarise the 250 outcomes at three honesty levels: the **typical** path (median, p50), a **bad market** (the 20th percentile — 1 in 5 lifetimes end below it), and the **worst markets** (the 5th percentile — the 1-in-20 disaster). The worst-markets line is the one your safety verdict hangs on.

THE STRATEGY ENGINE

5 · 22 glidepaths, one survivor

Your asset mix should not be a static pie chart — it should be a path through time. The engine tests 22 strategies: eleven starting mutual-fund weights (0% to 100%) across two proprietary glidepath families:

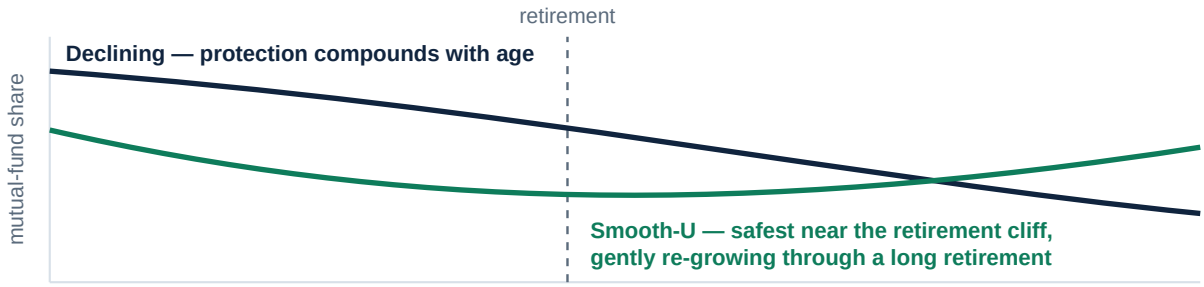

Two families, eleven starting mixes each. Curve anchors and floors are proprietary (§8).

Every strategy runs over your full ledger under all 250 futures. Selection is a two-step discipline: first a **failure gate** — a strategy is only eligible if the share of futures where the money runs out is within your risk style's tolerance (Conservative 1%, Moderate 2%, Aggressive 5%). Then, among the survivors, the engine ranks outcomes at your style's own percentile — a conservative investor is judged by bad-market outcomes, an aggressive one by typical outcomes — and recommends the winner. If nothing passes the gate, the engine says so and recommends the least-bad strategy honestly. The exact scoring and tie-break rules are proprietary (§8).

> **Why a gate, not a score**
>
> A single blended score lets a spectacular median bribe its way past a fatal tail. The gate is absolute: no expected brilliance can buy back a plan that starves in the worst 1-in-20 futures.

THE ANSWERS

6 · From simulation to the five numbers that matter

- **“You can spend Rs X/month”** — the die-with-zero number: the monthly lifestyle at which your money lasts exactly to the end of the plan and no further, judged at the 1-in-20 standard. Overfunded plans can honestly exceed income here — the corpus carries the difference.
- **The corpus you need** — the present value at retirement of every retirement-year expense, including the children’s cost tail, discounted at the recommended mix’s expected return — against the corpus you are heading for on the median simulated path.
- **Your retirement paycheck** — deliberately priced off the bad-market (p20) corpus, not the median: a paycheck that survives a poor market, quoted in today’s rupees.
- **When you’re free** — the earliest retirement age on expected markets, plus a worry-free age that adds a volatility cushion sized to your style and horizon.
- **What moves your plan** — instant scenarios (spend Rs 5,000/10,000/20,000 a month more or less; retire earlier or later). These answer in milliseconds because the engine pre-computes the scenario grid with every plan; the acceleration method is proprietary (§8).

PROTECTION

7 · Corpus-first insurance

We recommend insurance the way an economist would, not the way a seller would: cover only the gap your invested corpus cannot already fund. Term-life need = the present value of your family’s must-have spending for the remaining planning horizon (child-aware, year by year, at the real FD rate) plus outstanding loans and named goals, minus your invested corpus. As your corpus grows, the recommended cover shrinks — to zero, eventually. Pure term cover only; health cover benchmarked to income with a floor; a six-month emergency buffer. We earn nothing from any of it.

WHAT WE KEEP

8 · What we publish, and what we don’t

This paper publishes the full economic model: the ledger, the market mathematics, the safety standard, the selection discipline and every default assumption. Four implementation details are FinLaxmi’s intellectual property and are not published. We name them so the boundary is honest:

- **The glidepath curves.** You know both families’ shapes and intent; the exact anchors, floors and terminal weights that make them behave well across Indian market assumptions are ours.
- **The selection scoring.** The failure gate and style percentiles are public; the ranking metric and tie-break rules among surviving strategies are ours.
- **The scenario acceleration.** Instant what-ifs answer without re-running 250 futures; the pre-computation structure that makes that exact — not approximate — is ours.
- **The AI figure-audit internals.** That every AI number is verified against your plan is public; the matching tolerances and whitelist mechanics that make it hard to fool are ours.

Nothing withheld changes what a number means — only how efficiently or robustly it is produced.

KEEPING IT HONEST

9 · Verification, and the limits we admit

- **Independent replication.** The engine's simulation has been independently re-implemented from its published mathematics and reproduces the workbook's results to the rupee on the frozen futures.
- **Standing audits.** A five-profile validation suite (from over-funded to under-funded lives) and sensitivity checks run against every engine release; the AI audit's failure rate is tracked with an alert threshold.
- **Named errors, never guesses.** If the engine cannot supply a number, the report shows exactly which number is missing. It never silently substitutes an estimate.
- **Before you sign up, ranges.** The instant preliminary report deliberately shows ranges, not false precision — exact figures need your full details.

Known simplifications

- **Tax planning is out of scope by design.** Tax depends on regime choices and legislation; we would rather model nothing than model it glibly. Enter post-tax income and post-tax return assumptions.
- **Returns are taken as you enter them.** Defaults are long-run and deliberately moderate; fund costs are yours to reflect in the return assumption.
- **Gold is regime-modelled, not path-modelled;** property is tracked as an asset you own and a rent you receive, not simulated as a market.

APPENDIX

A · Default assumptions (all editable in your plan)

Assumption	Default	Note
Inflation	6.5% / yr	applied to spending, goals and deflation to today's rupees
Mutual-fund return (μ / σ)	13% / 20%	long-run nominal; lognormal, mean-preserving
FD / debt return (μ / σ)	7.5% / 4%	long-run nominal
Gold return / portfolio share	9% / 10%	per-lifetime regime model
Income growth	8% / yr	salary and bonus, until retirement
Life expectancy	85	plan runs to this age; horizon capped at ~80 years ahead
Child cost stages	0–5 / 5–17 / 17–22	infant / school / college; independence at 23
Child cost bumps	+10% / +20% / +30%	of base spending, per child, per stage
Emergency buffer target	6 months	of must-have spending plus EMIs
Health cover benchmark	max(Rs 10L, 6x monthly income)	floor and income rule, both editable
Safety standard	worst 1-in-20	the 5th percentile of 250 simulated futures

Disclaimer

FinLaxmi provides computational guidance for personal financial planning. It is not investment advice, and no output of the engine is a promise of returns. Simulated futures are models, not predictions; the worst 1-in-20 standard is a discipline, not a guarantee. · Methodology v1.0 · This document corresponds to engine release v16.