Linear, Geometric, Quadratic and Fuzzy Programming for Sectoral Portfolio Allocation in India: An Out-of-Sample Comparison Against the Equally Weighted Benchmark

Raashi Gaur[1], Reyansh Godha[2], Priyansh Motwani[3], Sanaa Ahluwalia[4], Rian Sorap[5]
[1,2,3,4,5]*Student, B.Sc. Finance, Anil Surendra Modi School of Commerce (ASMSOC), SVKM's NMIMS, Mumbai, Maharashtra, India*

***Abstract*—This study tests whether four operations research models can outperform a simple equal split (1/N) when allocating money across ten NIFTY sectoral indices. The models are a Linear Programming model that maximizes return under a mean-absolute-deviation ceiling, a Geometric Programming model that minimizes posynomial risk under a monomial reward constraint, the Markowitz maximum-Sharpe portfolio, and a Zimmermann max–min fuzzy model with triangular returns. All four share the same rules (fully invested, no short selling, at most 30% per sector) and are evaluated on unseen data from January 2019 to September 2026, using a rolling 24-month window, monthly rebalancing and trading costs of 10 basis points. The Linear Programming model achieved the highest Sharpe ratio (0.510 against 0.424 for 1/N) and compound growth (14.2% against 12.6% a year). The Geometric Programming model had the lowest volatility and drawdown and traded about a third as much as the other optimizers; we show that its solution allocates risk in proportion to return-based budgets. Markowitz ranked second and the fuzzy model trailed 1/N. No Sharpe-ratio difference is statistically significant ($p > 0.55$). Every optimizer beat 1/N in 2019–2021 and lagged it during the 2022–2024 rate-hike cycle. Across thirteen robustness checks, only Geometric Programming never fell more than 0.002 behind 1/N on Sharpe ratio while keeping a smaller drawdown.**

***Index Terms*—Fuzzy optimization, geometric programming, linear programming, Markowitz model, NIFTY sectoral indices, portfolio optimization.**

I. INTRODUCTION

Deciding how much money to place in each of several risky assets ranks among the earliest problems that operations research took on, and it remains among the most practically important. Once Markowitz [1] wrote the choice as a balance between average return and variance, the question became a proving ground for nearly every family of optimisation model linear, quadratic, integer, geometric, stochastic and fuzzy alike [2], [3]. In India the question has gained fresh urgency. Low-cost sector index funds and exchange-traded funds now let a small saver buy all of banking, all of IT or all of pharmaceuticals with one purchase, so a growing share of allocation decisions are taken between sectors instead of between individual companies.

Two complications stand in the way. The first is input quality. Optimisers trust their estimates completely, and modest mistakes in forecast returns can drive mean–variance weights into lopsided, jumpy positions [4], [5]. DeMiguel et al. [6] put fourteen optimisers against a simple 1/N split and found that none won reliably once tested on fresh data. The second complication is that Indian markets changed character repeatedly between 2017 and 2026. Cheap money gave way to the COVID-19 collapse and a steep rebound; from 2022 the Reserve Bank tightened policy just as the Russia–Ukraine war jolted commodity prices; and a calmer phase followed. A rule tuned to one of these phases can fail badly in the next.

The four OR approaches respond to noisy inputs in quite different ways. MAD-based LP [7] swaps variance for an absolute-value risk yardstick that outliers disturb less and that large solvers handle easily. GP [8], [9] is built for products and ratios and, as Section IV demonstrates, leads naturally to risk-budgeting allocations that lean only lightly on return forecasts. Fuzzy programming [10], [11] admits that returns are vague and searches for a compromise between competing aims. The Markowitz QP is still

the yardstick of modern portfolio theory. What is missing is a fair contest: each technique is usually tried on its own data, with its own constraints and its own test window, which makes it hard to say which one an investor should actually use.

We stage that contest. All four models are built on the same ten NIFTY sector series, all face the same practical limits, and all are scored on unseen months, after trading charges, against 1/N. Three questions guide the work:

- RQ1. Which of LP, GP, Markowitz QP and fuzzy LP gives the strongest risk-adjusted result on unseen data for an Indian sector portfolio, and does any of them get ahead of 1/N?
- RQ2. Do the gaps remain once we allow for fat tails and day-to-day dependence in returns?
- RQ3. Do the rankings hold across different macro phases and under altered modelling choices?

Our contribution is fourfold. First, the comparison is controlled: one data set, one rulebook and one evaluation protocol, extended through the 2022–2024 rate-hike phase that earlier Indian work could not cover. Second, we write portfolio choice as a GP in which a posynomial risk term is minimised under a monomial reward term, and we prove that the rescaled solution hands every sector a share of total risk equal to a return-driven budget, which ties GP to the risk-parity strand of the literature [12]. Third, we report the diagnostics an OR analyst expects dual prices, parametric runs and membership degrees next to the measures a fund manager expects. Fourth, we show that model rankings depend on the market phase and that what really separates the models is stability, not their best-case result.

Fig. 1 sets out the study design.

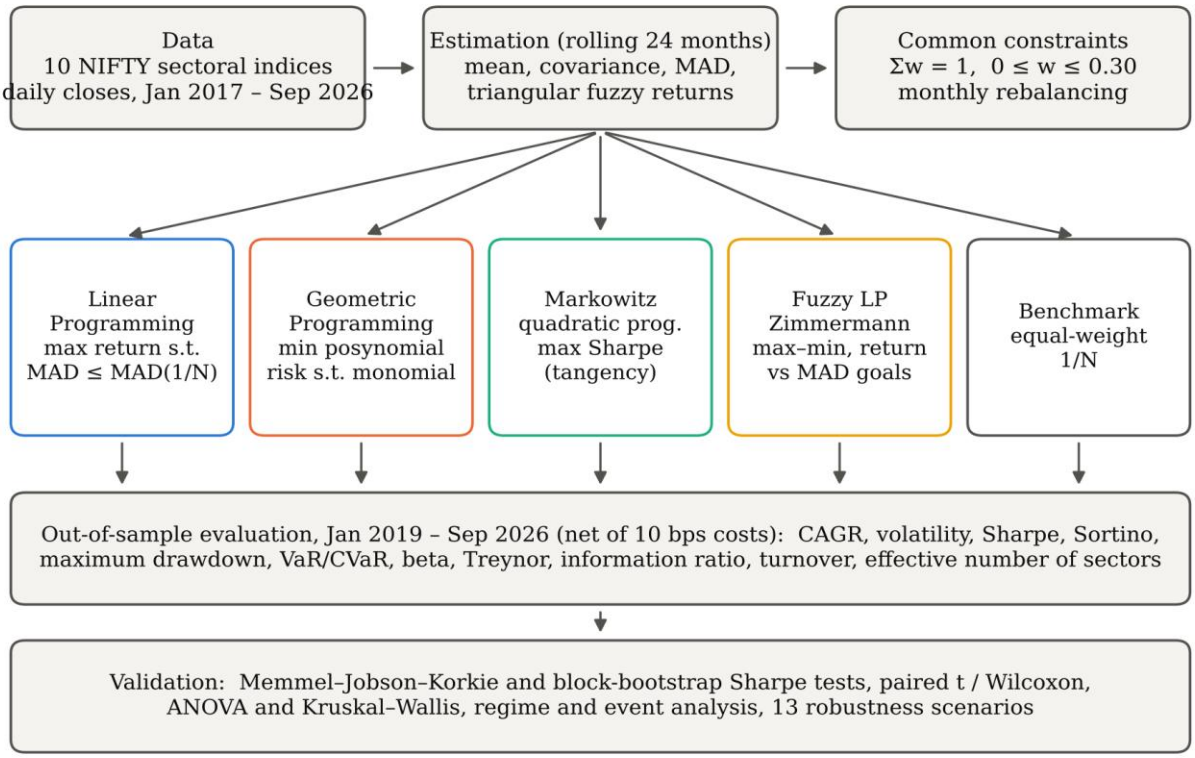

Fig. 1. Research framework.

The rest of the paper runs as follows. Section II surveys earlier work and examines two benchmark studies closely. Section III introduces the data. Section IV sets out the four models and the testing procedure. Section V reports the evidence, Section VI interprets it, and Section VII closes.

II. LITERATURE REVIEW

A. Mean–variance optimisation and its sensitivity to inputs

In Markowitz's [1] framework a portfolio is efficient when no other mix offers the same average return with less variance; tracing such mixes over all target returns gives the efficient frontier, and each point on it solves a convex QP. Allowing borrowing and lending at a riskless rate picks out one frontier point, the tangency portfolio, at which the Sharpe [13] ratio peaks. That ratio divides a linear term by the square root of a quadratic one, so it is normally optimised after the Charnes–Cooper [14] change of variables, which recasts it as a QP [15]. The approach has a familiar weak spot: slight errors in average-return inputs swing the weights sharply [4], and a mistaken mean does roughly ten times the damage of a comparably mistaken variance [5]. Kolm et al. [2] catalogue the remedies proposed over sixty years, among them weight limits, shrinkage estimators and robust formulations.

B. Linear-programming formulations

Interest in LP versions grew because they cope with very large asset lists and can carry integer and logical side conditions. Konno and Yamazaki [7] measured risk by mean absolute deviation instead of variance; the problem then becomes an LP, and when returns are close to normal its answers sit near those of the mean–variance model. Rockafellar and Uryasev [16] later showed that conditional value-at-risk can likewise be minimised through LP. Surveying two decades of such models, Mansini et al. [3] highlight their speed and their ability to absorb trading lots and limits on the number of holdings. Papahristodoulou and Dotzauer [17] set quadratic, maximin and MAD versions against one another and judged the linear ones competitive.

C. Geometric programming and finance

A GP minimises a posynomial while keeping other posynomials below one and monomials equal to one; taking logarithms of the variables makes the whole problem convex [8], [9]. An early finance use came from Jefferson et al. [18], who exploited generalised GP duality to study a bundle of risky ventures under exponential utility and collapsed the problem to a single dimension. Empirical GP work on portfolios has stayed scarce since then, in part because variance is a posynomial only when no two assets co-move negatively. Section III shows that this requirement is met by Indian sector indices, and present-day disciplined-GP software [19] has removed the old computational hurdle.

D. Fuzzy approaches to portfolio choice

Zadeh's [20] fuzzy sets, and Bellman and Zadeh's [21] account of choosing under fuzzy goals and constraints, let a model work with targets and coefficients that are only loosely known. Zimmermann [10] proved that several linear goals with linear membership curves can be pursued together by solving one max–min LP. Within portfolio work, Tanaka et al. [22] described returns through possibility distributions; Carlsson and Fullér [11] defined a possibilistic mean and variance for fuzzy numbers; Inuiguchi and Ramík [23] contrasted fuzzy with stochastic treatments; and Gupta et al. [24] used fuzzy mathematical programming to reflect the goals of individual investors. Treating a forecast as a band instead of a single figure is attractive when market conditions keep changing.

E. The 1/N rule and Indian studies

DeMiguel et al. [6] made 1/N the bar every optimiser has to clear: because it estimates nothing, it makes no estimation mistakes, and that advantage frequently outweighs the theoretical benefit of optimising when data windows are short compared with the number of assets. Indian applications include mean–variance work on NIFTY 50 stocks [25] and a study that pairs deep-learning price forecasts with optimisation for NIFTY thematic sectors [26]. Such papers usually test a single optimiser and seldom compare it formally with 1/N.

F. Close reading of two benchmark studies

Two earlier OR studies of stock selection shaped our design and are examined in Table I: the MAD model of Konno and Yamazaki [7], from which our LP and fuzzy formulations descend, and the integer programme of Alfares and Al-Marhoun [27], a recent application to a single stock-market sector.

Table I. Side-by-side assessment of two earlier or studies of stock selection

Point of comparison	Konno & Yamazaki [7]	Alfares & Al-Marhoun [27]
Practical question	Choosing among a very large list of shares when dense quadratic models were slow to solve	Deciding how many shares of each petrochemical firm to buy with a fixed sum
Setting	Japanese equities (Tokyo exchange)	Saudi equities, ten firms in one sector
Reason for using OR	A risk measure that an LP could handle for hundreds of shares	Many financial and logical conditions acting at once
Goal	Smallest average absolute deviation of portfolio return	Largest total expected return on the shares bought
Restrictions	Return target, budget, limits on each holding	Budget including 0.16% brokerage, dividend floor, ceiling on an average risk index, 30% per-firm ceiling, minimum lot, coverage of 12 or more of 22 products
Model class	Continuous LP	LP with 0–1 variable (mixed-integer)
Decision variables	Weights plus helper deviation terms	Number of shares X_i; 0–1 flag Y_i (firm used) and Z_j (product covered)
Main assumptions	Past return scenarios stand in for future ones; absolute deviation is a fair risk gauge	Three-year (2016–2018) averages stand in for the future; risk index partly supplied by analysts
Solution route	Simplex-type LP solution	Integer programming; software not named
Checking	Compared against mean–variance answers	One in-sample answer; promised sensitivity runs not shown
Headline outcome	MAD mixes close to mean–variance mixes, obtained far faster	Four firms chosen, spanning 14 products; expected gain of 142,669.87 on a 10-million budget
What works well	Linear, fast, needs no distributional assumption	Adds a second axis of spreading risk (firms and products); realistic limits
Weak points	One period only; no remedy for noisy means	Short history; no test on later data; no comparison rule; subjective risk index
How this paper builds on it	MAD-LP tested on unseen months with a risk ceiling matched to 1/N	Head-to-head with 1/N, significance tests, four model types, several market phases

Two points emerge. An LP can encode almost any practical rule, but whether it helps an investor can only be settled by testing it on later data and against a naïve alternative, and neither study does so. And the quality of the risk input matters as much as the solver: a hand-adjusted risk index estimated from three years of history leaves the Saudi solution resting on uncertain numbers. We therefore estimate risk from a rolling window of data, give four OR models identical constraints, and set each of them against 1/N.

G. What is missing

Single-model papers are plentiful, but controlled contests between model families are rare, empirical GP portfolio studies hardly exist, and Indian sector evidence spanning the rate increases after 2022 is thin. Few papers, moreover, pair OR sensitivity diagnostics with formal tests on unseen data. The present study is designed to close these gaps.

III. DATA

We use daily closing values of ten NIFTY sector indices compiled by NSE Indices Ltd. [28] Pharma, Oil and Gas, Media, Bank, Auto, Consumer Durables, IT, Metal, Realty and FMCG from 2 January 2017 through 28 September 2026, a span of 2,413 trading sessions. Fourteen scattered gaps (eleven in Oil and Gas, three in Consumer Durables) were filled with the previous close. Daily returns are computed as $r_{i,t} = P_{i,t}/P_{i,t-1} - 1$. We work at the sector level for two reasons: every index can be bought cheaply through an index fund or ETF, and ten assets is a small enough number for a 24-month window to estimate reliably.

Table II gives the summary figures. Realty, Metal and Consumer Durables grew fastest, at roughly 18% a year (Fig. 2); Media shrank by 5.4% a year and is the lone sector that ended below its starting value (−41.2% overall). FMCG swung least (15.8% annual volatility) and Realty most (28.7%). Tails are heavy everywhere,

with excess kurtosis between 3.1 and 17.1, and every sector except Pharma is skewed to the left. Jarque–Bera statistics reject normality for all ten series ($p < 0.001$), and augmented Dickey–Fuller statistics reject a unit root in each return series ($p < 0.001$). Because returns are so far from normal, we look at MAD, fuzzy and rank-based measures as well as variance.

Table II. Summary Statistics for Daily Sector Returns, January 2017 – September 2026 (Annual Figures Where Indicated)

Sector	Mean (%)	Vol. (%)	CAGR (%)	Cum. (%)	Sharpe	Skew	Ex. kurt.	Worst day (%)	Best day (%)
Pharma	11.76	18.94	10.48	159.6	0.30	0.11	5.57	−8.93	10.37
Oil and Gas	13.28	21.31	11.62	186.3	0.34	−0.83	10.94	−11.80	9.06
Media	−1.69	27.72	−5.40	−41.2	−0.28	−0.34	9.57	−16.37	14.40
Bank	13.93	21.56	12.28	203.1	0.37	−0.80	17.08	−16.73	10.51
Auto	13.32	21.87	11.54	184.4	0.33	−0.24	9.72	−13.85	10.41
Consumer Durables	18.50	19.03	18.16	393.7	0.66	−0.37	6.96	−11.35	7.69
IT	12.80	21.80	10.97	170.9	0.31	−0.23	5.00	−9.58	9.02
Metal	20.14	27.69	17.68	375.2	0.51	−0.53	3.60	−11.60	7.90
Realty	20.81	28.72	18.14	393.0	0.52	−0.33	3.12	−11.35	8.66
FMCG	9.43	15.84	8.51	118.6	0.22	−0.16	12.49	−10.60	8.32

Notes: Mean and Vol. are annualised; Cum. is the cumulative return over the whole sample. Sharpe ratios assume a riskless rate of 6% a year. For every sector, Jarque–Bera $p < 0.001$ and ADF $p < 0.001$.

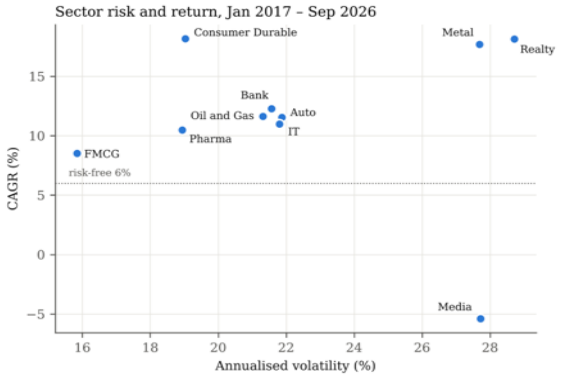

Fig. 2. Risk and compound return of the ten sectors, January 2017 – September 2026.

The correlation matrix appears in Fig. 3. Every one of the 45 sector pairs moves together positively; the weakest link is IT–Media (0.31), the strongest Auto–Consumer Durables (0.69), and the average is 0.49. IT is the most independent sector and so the best source of diversification. This absence of negative co-movement, which persisted in all 93 rolling windows, is exactly what allows variance to be written as a posynomial and keeps the GP model of Section IV-C valid.

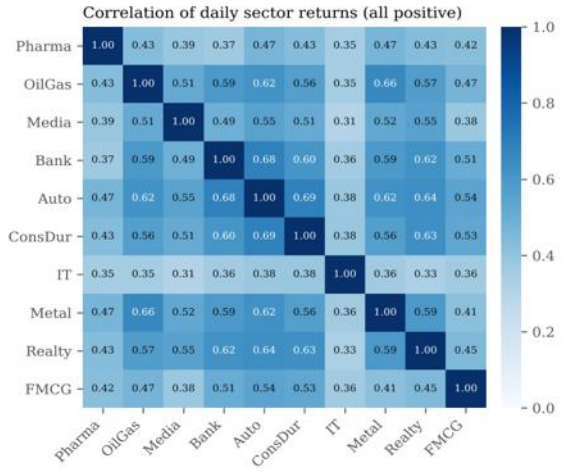

Fig. 3. Correlation matrix of daily sector returns.

IV. METHODOLOGY

A. Symbols and shared constraints

There are $n = 10$ sectors. $w \in \mathbb{R}^n$ holds the portfolio weights, $r_t \in \mathbb{R}^n$ the sector returns on day $t = 1, \dots, T$ of the estimation window, $\bar{r} = \frac{1}{T}\sum_t r_t$ their average, Σ their sample covariance matrix and r_f the riskless rate per day. All models choose weights from the same set

$$\mathcal{W} = \left\{ w \in \mathbb{R}^n : \sum_{i=1}^{n} w_i = 1,\ 0 \le w_i \le u,\ i = 1, \dots, n \right\}, \qquad u = 0.30. \tag{1}$$

The first condition puts all money to work. The second forbids short positions, which retail buyers of Indian sector funds generally cannot take. The third stops any sector from exceeding u, in the same spirit as the 30% per-firm limit of Alfares and Al-Marhoun [27].

B. Model 1: Linear Programming with a MAD risk ceiling

As in Konno and Yamazaki [7], risk is the average size of the portfolio's departures from its mean, $\mathrm{MAD}(w) = \frac{1}{T}\sum_t |(r_t - \bar{r})^\top w|$. Auxiliary variables $d_t \ge 0$ remove the absolute values. We let the LP chase the highest expected return while its MAD may not exceed that of the 1/N mix in effect, the model must out-earn the benchmark without taking more risk:

$$\begin{aligned} \max_{w,\,d} \quad & \bar{\boldsymbol{r}}^\top \boldsymbol{w} \\ & \boldsymbol{d_t} \ge (\boldsymbol{r_t} - \bar{\boldsymbol{r}})^\top \boldsymbol{w}, \quad \boldsymbol{d_t} \ge -(\boldsymbol{r_t} - \bar{\boldsymbol{r}})^\top \boldsymbol{w}, \quad \boldsymbol{t} = \boldsymbol{1}, \dots, \boldsymbol{T}, \\ \text{s.t.} \quad & \frac{\boldsymbol{1}}{\boldsymbol{T}} \sum_{\boldsymbol{t}=\boldsymbol{1}}^{\boldsymbol{T}} \boldsymbol{d_t} \le \mathrm{MAD}(\boldsymbol{1/n}\,\boldsymbol{1}), \\ & \boldsymbol{w} \in \boldsymbol{\mathcal{W}},\ \boldsymbol{d} \ge \boldsymbol{0}. \end{aligned} \tag{2}$$

A 24-month window holds about 490 trading days ($T \approx 490$), so each LP has roughly 500 variables and 980 constraints; we solve it with the dual simplex routine in HiGHS [29]. The dual solution prices both the risk ceiling and every sector limit.

C. Model 2: Geometric Programming with posynomial risk and monomial reward

A monomial is a positive constant time a product of positive variables raised to real powers, $c \prod_i x_i^{a_i}$ with $c > 0$; a posynomial adds several of them. If no covariance σ_{ij} is negative, the variance $w^\top \Sigma w = \sum_i \sum_j \sigma_{ij} w_i w_j$ is a posynomial for $w > 0$. We attach to each sector a reward budget $b_i > 0$, with $\sum_i b_i = 1$, that grows with its past risk-adjusted return:

$$b_i = \frac{\exp(\gamma\, SR_i)}{\sum_j \exp\left(\gamma\, SR_j\right)}, \qquad SR_i = \frac{\bar{r}_i - r_f}{s_i}\sqrt{252}, \tag{3}$$

using $\gamma = 1$ in the main runs; $\gamma = 0$ makes all budgets equal. The geometric programme reads

$$\begin{aligned} \min_{w>0} \quad & \sum_{i=1}^{n}\sum_{j=1}^{n} \sigma_{ij}\, w_i w_j \\ \text{s.t.} \quad & \prod_{i=1}^{n} w_i^{-b_i} \le 1 \end{aligned} \tag{4}$$

with a posynomial to minimise and a monomial restriction. Writing $\hat{w}$ for its optimum, the invested portfolio is the rescaled vector $w^{GP} = \hat{w}/1^\top \hat{w}$. In logarithmic form the restriction says $\sum_i b_i \ln w_i \ge 0$, so it pays the model to spread capital along the budgets, while the objective charges it for risk. Replacing each w_i by e^{y_i} yields a convex problem whose global optimum is found exactly.

Proposition 1. Suppose Σ is positive definite with no negative entries. Then the rescaled GP solution satisfies

$$\frac{w_i^{GP}\,(\Sigma w^{GP})_i}{(w^{GP})^\top \Sigma\, w^{GP}} = b_i, \qquad i = 1, \dots, n, \tag{5}$$

so the fraction of portfolio risk coming from sector i equals its reward budget.

Proof. The objective rises in every w_i, so the restriction must hold with equality at the optimum. Stationarity gives $2(\Sigma\hat{w})_i = \lambda b_i/\hat{w}_i$ for some $\lambda > 0$, that is $\hat{w}_i(\Sigma\hat{w})_i = \lambda b_i/2$. Adding over i and using $\sum_i b_i = 1$ yields $\hat{w}^\top \Sigma \hat{w} = \lambda/2$; dividing the two expressions gives $\hat{w}_i(\Sigma\hat{w})_i/\hat{w}^\top\Sigma\hat{w} = b_i$. Numerator and denominator both scale with the square of w, so rescaling leaves the ratio intact. ▫

The proposition identifies the GP allocation as a risk-budgeting portfolio [12], [30], and with $\gamma = 0$ as the equal-risk-contribution, or risk-parity, portfolio. Returns enter only through a smooth and bounded map into the budgets, so noisy return estimates disturb this model far less than they disturb Markowitz weights. If a rescaled weight breaks the cap u, we pin that sector at u and re-solve the GP over the others with whatever weight is left; this was needed once in 93 rebalances. We use the disciplined-GP layer of CVXPY [19] with the Clarabel interior-point code. On the full sample the computed risk shares agreed with the budgets to within 7.5×10^{-6}, as Proposition 1 predicts.

D. Model 3: Markowitz tangency portfolio

The Markowitz model looks for the mix in $\mathcal{W}$ with the largest ratio $(\bar{r} - r_f 1)^\top w / \sqrt{w^\top \Sigma w}$. Setting $y = \kappa w$ with $\kappa > 0$ [14], [15] turns this into the convex QP

$$\begin{aligned} \min_{y,\kappa} \quad & y^\top \Sigma\, y \\ \text{s.t.} \quad & (\bar{r} - r_f 1)^\top y = 1, \quad 1^\top y = \kappa, \quad 0 \le y_i \le u\,\kappa, \quad \kappa \ge 0, \end{aligned} \tag{6}$$

and the weights are recovered as $w = y/\kappa$. When no admissible mix beats the riskless rate on expected return, no tangency point exists and the model falls back on the minimum-variance mix. That happened at three month-ends (March, April and May 2020), when the COVID-19 slump dominated the 24-month window. We solve the QP with Clarabel; the accompanying workbook also sets the full-sample problem up for Excel Solver, which returns the same optimal Sharpe ratio (0.6450).

E. Model 4: Fuzzy max–min programme with possibilistic returns

No investor knows future returns exactly, so we describe sector i's daily return as a triangular fuzzy number $\tilde{r}_i = (a_i, m_i, c_i)$. Its low, central and high points are the lower quartile, median and upper quartile of the window's returns. Following Carlsson and Fullér [11], its possibilistic mean is

$$\bar{E}(\tilde{r}_i) = \frac{a_i + 4m_i + c_i}{6}. \tag{7}$$

Two loosely stated aims are pursued at once: a high possibilistic return $R(w) = \sum_i \bar{E}(\tilde{r}_i) w_i$ and a low MAD(w). Two preliminary LPs fix the range of each aim. Pushing $R(w)$ to its peak over $\mathcal{W}$ gives R^{max} and the accompanying risk M^{max}; pushing MAD to its floor gives M^{min} and the accompanying return R^{min}. Straight-line membership functions then score how fully each aim is reached,

$$\mu_R(w) = \frac{R(w) - \mathrm{R^{min}}}{\mathrm{R^{max}} - \mathrm{R^{min}}}, \qquad \mu_{\mathrm{M}}(\mathrm{w}) = \frac{\mathrm{M^{max}} - \mathrm{MAD(w)}}{\mathrm{M^{max}} - \mathrm{M^{min}}}, \tag{8}$$

and Zimmermann's [10] max–min rule reduces the two aims to one ordinary LP:

$$\begin{aligned} \max_{\mathrm{w,\,d,\,\lambda}} \quad & \lambda \\ \text{s.t.} \quad & \mu_{\mathrm{R}}(\mathrm{w}) \ge \lambda, \quad \mu_{\mathrm{M}}(\mathrm{w}) \ge \lambda, \quad 0 \le \lambda \le 1, \\ & \mathrm{d_t} \ge \pm(\mathrm{r_t} - \bar{\mathrm{r}})^\top \mathrm{w}, \ \mathrm{t} = 1, \dots, \mathrm{T}, \quad \mathrm{w} \in \mathcal{W}. \end{aligned} \tag{9}$$

At the optimum, $\hat{\lambda}$ measures how far both aims are met together [21]. Over the 93 month-ends its average was 0.646, with a low of 0.581 and a high of 0.735.

F. Benchmark: the 1/N portfolio

The benchmark puts $\mathrm{w_i} = 1/\mathrm{n} = 10\%$ in each sector and restores those weights every month. It estimates nothing and is the usual hurdle for optimised portfolios [6].

Table III sets the five strategies side by side.

Table III. The five strategies at a glance

Strategy	Model family	What it optimises	How risk enters	Role of return forecasts	Solver	Solve time per month-end (mean)
LP	Linear programme	Highest expected return	MAD ceiling	Sample average	HiGHS simplex	0.12 s
GP	Geometric programme	Lowest variance under a monomial reward	Posynomial variance	Only via budgets $\mathrm{b_i}$	Clarabel (log-convex)	0.06 s
Markowitz	Quadratic programme	Highest Sharpe ratio	Variance	Sample average	Clarabel	0.01 s
Fuzzy	Fuzzy two-goal LP	Joint satisfaction λ	MAD as a fuzzy goal	Possibilistic mean	HiGHS simplex (three LPs)	0.18 s
1/N	Rule of thumb	—	—	None	—	—

G. Testing on unseen data

On the final trading day of every month from December 2018 to August 2026 93 occasions in all each model is fitted to the previous 24 months of daily returns. The resulting weights are then left to drift with prices until the following month-end. Evaluation therefore covers 1 January 2019 to 28 September 2026, or 1,919 trading days, none of which is used for fitting. At every rebalance we measure turnover as $\tau = \sum_i \left| \mathrm{w_i^{new}} - \mathrm{w_i^{drift}} \right|$ and subtract 10 basis points per unit of turnover, roughly what trading index ETFs costs. The riskless rate is fixed at 6% a year, close to the

average 91-day Treasury-bill yield reported by the Reserve Bank of India over the sample [31]. Table IV lists every input, its value and where it comes from; inputs that rest on an assumption rather than a published figure are varied in Section V-F.

Table IV. Model inputs and their sources

Input	Value	Source or basis
Sector index closing levels	10 NIFTY sectoral indices, daily, 2 Jan 2017 – 28 Sep 2026 (2,413 days)	NSE Indices Ltd. [28]
Missing closes	14 cells, forward-filled	Authors' data check
Daily return	$P_{i,t}/P_{i,t-1} - 1$	Standard definition
Estimation window	24 months (about 490 trading days)	Design choice; 12 and 36 months tested
Rebalancing	Last trading day of each month, 93 times	Design choice
Sector cap u	30%	Mirrors Alfares and Al-Marhoun [27]; 20%, 40% and no cap tested
Trading cost	10 bps per unit of turnover	Assumption for index ETFs; 0, 25 and 50 bps tested
Risk-free rate	6% a year	Approximate 2017–2026 average 91-day T-bill yield [31]; 4% tested
GP budget tilt γ	1	Design choice; 0 and 2 tested
Fuzzy return bounds	Lower quartile, median, upper quartile of window returns	Design choice following Carlsson and Fullér [11]
Bootstrap	5,000 resamples, 21-day circular blocks	Politis and Romano [32]; Ledoit and Wolf [33]

H. Scoring the strategies

Let $r_{p,t}$ be a strategy's daily return after costs, $\bar{e}$ its average return above the riskless rate and σ_p its daily standard deviation. We record: compound annual growth (CAGR); annual volatility; the Sharpe ratio $SR = 252\,\bar{e}/\left(\sqrt{252}\,\sigma_p\right)$; the Sortino ratio, whose denominator counts only shortfalls below r_f [34]; the deepest peak-to-trough loss, or maximum drawdown (MDD); the Calmar ratio, CAGR divided by |MDD|; one-day 95% value-at-risk and expected shortfall from the empirical distribution [16]; beta against 1/N and the Treynor [35] ratio built on it; the information ratio and tracking error against 1/N; mean turnover; and the effective number of sectors, $N_{eff} = 1/\sum_i w_i^2$.

I. Significance testing

To ask whether a model's Sharpe ratio really differs from that of 1/N we use the Jobson and Korkie [36] statistic with Memmel's [37] correction. Daily returns have heavy tails and short-range dependence, so we add a circular block bootstrap [32] 5,000 resamples of 21-day blocks in the manner of Ledoit and Wolf [33]. Monthly differences from 1/N are checked with paired t and Wilcoxon signed-rank tests, and whether the five strategies share a common monthly mean is checked with one-way ANOVA and the Kruskal and Wallis [38] test.

J. Software and simplifying assumptions

Everything is coded in Python 3 using pandas, NumPy, SciPy and CVXPY 1.9 [39], with HiGHS for LPs and Clarabel for conic problems. The code, the data and an Excel workbook driven by live formulas accompany the paper. We assume that (i) the past 24 months are a fair guide to the next month; (ii) sector exposure can be bought at the closing level through funds or ETFs; (iii) trading costs scale with the amount traded; (iv) taxes, fund tracking error and price impact can be set aside; and (v) the riskless rate stays fixed.

V. RESULTS

A. Full-sample solutions and OR diagnostics

Table V lists the weights each model picks when fitted to the whole sample, and the four answers differ noticeably. LP, Markowitz and fuzzy all fill Consumer Durables the best Sharpe ratio among the sectors to the 30% ceiling and leave Media out. Markowitz is the most concentrated: four sectors account for 97% of the money ($N_{eff} = 4.0$), with Metal and Realty included for their high averages. GP owns every sector, from 4.5% in Media to 14.0% in Consumer Durables, giving $N_{eff} = 9.3$, close to 1/N's 10.0.

Table V. Weights (%) chosen on the full sample, with in-sample statistics

Sector	LP	GP	Markowitz	Fuzzy	1/N
Pharma	11.6	12.3	2.8	0.0	10.0
Oil and Gas	1.5	9.5	0.0	14.0	10.0
Media	0.0	4.5	0.0	0.0	10.0
Bank	10.9	9.6	0.0	15.8	10.0
Auto	0.0	8.6	0.0	0.0	10.0
Consumer Durables	30.0	14.0	30.0	30.0	10.0
IT	16.2	11.9	15.1	7.6	10.0
Metal	14.8	8.6	27.3	17.7	10.0
Realty	9.9	8.3	24.9	5.6	10.0
FMCG	5.0	12.8	0.0	9.3	10.0
Average return (% a year)	16.23	13.87	18.48	16.19	13.23
Volatility (% a year)	16.51	15.92	19.35	17.13	16.69
Sharpe ratio	0.620	0.494	0.645	0.595	0.433
Effective number of sectors	5.7	9.3	4.0	5.4	10.0

Notes: Bold shows the most favourable figure in each row (the smallest, for risk measures). Without trading costs, the Sharpe ratios would be 0.526 (LP), 0.453 (GP), 0.485 (Markowitz), 0.398 (fuzzy) and 0.427 (1/N).

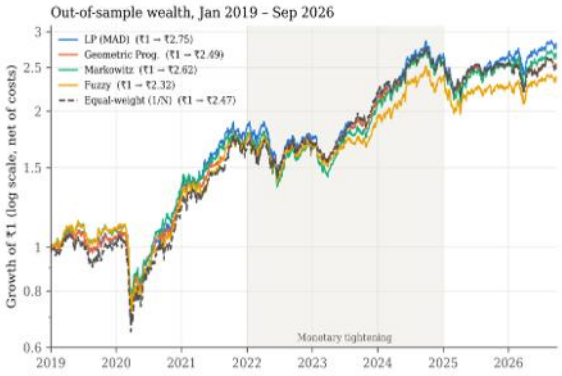

Fig. 6. Out-of-sample growth of ₹1, net of transaction costs (log scale). The shaded band marks the 2022–2024 tightening cycle.

Fig. 6 makes clear that the optimisers built most of their lead in 2019–2021, when they leaned towards the sectors driving the post-COVID rebound and away from Media. 1/N closed the gap during 2022–2024, the lines had met by late 2024, and LP and Markowitz edged ahead once more in 2025–2026. Fig. 7 plots losses from previous peaks. Every strategy hit its low in the March 2020 COVID sell-off, yet each optimiser lost less than 1/N, and GP spent most of the period in a shallower hole than the benchmark.

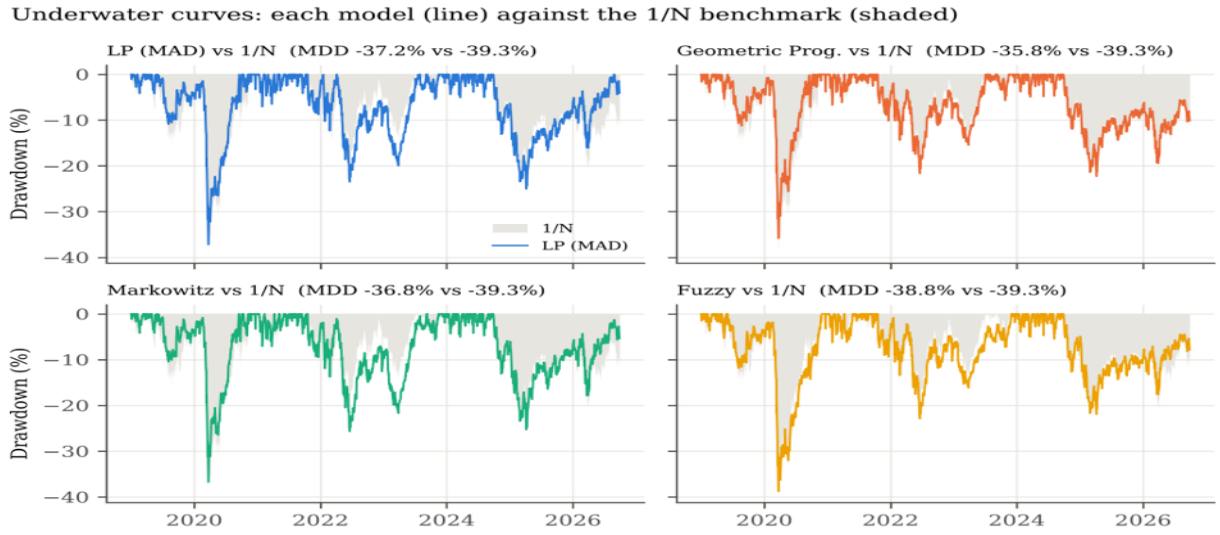

Fig. 7. Drawdowns of each model (line) against the 1/N benchmark (shaded).

C. How the models allocate

Average weights over the test period appear in Fig. 8. LP and Markowitz look almost alike: Pharma, IT and Consumer Durables get the most, while Media (0.2%) and Oil and Gas (about 3%) are all but ignored. Both pressed against the 30% limit at 92 of the 93 month-ends and so held only about 4.5 sectors on average. The fuzzy model put the most into Consumer Durables (21.3% on average). GP kept all ten sectors at every month-end, with average weights from 3.8% (Media) to 13.8% (IT), and reached the limit just once.

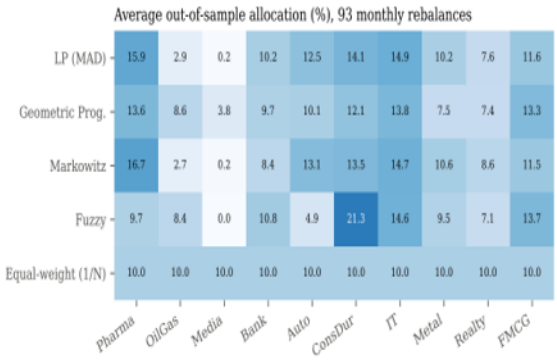

Fig. 8. Average out-of-sample allocation (%) across 93 monthly rebalances.

Fig. 9 follows diversification month by month. For LP, Markowitz and fuzzy the effective number of sectors moved between 3.6 and 5.1 at or near 3.6, the lowest value a 30% limit permits. For GP it ranged from 6.0 to 9.3. This gap in breadth accounts for GP's light trading, its closeness to 1/N and its steadiness when market leadership passed from one sector to another.

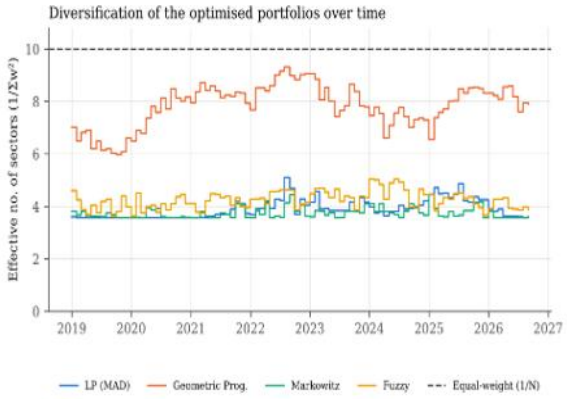

Fig. 9. Effective number of sectors held by each model over time.

D. Are the differences real?

The tests are collected in Table VII. No Sharpe-ratio gap relative to 1/N is statistically meaningful: Memmel-corrected Jobson–Korkie p-values lie between 0.56 and 0.78, and bootstrap p-values between 0.63 and 0.82. LP's lead has a 95% bootstrap interval of [−0.27, 0.43], which easily includes zero. GP's lead is tiny (0.023) but pinned down most tightly a bootstrap standard error of 0.074, against 0.16 to 0.19 for the rest because GP stays so close to 1/N. Across the five strategies, monthly averages cannot be told apart (ANOVA F = 0.017, p = 0.999; Kruskal–Wallis H = 0.567, p = 0.967), and nor can monthly variances (Levene W = 0.128, p = 0.972).

Table VII. Significance Of Each Model's Performance Relative To 1/N

Model	Sharpe gap vs 1/N	JK–Memmel z	p-value	Bootstrap s.e.	Bootstrap p	95% interval	Mean monthly gap (%)	Paired t p	Wilcoxon p
LP	0.086	0.586	0.558	0.179	0.633	[−0.275, 0.433]	0.11	0.672	0.482
GP	0.023	0.330	0.741	0.074	0.761	[−0.122, 0.171]	−0.02	0.879	0.940
Markowitz	0.043	0.285	0.776	0.186	0.816	[−0.327, 0.411]	0.07	0.800	0.496
Fuzzy	−0.043	−0.339	0.735	0.155	0.783	[−0.362, 0.247]	−0.08	0.740	0.962

Notes: Based on 1,919 days and 93 months. Bootstrap uses circular 21-day blocks and 5,000 draws.

That nothing is significant is a finding in its own right. Under eight years of unseen data simply cannot separate a Sharpe gap of 0.09 from chance, which fits DeMiguel et al.'s [6] estimate that sample-based mean–variance rules would need thousands of months of data before they beat 1/N dependably.

E. Market phases and stress episodes

Table VIII and Fig. 10 show that which model wins depends on the market phase. During the low-rate and COVID years (2019–2021) all four optimisers finished ahead of 1/N, with LP reaching a Sharpe ratio of 0.852 and a cumulative gain of 84.7% against 0.638 and 66.3% for 1/N. The 2022–2024 rate-hike years reversed the order: 1/N now had the best Sharpe ratio (0.626) and the mildest drawdown (−18.1%), GP was next (0.563), and LP (0.467), Markowitz (0.475) and fuzzy (0.364) trailed. In the calmer stretch from 2025 to September 2026, only LP and Markowitz ended with gains (4.5% and 2.6%).

Table VIII. Results by market phase and stress episode, after costs

Window	Measure	LP	GP	Markowitz	Fuzzy	1/N
2019–2021 (low rates, COVID)	Sharpe	0.852	0.754	0.791	0.727	0.638
	Cum. return (%)	84.7	73.5	77.6	73.5	66.3
2022–2024 (rate hikes)	Sharpe	0.467	0.563	0.475	0.364	0.626
	Cum. return (%)	42.5	46.8	43.7	35.6	52.0
2025–Sep 2026 (calmer phase)	Sharpe	−0.144	−0.434	−0.188	−0.391	−0.412
	Cum. return (%)	4.5	−2.2	2.6	−1.3	−2.1
COVID sell-off (20 Feb–23 Mar 2020)	Cum. return (%)	−36.3	−34.9	−36.3	−37.7	−36.9
COVID rebound (24 Mar–31 Dec 2020)	Cum. return (%)	91.9	85.2	90.4	82.7	88.1
Russia–Ukraine shock (24 Feb–30 Jun 2022)	Cum. return (%)	−13.5	−9.4	−15.1	−11.8	−7.7

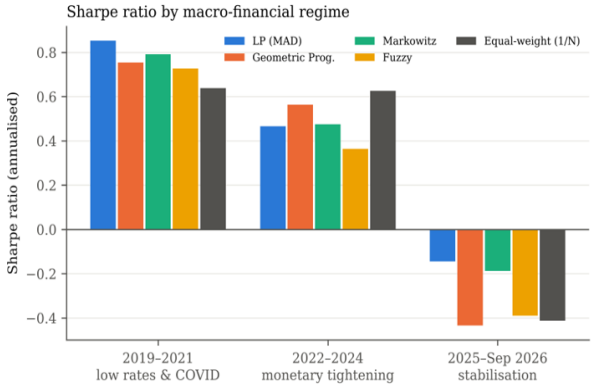

Fig. 10. Sharpe ratio by macro-financial regime.

The short stress windows sharpen this picture. GP lost the least in the COVID sell-off (−34.9%), and LP gained the most in the rebound that followed (91.9%). When the Russia–Ukraine war hit commodity prices, the concentrated optimisers fared worst Markowitz fell 15.1% and LP 13.5%, while 1/N fell 7.7%. Their 24-month windows at that point were filled with the 2020–2021 rally, so they were heavy in exactly the sectors that turned down once rates and input costs climbed. GP's broader spread held its loss to 9.4%.

F. Stress-testing the design

In Table IX the backtest is rerun thirteen times, each time altering a single setting. Three regularities stand out.

Table IX. Sharpe ratios and drawdown ranges when one setting is changed at a time

Variant	LP	GP	Markowitz	Fuzzy	1/N	Top Sharpe	GP at least level with 1/N?
Main settings (cap 30%, 24 m, 10 bps)	0.510	0.447	0.468	0.382	0.424	LP	Yes
Cap 20%	0.479	0.457	0.518	0.446	0.424	Markowitz	Yes
Cap 40%	0.486	0.447	0.436	0.388	0.424	LP	Yes
No cap (100%)	0.382	0.447	0.371	0.404	0.424	GP	Yes
12-month window	0.438	0.450	0.479	0.534	0.424	Fuzzy	Yes
24-month window, tested from 2020	0.572	0.496	0.489	0.387	0.498	LP	≈ (−0.002)
36-month window, tested from 2020	0.280	0.499	0.286	0.351	0.498	GP	Yes
Costs 0 bps	0.526	0.453	0.485	0.398	0.427	LP	Yes
Costs 25 bps	0.486	0.437	0.442	0.357	0.420	LP	Yes

Costs 50 bps	0.445	0.422	0.399	0.317	0.414	LP	Yes
GP budgets with $\gamma = 0$ (risk parity)	0.510	0.446	0.468	0.382	0.424	LP	Yes
GP budgets with $\gamma = 2$	0.510	0.427	0.468	0.382	0.424	LP	Yes
Riskless rate 4%	0.626	0.570	0.564	0.498	0.538	LP	Yes
Eight sectors (Oil & Gas, Cons. Durables dropped)	0.630	0.439	0.579	0.488	0.403	LP	Yes
Range of deepest drawdown over all variants (%)	−35.8 to −41.6	−35.4 to −37.9	−35.0 to −41.2	−33.9 to −39.4	−38.4 to −41.0		

Notes: Each line alters one setting and keeps the rest at their main values. The 36-month window can only be tested from January 2020, so the 24-month window is also shown over that shorter span for a like-for-like comparison.

To begin with, LP posts the top Sharpe ratio in ten of the fourteen lines, including all three alternative cost levels and both alternative universes, so its lead does not hinge on the cost assumption. Next, LP, Markowitz and fuzzy all react strongly to the window length and the cap. Stretching the window to 36 months which lets the 2020–2021 boom dominate the portfolios formed for 2022–2023 drags the Sharpe ratios of LP and Markowitz down to 0.280 and 0.286, far under 1/N's 0.498; removing the cap lets them concentrate even more, and they fall behind 1/N. The fuzzy model, by contrast, comes out on top with a 12-month window, where its median-based estimates update fastest. Finally, GP is the one model that is never beaten by 1/N on Sharpe ratio by more than 0.002 and that has a shallower worst drawdown than 1/N in every line. Switching off its return tilt ($\gamma = 0$, plain risk parity) barely changes its results, which tells us its strength lies in how it budgets risk, not in any return forecast.

G. Computational performance

Each estimation window holds between 490 and 500 trading days. The LP therefore has about 500 continuous variables and 1,000 constraints, the fuzzy model solves three LPs of that size, while the GP and Markowitz problems have only 10 and 11 variables. On a two-core Intel Xeon 2.1 GHz cloud machine, mean solve times per month-end were 0.12 s for LP (maximum 0.25 s), 0.06 s for GP (0.17 s), 0.01 s for Markowitz (0.02 s) and 0.18 s for the fuzzy model (0.24 s). All 372 optimisations of the main backtest finished in about 35 seconds, and the full study, including the thirteen robustness reruns, in under eight minutes. Every solver used is free and open source, and the Markowitz optimum is reproduced exactly in Excel Solver, so the models can be re-run by a student or adviser without commercial software.

VI. DISCUSSION

Which model should an investor pick? It depends on what the investor cares about. Someone after the best reward per unit of risk over the whole period would have done best with LP: an absolute-deviation risk gauge plus a risk ceiling matched to 1/N gave it the top Sharpe, Sortino, Calmar and Treynor ratios, and it stayed in front in most variants. Someone who mainly wants never to fall well behind the simple benchmark, with little trading and modest losses, would have chosen GP. Markowitz added nothing beyond LP; its holdings were similar but a little more volatile and more concentrated. The fuzzy model treats uncertainty in an appealing way, yet under the main settings its sturdier return estimate did not lead to better results though it did come first when re-estimated on a 12-month window.

Why is 1/N so hard to beat? All four optimisers carried less risk than 1/N, but their extra return arrived mainly while yesterday's leading sectors kept leading. When leadership flipped abruptly, as it did in 2022 when higher rates and commodity costs hit the earlier winners, models fitted to the past 24 months were holding the wrong sectors. 1/N, which forecasts nothing, could not make that mistake. Our tests show that over roughly eight years the optimisers' gains are large enough to matter in money terms yet too small to distinguish from luck. This echoes DeMiguel et al.'s [6] international evidence and carries it over to Indian sector data and to LP, GP and fuzzy formulations.

Does the model family matter? The four formulations differ less in their solution algorithms than in how far

they rely on return forecasts. LP and Markowitz feed sample averages straight in and end up concentrated; the fuzzy model uses a median-based figure but still concentrates because both of its aims are linear; GP lets returns in only through bounded budgets and so spreads its risk. The robustness ranking in Table IX follows this order of reliance exactly. The practical message for OR modellers is that deciding what the model may exploit the shape of its objective and constraints matters more than choosing between LP, QP and GP solvers, all of which finished each month's problem in well under a second.

What this means in practice. Three lessons follow for Indian savers and advisers who use sector funds. A per-sector limit is indispensable: without one, the return-seeking models piled into a few sectors and lagged 1/N. A risk-budgeting rule such as the GP model is a sound starting point, since it reduced both losses and trading compared with 1/N and with the other optimisers while giving up nothing in Sharpe ratio. Return-seeking models such as the MAD-LP can help, but they carry phase risk; combining them with a risk-budgeting core, or shortening the estimation window after a change of phase, would temper that risk.

VII. CONCLUSION

We pitted Linear Programming, Geometric Programming, Markowitz Quadratic Programming and Fuzzy Optimisation against one another and against an equal split for dividing money among ten Indian sector indices, using one set of rules and a cost-adjusted test on data from January 2019 to September 2026 that none of the models had seen.

The evidence answers our three questions as follows. On RQ1, MAD-based LP gave the best reward for risk (Sharpe 0.510 against 0.424 for 1/N; CAGR 14.2% against 12.6%), GP gave the lowest risk and least trading, Markowitz fell between the two, and the fuzzy model lagged 1/N; all four-cut volatility, expected shortfall and drawdown compared with 1/N. On RQ2, no Sharpe gap is statistically significant, so the case that optimisation beats equal weighting is economically suggestive but not proven. On RQ3, the winner changes with the market phase the optimisers led in 2019–2021 and 2025–2026, 1/N led during the 2022–2024 rate increases and GP alone kept level with or ahead of 1/N (never more than 0.002 behind in Sharpe ratio) in every variant. On the method side, we showed that a geometric programme with posynomial risk and monomial reward produces a risk-budgeting portfolio in closed form, which gives GP an intuitive financial meaning.

Limitations. We work with index levels, not with real funds, and so ignore ETF expense ratios and tracking error. The riskless rate is held constant and trading costs are taken to be proportional. Only one investor profile, with a 30% limit, is examined. The test period, while spanning three market phases, is still too brief for firm statistical conclusions. Finally, the GP model needs non-negative covariances; they held here but may not for other asset mixes.

Next steps. Promising extensions include adding bonds and gold, whose negative co-movement with equities would call for signomial programming; feeding regime-switching or machine-learning forecasts into the LP and fuzzy models; LP formulations that cap conditional value-at-risk; multi-period stochastic programmes that model trading costs explicitly; and combined portfolios built around a GP risk-budgeting core with an LP return-seeking satellite.

DECLARATIONS

Funding: No dedicated grant supported this research.
Competing interests: None.
Data and code: The index series can be obtained from NSE Indices Ltd. Our data file, Python scripts and Excel workbook will be shared by the corresponding author on reasonable request [or: are posted at repository URL].
Author contributions: Raashi Gaur: conceptualisation, methodology, project administration, writing – original draft, writing – review and editing. Reyansh Godha: data curation, formal analysis (Linear and Geometric Programming models). Priyansh Motwani: software, formal analysis (Markowitz and Fuzzy models). Sanaa Ahluwalia: investigation (literature review), writing – original draft. Rian Sorap: visualisation, validation (statistical tests and robustness checks), writing – review and editing. All authors read and approved the final manuscript.

REFERENCES

[1] H. Markowitz, "Portfolio selection," *J. Finance*, vol. 7, no. 1, pp. 77–91, 1952, doi: 10.1111/j.1540-6261. 1952.tb01525. x.

[2] P. N. Kolm, R. Tütüncü, and F. J. Fabozzi, "60 years of portfolio optimization: Practical challenges and current trends," *Eur. J. Oper. Res.*, vol. 234, no. 2, pp. 356–371, 2014.

[3] R. Mansini, W. Ogryczak, and M. G. Speranza, "Twenty years of linear programming-based portfolio optimization," *Eur. J. Oper. Res.*, vol. 234, no. 2, pp. 518–535, 2014, doi: 10.1016/j.ejor.2013.08.035.

[4] R. O. Michaud, "The Markowitz optimization enigma: Is 'optimized' optimal?" *Financ. Anal. J.*, vol. 45, no. 1, pp. 31–42, 1989.

[5] V. K. Chopra and W. T. Ziemba, "The effect of errors in means, variances, and covariances on optimal portfolio choice," *J. Portfolio Manage.*, vol. 19, no. 2, pp. 6–11, 1993, doi: 10.3905/jpm.1993.409440.

[6] V. DeMiguel, L. Garlappi, and R. Uppal, "Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?" *Rev. Financ. Stud.*, vol. 22, no. 5, pp. 1915–1953, 2009, doi: 10.1093/rfs/hhm075.

[7] H. Konno and H. Yamazaki, "Mean-absolute deviation portfolio optimization model and its applications to Tokyo stock market," *Manage. Sci.*, vol. 37, no. 5, pp. 519–531, 1991, doi: 10.1287/mnsc.37.5.519.

[8] R. J. Duffin, E. L. Peterson, and C. Zener, *Geometric Programming: Theory and Application*. New York, NY, USA: Wiley, 1967.

[9] S. Boyd, S.-J. Kim, L. Vandenberghe, and A. Hassibi, "A tutorial on geometric programming," *Optim. Eng.*, vol. 8, no. 1, pp. 67–127, 2007, doi: 10.1007/s11081-007-9001-7.

[10] H.-J. Zimmermann, "Fuzzy programming and linear programming with several objective functions," *Fuzzy Sets Syst.*, vol. 1, no. 1, pp. 45–55, 1978.

[11] C. Carlsson and R. Fullér, "On possibilistic mean value and variance of fuzzy numbers," *Fuzzy Sets Syst.*, vol. 122, no. 2, pp. 315–326, 2001.

[12] S. Maillard, T. Roncalli, and J. Teïletche, "The properties of equally weighted risk contribution portfolios," *J. Portfolio Manage.*, vol. 36, no. 4, pp. 60–70, 2010, doi: 10.3905/jpm.2010.36.4.060.

[13] W. F. Sharpe, "Mutual fund performance," *J. Bus.*, vol. 39, no. 1, pp. 119–138, 1966.

[14] Charnes and W. W. Cooper, "Programming with linear fractional functionals," *Naval Res. Logist. Quart.*, vol. 9, nos. 3–4, pp. 181–186, 1962.

[15] G. Cornuéjols and R. Tütüncü, *Optimization Methods in Finance*. Cambridge, U.K.: Cambridge Univ. Press, 2007.

[16] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *J. Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[17] C. Papahristodoulou and E. Dotzauer, "Optimal portfolios using linear programming models," *J. Oper. Res. Soc.*, vol. 55, no. 11, pp. 1169–1177, 2004.

[18] T. R. Jefferson, C. H. Scott, and J. M. Cozzolino, "The analysis of risky portfolios by geometric programming," *Z. Oper. Res.*, vol. 23, no. 5, pp. 207–217, 1979, doi: 10.1007/BF01919485.

[19] Agrawal, S. Diamond, and S. Boyd, "Disciplined geometric programming," *Optim. Lett.*, vol. 13, no. 5, pp. 961–976, 2019.

[20] L. A. Zadeh, "Fuzzy sets," *Inf. Control*, vol. 8, no. 3, pp. 338–353, 1965, doi: 10.1016/S0019-9958(65)90241-X.

[21] R. E. Bellman and L. A. Zadeh, "Decision-making in a fuzzy environment," *Manage. Sci.*, vol. 17, no. 4, pp. B141–B164, 1970.

[22] H. Tanaka, P. Guo, and I. B. Türksen, "Portfolio selection based on fuzzy probabilities and possibility distributions," *Fuzzy Sets Syst.*, vol. 111, no. 3, pp. 387–397, 2000.

[23] M. Inuiguchi and J. Ramík, "Possibilistic linear programming: A brief review of fuzzy mathematical programming and a comparison with stochastic programming in portfolio selection problem," *Fuzzy Sets Syst.*, vol. 111, no. 1, pp. 3–28, 2000.

[24] P. Gupta, M. K. Mehlawat, and A. Saxena, "Asset portfolio optimization using fuzzy mathematical programming," *Inf. Sci.*, vol. 178, no. 6, pp. 1734–1755, 2008.

[25] S. Agarwal and N. B. Muppalaneni, "Portfolio optimization in stocks using mean–variance optimization and the efficient frontier," *Int. J.*

Inf. Technol., vol. 14, no. 6, pp. 2917–2926, 2022, doi: 10.1007/s41870-022-01052-2.

[26] J. Sen, S. Mondal, and S. Mehtab, "Portfolio optimization on NIFTY thematic sector stocks using an LSTM model," *arXiv preprint arXiv:2202.02723*, 2021. [Online]. Available: https://arxiv.org/abs/2202.02723

[27] H. K. Alfares and M. A. Al-Marhoun, "Stock market portfolio selection using linear programming," in *Proc. 51st Annu. Conf. Decision Sciences Institute*, 2020, pp. 1096–1104.

[28] NSE Indices Ltd., "Historical index data: NIFTY sectoral indices," National Stock Exchange of India, 2026. [Online]. Available: https://www.nseindia.com/reports-indices-historical-index-data

[29] Q. Huangfu and J. A. J. Hall, "Parallelizing the dual revised simplex method," *Math. Program. Comput.*, vol. 10, no. 1, pp. 119–142, 2018.

[30] T. Roncalli, *Introduction to Risk Parity and Budgeting*. Boca Raton, FL, USA: Chapman & Hall/CRC, 2013.

[31] Reserve Bank of India, *Handbook of Statistics on the Indian Economy 2024–25*. Mumbai, India: RBI, 2025.

[32] D. N. Politis and J. P. Romano, "A circular block-resampling procedure for stationary data," in *Exploring the Limits of Bootstrap*, R. LePage and L. Billard, Eds. New York, NY, USA: Wiley, 1992, pp. 263–270.

[33] O. Ledoit and M. Wolf, "Robust performance hypothesis testing with the Sharpe ratio," *J. Empir. Finance*, vol. 15, no. 5, pp. 850–859, 2008.

[34] F. A. Sortino and L. N. Price, "Performance measurement in a downside risk framework," *J. Investing*, vol. 3, no. 3, pp. 59–64, 1994.

[35] J. L. Treynor, "How to rate management of investment funds," *Harvard Bus. Rev.*, vol. 43, no. 1, pp. 63–75, 1965.

[36] J. D. Jobson and B. M. Korkie, "Performance hypothesis testing with the Sharpe and Treynor measures," *J. Finance*, vol. 36, no. 4, pp. 889–908, 1981, doi: 10.1111/j.1540-6261. 1981.tb04891. x.

[37] C. Memmel, "Performance hypothesis testing with the Sharpe ratio," *Finance Lett.*, vol. 1, pp. 21–23, 2003.

[38] W. H. Kruskal and W. A. Wallis, "Use of ranks in one-criterion variance analysis," *J. Amer. Statist. Assoc.*, vol. 47, no. 260, pp. 583–621, 1952.

[39] S. Diamond and S. Boyd, "CVXPY: A Python-embedded modeling language for convex optimization," *J. Mach. Learn. Res.*, vol. 17, no. 83, pp. 1–5, 2016.