

Stationary Portfolio Goal-Based Optimization with Regime-Switching

Ahmed AlQubaisi^{1,2,3}, Alexandre Antonov^{1,2}, Dmitry Muravey¹

¹ Strategy and Planning Department, ADIA, Abu Dhabi, UAE

² ADIA Lab, Abu Dhabi, UAE

³ Mathematics Department, Khalifa University, Abu Dhabi, UAE

September 30, 2026

Abstract

We formulate and solve a stationary goal-based portfolio optimization problem under regime-switching market drivers. Portfolio weights depend only on the current regime, so the weight and return processes are strictly stationary; ergodicity then aligns the along-path frequency with which a target return is attained with the model-implied utility, turning historical backtests into direct empirical estimators of the objective — an identity that time-dependent frameworks lack.

Analytically, the utility reduces in closed form to a low-dimensional integral over regime occupation times, yielding a two-level optimization scheme whose complexity scales with the number of regimes rather than the number of assets. The inner sub-problem admits a non-convex hard-equality form and a convex soft-inequality relaxation; a sign-analysis criterion in the target rate r characterises when a plain convex QP suffices in place of the DCCP fallback. Empirically, in a proof-of-concept long-only setting on 25 years of GICS-sector returns, the regime-conditional strategy exceeds the passive S&P 500 baseline by 15–40% of goal-reaching probability at multi-year horizons and aggressive targets, 2σ -significant under paired Monte-Carlo standard errors derived in closed form; a state-blind single-regime collapse captures only part of the lift, isolating regime conditioning as the source of the gain. The unconstrained optimum is intrinsically concentrated — Browne’s phenomenon for probability-of-reaching utilities — and a tracking-and-turnover-bounded *Hedged* variant recovers most of the gain in an institutionally tradable form. Analytical robustness bounds indicate that the headline magnitudes survive natural perturbations of the VIX thresholds, calibration windows, and sample length.

Keywords: goal-based portfolio optimization; regime switching; stationary allocation; occupation times; moment generating function; goal probability; backtesting.

JEL classification: G11, G17, C58, C61.

MSC 2020: 91G10, 91G60, 60J27.

1 Introduction

Our main empirical finding is stated first, before the machinery that produces it. On 25 years of GICS-sector daily returns, in the long-only class, a stationary regime-conditional goal-based allocation beats a passive S&P 500 baseline by 15–40% of goal-reaching probability at multi-year horizons and aggressive target rates, peaking at $\approx 40\%$ for $(r = 15\%, T \in \{3, 5\}Y)$, and is 2σ -significant under the paired Monte-Carlo standard errors derived in Appendix C. Two qualifications should be borne in mind from the outset. First, this is a proof of concept on a small universe ($N = 11$ sectors) rather than a directly deployable trading strategy: the unconstrained *Optimal* weights are highly concentrated, and institutional deployment would proceed through the tracking-and-turnover-bounded *Hedged* variant developed below, which retains most of the probability gain in a tradable form. Second, a state-blind single-regime Gaussian collapse of the same optimization captures only part of the lift, so it is the regime conditioning — not the optimization apparatus by itself — that is the source of the gain.

The generalisation of this proof-of-concept magnitude to other samples rests on a structural feature of the stationary formulation that dynamic frameworks lack. When the policy $u(z)$ depends only on the current

regime label — not on wealth, calendar time, or path history — both the weight process and the induced return process are strictly stationary; by ergodicity of the regime chain, the along-path frequency with which the target is attained is a direct empirical estimator of the model-implied utility itself. Historical performance and model expectation are, in this framework, the same object. Time-dependent HJB policies break this identity: the same (t, Π, z) configuration is visited only once along a historical path, so the backtest and the objective drift apart. In the stationary regime-switching setup they do not, and this is the sense in which the proof-of-concept magnitudes are informative — the mechanism carries over to larger universes and different partitions, even where the specific numbers do not.

The goal-based utility is not merely a reparameterisation of Sharpe: it produces two effects that a mean-variance perspective does not anticipate. First, it deforms the return distribution rightward, pushing mass past the threshold rT . Sharpe, being a two-moment summary, averages away precisely the distributional detail that the objective is optimising, and the same strategy that exhibits a 2σ -significant probability gain over the S&P 500 shows only a $\approx 0.8\sigma$ Sharpe gain (see the remark in Section 5). Operationally, the statistical evidence for a goal-based strategy therefore lives in the probability channel rather than the Sharpe channel, which motivates the closed-form paired standard errors of Appendix C and the paired Monte-Carlo estimator we use to handle overlapping-window autocorrelation. Second, the same objective concentrates: under long-only constraints with moderate cross-sectional Sharpe dispersion, the optimum places nearly 100% of the portfolio on the top-Sharpe sector per regime. This is Browne’s [Browne, 1999] phenomenon in a stationary setting, intrinsic to probability-of-reaching utilities and unrelated either to our particular reduction or to the small size of the universe. The concentration is what motivates the *Hedged* and box-bounded variants developed below — not because the unconstrained answer is incorrect, but because it is not one an institution can trade. The paper therefore must deliver on two fronts: inference machinery that makes the probability signal statistically legible, and a bounded variant that recovers a tradable allocation from the corner solution.

Regime-switching itself is a deliberate coarse-graining of a continuous stochastic drift-and-volatility model: rather than let drifts and volatilities respond smoothly to a continuous factor such as the VIX, one collapses the state space into a handful of discrete regimes carrying intuitive labels (*expansion*, *neutral*, *crisis*). This choice is favoured in practice because each regime has an economic interpretation, calibration accepts either historical data or expert views, and structural breaks and tail behaviour are modelled explicitly rather than smoothed. Analytically the discreteness comes at a cost — states are latent, transitions are discrete, and the objective is non-convex in u — and each element is addressed in turn: the regime label is treated as observed via VIX thresholds; the Continuous-Time Markov Chain (CTMC) generator Λ is estimated from empirical transitions; and the non-convexity is decomposed in Section 3 into a K -dimensional outer search over per-regime variances coupled with a convex inner QP. One technical remark: the regime process z_t is assumed independent of the Brownian motion W_t at the innovation level, but long-horizon correlations between regimes and returns arise endogenously through the regime-conditional drifts; calibration details are deferred to Section 5.

Regime-switching portfolio optimization has an extensive literature against which we position our approach. On the identification side, historical-data and sliding-window methods [Nystrup et al., 2016] treat the regime label as an object to be recovered from returns and subsequently plugged into a downstream allocation problem. On the control side, [Zhou and Yin, 2003] formulates observed-regime Markowitz allocation as a linear-quadratic problem solved via a system of HJB equations, [Sass and Haussmann, 2004] infers regimes from data within the control loop, and [Bielecki and Pliska, 1999] treats an infinite-horizon growth-versus-risk-aversion trade-off in continuous time. More recent work uses deep neural networks: [Bauman et al., 2024] learns a policy mapping regime-probability vectors, current wealth, and time-to-horizon into portfolio weights. None of these strands combines the two elements that characterise our formulation: the policy depends only on the current regime label (so the induced process is strictly stationary and the backtest of Section 5 directly estimates the model utility), and the objective is the historical probability of reaching a target return (rather than a smooth utility of terminal wealth). Both choices motivate the machinery developed below.

The technical contributions are organised around the two-level structure that the objective naturally admits. The *outer* search runs over the per-regime variance vector $v \in \mathbb{R}^K$ — a low-dimensional non-convex problem — and is solved by generalized simulated annealing (GSA) [Tsallis and Stariolo, 1996, Xiang et al., 1997]; at $K = 3$ the global optimum is reached reliably and at negligible cost, and this level is

treated as solved. The *inner* sub-problem, per regime, maximises expected return at a strictly fixed variance under unit-budget and box constraints. Its natural hard-equality form is non-convex and calls for a difference-of-convex-programming (DCCP) solver, which internally reduces to a small number of convex-QP iterations — fast in absolute terms but only locally convergent. Our main technical result is that a soft (inequality) relaxation of the same problem is a plain convex QP with a global optimum, and yields the same outer optimum for every target rate r below an explicit sign-analysis threshold. Every (T, r) cell in the proof-of-concept experiments lies within this range, so the fast convex path is used throughout; aspirational targets above the threshold require the DCCP fallback, applied per-regime. On the empirical side, we contribute closed-form paired standard errors for Sharpe and goal-probability differences together with the paired Monte-Carlo estimator noted above; the two components jointly make the probability-channel signal statistically legible at proof-of-concept scale. Finally, a diversification-sensitivity analysis prices governance-driven weight bounds against goal-reaching probability, providing the framework-level response to the deployment question raised by concentration.

A note on scope and organisation. The empirical block should be read as a proof of concept of the framework rather than a live trading claim: the sector-ETF universe is small ($N = 11$), the calibrated regime model is a deliberate simplification of the true market dynamics, transaction costs are reported but not priced explicitly, and the unconstrained *Optimal* allocations would require the tracking, turnover, or diversification constraints of the Hedged and bounded variants before institutional deployment. Section 5 closes with an extensions checklist — larger universes, long-short and leverage, custom regime labels, additional convex constraints via Ω_n , transaction-cost penalties — under which the framework carries over unchanged. The paper is organised as follows: Section 2 sets up the regime-switching market model; Section 3 formulates the stationary goal-based objective and gives the occupation-time reduction; Section 4 develops the two-level scheme with the soft/hard analysis and the numerical evaluation of the utility; Section 5 reports the empirical results, discusses concentration, and prices the diversification–probability trade-off. Appendix A collects the moment-generating-function derivation, Appendix B the sign criterion for the soft form at general K , Appendix C the paired-difference standard-error formulas, and Appendix D analytical robustness bounds on the empirical findings.

2 Regime-Switching Market Model

The market comprises N risky assets driven by an M -dimensional Brownian motion W_t and modulated by a discrete regime process $z_t \in \{1, \dots, K\}$, whose states represent distinct macroeconomic environments (e.g. expansion, recession, crisis). Each regime n carries a drift vector $\mu^{(n)} \in \mathbb{R}^N$ and a volatility matrix $\sigma^{(n)} \in \mathbb{R}^{N \times M}$, with associated instantaneous covariance $C^{(n)} \equiv \sigma^{(n)} \sigma^{(n)\top} \in \mathbb{R}^{N \times N}$. Using the shorthand

$$\mu(z) = \sum_{n=1}^K \mu^{(n)} \mathbb{1}_{z=n}, \quad \sigma(z) = \sum_{n=1}^K \sigma^{(n)} \mathbb{1}_{z=n}, \quad C(z) = \sum_{n=1}^K C^{(n)} \mathbb{1}_{z=n},$$

we specify regime dynamics, asset dynamics, and portfolio return conventions in turn.

2.1 Regime Dynamics

The regime process z_t is a time-homogeneous, irreducible continuous-time Markov chain with generator $\Lambda \in \mathbb{R}^{K \times K}$ satisfying $\Lambda_{ij} \geq 0$ for $i \neq j$ and $\Lambda \mathbb{1} = 0$. The transition matrix over horizon Δt is

$$(1) \quad P_{ij}^{(\Delta t)} = \mathbb{P}[z_{t+\Delta t} = j \mid z_t = i] = (e^{\Lambda \Delta t})_{ij}, \quad e^{\Lambda \Delta t} \mathbb{1} = \mathbb{1}.$$

Irreducibility guarantees a unique stationary distribution $\pi_\infty \in \mathbb{R}^K$ (row vector, $\pi_\infty \Lambda = 0$, $\pi_\infty \mathbb{1} = 1$), and the limiting transition matrix has identical rows:

$$(2) \quad \lim_{t \rightarrow \infty} e^{\Lambda t} = \mathbb{1} \pi_\infty^\top.$$

π_∞ encodes the long-run occupation probabilities of the regimes and is the natural initial law for a time-invariant policy: under $z_0 \sim \pi_\infty$ the chain z_t is strictly stationary, so any state-only weight map $u(z)$ induces strictly stationary weight and return processes.

To fix the picture, a representative $K = 3$ generator on our sector data (units of year⁻¹; regimes ordered as expansion, neutral, crisis) looks like

$$(3) \quad \Lambda \approx \begin{pmatrix} -0.5 & 0.4 & 0.1 \\ 0.6 & -1.0 & 0.4 \\ 0.2 & 1.3 & -1.5 \end{pmatrix}, \quad \pi_\infty \approx (0.60, 0.25, 0.15),$$

so that expansion has an average duration of $1/|\Lambda_{11}| = 2$ years and, on leaving, transitions to neutral 80% of the time and directly to crisis 20%; crisis has an average duration of ≈ 8 months and almost always exits to neutral. The stationary weights match the empirical VIX-based occupation frequencies reported in Section 5.

2.2 Asset Dynamics

W_t is an M -dimensional standard Brownian motion, taken to be independent of z_t at the innovation level (the shocks driving regime transitions and those driving asset returns are orthogonal). Per-asset arithmetic returns follow the regime-dependent SDE

$$(4) \quad \frac{dS_t^i}{S_t^i} = \mu(z_t)_i dt + [\sigma(z_t) dW_t]_i, \quad i = 1, \dots, N,$$

with instantaneous covariance matrix $C(z_t)$ per regime.

2.3 Portfolio Value and Return Conventions

Given regime-dependent weights $u(z_t) = \sum_n u^{(n)} \mathbb{1}_{z_t=n}$ with $u^{(n)} \in \mathbb{R}^N$, the portfolio value Π_t satisfies

$$(5) \quad \frac{d\Pi_t}{\Pi_t} = u(z_t)^\top \mu(z_t) dt + u(z_t)^\top \sigma(z_t) dW_t.$$

We work throughout with two return conventions over horizon T (with $X_0 = 0$ and $\Pi_0 > 0$ fixed): the **simple** cumulative return $X_T = \int_0^T d\Pi_s/\Pi_s$ and the **log** return $R_T = \log(\Pi_T/\Pi_0)$. From (5) and Itô's lemma,

$$(6) \quad X_T = \int_0^T u(z_s)^\top \mu(z_s) ds + \int_0^T u(z_s)^\top \sigma(z_s) dW_s$$

and

$$(7) \quad R_T = \int_0^T \left(u(z_s)^\top \mu(z_s) - \frac{1}{2} u(z_s)^\top C(z_s) u(z_s) \right) ds + \int_0^T u(z_s)^\top \sigma(z_s) dW_s.$$

Fixing a target rate $r \in \mathbb{R}$ and horizon $T > 0$, the object of Section 3 is the historical probability that either return exceeds rT .

Remark 1 (Effective regime/return correlation). *Although z_t and W_t are independent at the innovation level, an effective correlation between the regime path and cumulative returns emerges over finite horizons through the drift channel: when a regime carries a large $|\mu^{(n)}|$ (as is typical for boom or crisis states), realized returns become strongly informative about the sequence of regimes visited. Conditioning on the observed X_T tightens the posterior over the regime path, even though the underlying Brownian and Markov-chain innovations remain orthogonal. This is the mechanism by which the discrete-state model reproduces empirically observed return/regime co-movement without introducing explicit instantaneous correlation.*

2.4 Notation summary

For the reader's convenience we collect the symbols used throughout the paper.

Symbol	Meaning
<i>Model and market</i>	
N, K	number of risky assets; number of regimes
$z_t \in \{1, \dots, K\}$	regime label at time t (CTMC with generator Λ)
Λ	CTMC generator; $\Lambda_{ij} \geq 0$ ($i \neq j$), rows sum to zero
π_∞	stationary distribution of z_t : $\pi_\infty \Lambda = 0$, $\pi_\infty \mathbf{1} = 1$
W_t	M -dimensional Brownian motion, independent of z_t
$\mu^{(n)}, \sigma^{(n)}, C^{(n)}$	per-regime drift vector, volatility matrix, and covariance $C^{(n)} = \sigma^{(n)} \sigma^{(n)\top}$
Π_t, X_T, R_T	portfolio value; simple cumulative return $X_T = \int_0^T d\Pi_s / \Pi_s$; log return $R_T = \log(\Pi_T / \Pi_0)$
<i>Policy, expectations, utility</i>	
$u^{(n)} \in \mathbb{R}^N$	per-regime weight vector; $u = \{u^{(n)}\}_{n=1}^K$
Ω_n, Ω	per-regime feasible set (linear/box constraints); $\Omega = \prod_n \Omega_n$
$\mathbb{E}_\infty[\cdot], \mathbb{P}_\infty[\cdot]$	expectation / probability under $z_0 \sim \pi_\infty$ (hot start)
r, T	annualized target return rate; horizon
$U_T[u], U_T^{(L)}[u]$	goal-based utility (simple / log return): probability of hitting rT over T
$\mathcal{S}_T[u]$	path-conditional Sharpe ratio in the closed-form utility $U_T = \mathbb{E}_\infty[\mathcal{N}(\mathcal{S}_T)]$
<i>Occupation times and two-level reduction</i>	
$\mathcal{T}_n, \tau_n$	regime local time $\mathcal{T}_n = \int_0^T \mathbf{1}_{z_s=n} ds$; normalized fraction $\tau_n = \mathcal{T}_n / T$
ν_T	law of $\tau = (\tau_1, \dots, \tau_K)$ on the simplex Δ^{K-1}
m_n, v_n	per-regime mean $u^{(n)\top} \mu^{(n)}$ and variance $u^{(n)\top} C^{(n)} u^{(n)}$
$\bar{\rho}_n(v_n), \rho_n(v_n)$	soft-form (inequality QP) and hard-form (equality DCCP) inner frontier drift at variance v_n
$\bar{M}(\tau, v), \bar{V}(\tau, v)$	τ -weighted per-path drift $\sum_n \tau_n \bar{\rho}_n(v_n)$ and variance $\sum_n \tau_n v_n$
$\bar{M}^L(\tau, v)$	log-return analogue $\bar{M} - \frac{1}{2} \bar{V}$
$v_n^{\min}, v_n^{\max}$	admissible variance range on Ω_n ; hyper-box $Q^{(v)} = \prod_n [v_n^{\min}, v_n^{\max}]$
$\bar{v}_n^{\max}, \bar{m}_n^{\max}$	argmax and max of $\bar{\rho}_n(v_n)$ on $[v_n^{\min}, v_n^{\max}]$
<i>Estimation and inference</i>	
$P_{T,r}[\Pi]$	historical rolling-window frequency of goal attainment (35)
ρ (correlation)	correlation between two paired return series in Appendix C (distinct from the frontier drift)
CV	coefficient of variation, robust std over mean

3 Stationary Goal-Based Objective

We now formulate the main problem of the paper: choose the regime-dependent weights $u = \{u^{(n)}\}_{n=1}^K \in \Omega$ that maximize the probability of exceeding a target return rate r over horizon T , where the admissible set $\Omega = \Omega_1 \times \dots \times \Omega_K$ encodes per-regime linear or box constraints (long-only, tracking, turnover, ...). Because the policy is a deterministic map of the current regime and z_t (started from π_∞) is strictly stationary, the weight process $u(z_t)$ and the induced return process are *strictly stationary by construction* — the structural property from which backtestability under ergodicity later follows (Section 5). The independence of z_t from the Brownian motion W_t makes the discrete-state case both analytically and numerically more tractable than the continuous-driver setup of [Alqubaisi et al., 2026]: it lets us integrate out the Brownian motion in closed form and reduces the remaining path integral to a $(K-1)$ -dimensional integral over regime occupation times.

3.1 Problem Statement and Conditional Gaussianity

The goal-based utility is

$$(8) \quad U_T[u] = \mathbb{P}_\infty[X_T > rT] = \mathbb{E}_\infty[\mathbf{1}_{X_T > rT}],$$

and the primary problem is $\max_{u \in \Omega} U_T[u]$.

The subscript ∞ on $\mathbb{E}_\infty[\cdot]$ signals the *hot-start* averaging convention: the regime process is initialized from its stationary law, $z_0 \sim \pi_\infty$, rather than from a fixed present state. Under $z_0 \sim \pi_\infty$ the CTMC z_t is strictly stationary and the state-only policy $u(z)$ induces a strictly stationary return process. By ergodicity of z_t , this model-based expectation coincides with a long time-series average of the running goal indicator: for a single realized joint path of (z, W) ,

$$(9) \quad U_T[u] = \mathbb{E}_\infty[\mathbf{1}_{X_T > rT}] = \lim_{T' \rightarrow \infty} \frac{1}{T'} \int_0^{T'} \mathbf{1}_{X_T(\tau, u) > rT} d\tau,$$

where $X_T(\tau, u) \equiv \int_{\tau}^{\tau+T} u(z_s)^{\top} dY_s$ is the running T -window portfolio return starting at τ (in particular $X_T(0, u) = X_T$). Operationally, this identity is what makes the framework natively backtestable: a single fitted policy $u(z)$ rolled forward on the historical time series and averaged over the running T -windows is a direct empirical estimator of $U_T[u]$, with no starting-date re-optimization required. Section 5 exploits exactly this equivalence.

The key structural fact exploited throughout is that, once the regime trajectory is fixed, portfolio dynamics reduce to a Gaussian integrator with piecewise-constant coefficients. Because z_t is independent of W_t , conditioning on the whole path $\mathcal{F}_T^z = \sigma\{z_s : 0 \leq s \leq T\}$ leaves (6) as a linear functional of W , whence

$$(10) \quad X_T \mid \mathcal{F}_T^z \sim \mathcal{N}(\bar{m}_T[u, z], \bar{v}_T[u, z]),$$

with

$$\bar{m}_T[u, z] = \int_0^T u(z_s)^{\top} \mu(z_s) ds, \quad \bar{v}_T[u, z] = \int_0^T u(z_s)^{\top} C(z_s) u(z_s) ds.$$

Integrating out the Brownian motion by the tower property turns the objective into a single expectation of a normal CDF,

$$(11) \quad U_T[u] = \mathbb{E}_{\infty}[\mathcal{N}(\mathcal{S}_T[u])], \quad \mathcal{S}_T[u] = \frac{\bar{m}_T[u, z] - rT}{\sqrt{\bar{v}_T[u, z]}},$$

where $\mathcal{N}$ is the standard normal cdf and $\mathcal{S}_T[u]$ is the path-conditional Sharpe ratio. The outer expectation runs only over the regime process, so the infinite-dimensional Brownian integral has been eliminated exactly. The log-return objective is treated in the next subsection alongside the occupation-time reduction.

3.2 Occupation-Time Reduction

Because the weights $u(z_s) = u^{(z_s)}$ are piecewise constant on the finite state space $\{1, \dots, K\}$, both $\bar{m}_T$ and $\bar{v}_T$ are additive functionals of the regime path that depend on it only through the *regime local times* (occupation times)

$$(12) \quad \mathcal{T}_n \equiv \int_0^T \mathbb{1}_{z_s=n} ds, \quad \mathcal{T}_n \geq 0, \quad \sum_{n=1}^K \mathcal{T}_n = T.$$

Formally, $\bar{m}_T[u, z] = \sum_n \mathcal{T}_n u^{(n)\top} \mu^{(n)}$ and $\bar{v}_T[u, z] = \sum_n \mathcal{T}_n u^{(n)\top} C^{(n)} u^{(n)}$, so $\mathcal{S}_T$ is $\mathcal{F}_T^z$ -measurable through $\{\mathcal{T}_n\}$ alone:

$$(13) \quad \mathcal{S}_T[u] = \frac{\sum_n \mathcal{T}_n u^{(n)\top} \mu^{(n)} - rT}{\sqrt{\sum_n \mathcal{T}_n u^{(n)\top} C^{(n)} u^{(n)}}}.$$

The vector $\mathcal{T} = (\mathcal{T}_1, \dots, \mathcal{T}_K)$ takes values in the scaled simplex $T \cdot \Delta^{K-1}$, so (11) reduces an integral over the infinite-dimensional path space to a $(K-1)$ -dimensional integral against the law ν_T of the local-time vector:

$$(14) \quad U_T[u] = \int_{T \cdot \Delta^{K-1}} \mathcal{N}(\mathcal{S}_T[u; \tau]) \nu_T(d\tau).$$

This is the discrete-state analogue of a Kac/Feynman–Kac occupation-time formula and drives every downstream computational advantage: (i) the outer randomness has dimension $K-1$ rather than ∞ ; (ii) the dependence on the optimization variables u is fully explicit inside the CDF; and (iii) ν_T admits a closed-form Laplace transform via matrix exponentials of Λ (Section 4.5.2).

Log-return variant. The same argument applies verbatim to the log-return process (7): conditioning on $\mathcal{F}_T^z$ leaves R_T Gaussian with the *same* conditional variance $\bar{v}_T[u, z]$ but with drift shifted by the Itô correction, so

$$(15) \quad U_T^{(L)}[u] = \mathbb{P}_\infty[R_T > rT] = \mathbb{E}_\infty \left[\mathcal{N}(\mathcal{S}_T^{(L)}[u]) \right],$$

where

$$(16) \quad \mathcal{S}_T^{(L)}[u] = \frac{\sum_n \mathcal{T}_n \left(u^{(n)\top} \mu^{(n)} - \frac{1}{2} u^{(n)\top} C^{(n)} u^{(n)} \right) - rT}{\sqrt{\sum_n \mathcal{T}_n u^{(n)\top} C^{(n)} u^{(n)}}}.$$

Only the numerator changes; the entire pipeline of Section 4 — two-level decomposition, Monte Carlo, and Bromwich inversion — applies unchanged.

3.3 Gaussian Limit ($K = 1$)

When the regime process degenerates to a single state ($K = 1$), the local-time constraint $\sum_n \mathcal{T}_n = T$ forces $\mathcal{T}_1 = T$ deterministically, and the outer expectation in (11) disappears. The path-conditional Sharpe becomes deterministic,

$$(17) \quad \mathcal{S}_T[u] = \sqrt{T} S(u), \quad S(u) = \frac{u^\top \mu - r}{\sqrt{u^\top C u}},$$

and the goal-based utility collapses to the closed-form Gaussian expression

$$(18) \quad U_T[u] = \mathcal{N}(\sqrt{T} S(u)).$$

Because $\mathcal{N}$ is strictly monotone, maximizing U_T is equivalent to maximizing the excess Sharpe ratio $S(u)$ subject to $\mathbb{1}^\top u = 1$ and $u \in \Omega$. The associated single-regime algorithm is developed in Section 4.3 as a corollary of the general procedure of Section 4.

4 Solving the Optimization Problem

Three ingredients turn (11) into a practical algorithm. Section 4.1 reduces the search from $u \in \mathbb{R}^{NK}$ to a K -dimensional non-convex outer problem coupled with a per-regime inner max-return-at-fixed-variance sub-problem: exactly, in its hard-equality form (non-convex, requiring DCCP); fast, in its soft-inequality relaxation (a convex QP with a global optimum). Section 4.2 illustrates what the inner sub-problem produces on the calibrated data and installs the soft/hard-frontier vocabulary. Sections 4.3 and 4.4 give the target-rate criterion under which the two inner forms yield the same outer optimum, so that the fast soft-form path can safely replace the slow hard-form DCCP — and collect the resulting two-level procedure as Algorithm 2. Section 4.5 finally evaluates the outer objective by Monte Carlo over regime local times, or — orders of magnitude faster — by a semi-analytical Bromwich inversion of the moment generating function.

Schematically, the pipeline is

$$(19) \quad \begin{array}{c} v \in Q^{(v)} \\ \downarrow \text{inner: soft QP / hard DCCP} \\ \bar{\rho}(v) \text{ or } \rho(v) \\ \downarrow \text{utility: MC / Bromwich} \\ U_T(v) \\ \downarrow \text{outer: GSA} \\ v^* \end{array}$$

Reading the diagram. The first arrow is the per-regime inner solve at fixed variance vector v : the fast convex QP (22) when the sign-analysis criterion of Sections 4.3–4.4 passes, otherwise the hard-equality DCCP (23) as a per-regime fallback. The second arrow evaluates the outer utility at the chosen v , either by Monte Carlo over regime local times (Section 4.5.1) or by the semi-analytical Bromwich inversion of the moment generating function (Section 4.5.2). The third arrow closes the loop: generalized simulated annealing (GSA) proposes new v ’s until the outer utility converges to v^* . The whole procedure is collected as Algorithm 2.

4.1 Two-Level Optimization Scheme

Direct maximization of (11) over the full weight vector $u \in \mathbb{R}^{NK}$ is uninviting: the utility is non-convex in u , its dimension scales with NK , and box or tracking constraints preclude closed-form updates. Formulation (13) nonetheless admits a clean two-level decomposition.

The key observation is that, per regime, only two scalars enter the path-conditional Sharpe: the expected return $m_n = u^{(n)\top} \mu^{(n)}$ and the variance $v_n = u^{(n)\top} C^{(n)} u^{(n)}$. Introducing the normalised occupation fractions

$$(20) \quad \tau_n \equiv \frac{\mathcal{T}_n}{T} \in [0, 1], \quad \sum_{n=1}^K \tau_n = 1,$$

the utility factors through the pair (m, v) as

$$(21) \quad U_T(m, v) = \mathbb{E}_\infty \left[\mathcal{N} \left(\sqrt{T} \frac{\sum_n \tau_n m_n - r}{\sqrt{\sum_n \tau_n v_n}} \right) \right], \quad U_T^{(L)}(m, v) = \mathbb{E}_\infty \left[\mathcal{N} \left(\sqrt{T} \frac{\sum_n \tau_n (m_n - \frac{1}{2} v_n) - r}{\sqrt{\sum_n \tau_n v_n}} \right) \right],$$

with (m, v) constrained to the admissible surface $S^{(m, v)}$ carved out by the weights and Ω_n .

Because U_T is monotone-increasing in each m_n at fixed v , we can eliminate the mean regime-by-regime by pushing m_n to its efficient-frontier value at fixed variance v_n . Two per-regime formulations of this inner sub-problem are available:

$$(22) \quad \bar{\rho}_n(v_n) \equiv \max_{u^{(n)}} u^{(n)\top} \mu^{(n)} \quad \text{s.t.} \quad u^{(n)\top} C^{(n)} u^{(n)} \leq v_n, \quad \mathbb{1}^\top u^{(n)} = 1, \quad u^{(n)} \in \Omega_n,$$

$$(23) \quad \rho_n(v_n) \equiv \max_{u^{(n)}} u^{(n)\top} \mu^{(n)} \quad \text{s.t.} \quad u^{(n)\top} C^{(n)} u^{(n)} = v_n, \quad \mathbb{1}^\top u^{(n)} = 1, \quad u^{(n)} \in \Omega_n,$$

where Ω_n collects any additional linear or box constraints. The *hard-risk* form (23) pins the variance identically to v_n ; substituting its optimum $\rho(v)$ into the outer utility yields the *exact* reduction of the original problem to a K -dimensional non-convex search over the variance vector,

$$(24) \quad \max_{u \in \Omega} U_T[u] = \max_{v \in Q^{(v)}} U_T(\rho(v), v),$$

where $Q^{(v)}$ is the hyper-box of admissible variances (per-regime bounds are the min and max of $u^{(n)\top} C^{(n)} u^{(n)}$ over Ω_n). Equation (24) is a formal identity — no relaxation, no approximation — and in this sense (23) is the exact primitive of the two-level reduction. The catch is that the equality $u^{(n)\top} C^{(n)} u^{(n)} = v_n$ is a difference-of-convex constraint, so (23) is non-convex. It is handled by disciplined convex-concave programming (DCCP), which linearises the concave side of the constraint at the current iterate and thereby reduces *each iteration* to a convex QP — structurally the same object as (22) and solved by the same interior-point machinery in microseconds. A single DCCP call is therefore fast in absolute terms; the price against (22) is only a small constant number of such QP iterations per call, together with the loss of the global-optimality guarantee (DCCP converges to a local optimum from its initial point).

A faster, globally-convergent alternative is the *soft-risk* relaxation (22): replacing the equality by the inequality $u^{(n)\top} C^{(n)} u^{(n)} \leq v_n$ turns the inner problem into a plain convex QP (quadratic objective, one quadratic inequality, one linear equality, convex Ω_n) whose global optimum interior-point solvers dispatch in microseconds. Its optimum $\bar{\rho}_n(v_n) \geq \rho_n(v_n)$ is the concave envelope of the hard frontier. Substituting $\bar{\rho}$ into the outer objective gives the fast surrogate

$$(25) \quad \max_{v \in Q^{(v)}} U_T(\bar{\rho}(v), v),$$

which coincides with the exact (24) whenever the soft and hard frontiers agree at the outer optimum. Whether this holds — and hence whether the fast QP can safely replace the slow DCCP — is a target-rate-dependent question. Sections 4.3 and 4.4 give a sign-analysis criterion in r (and in the tightness of the per-regime box constraints) that flags exactly when the soft form is guaranteed to reproduce the hard-form outer optimum. On the calibrated data of Section 5 every (T, r) cell satisfies this criterion, so the fast path (25) is used throughout the empirical results; only rare aspirational cells above the threshold call for the DCCP fallback of (24), applied per-regime.

The outer objective in either form has only K variables — $K = 3$ in Section 5 — but is non-convex. We use *generalized simulated annealing* [Tsallis and Stariolo, 1996, Xiang et al., 1997], which combines heavy-tailed Tsallis proposals with a local-search polishing step and converges to the global optimum in the limit of infinite cooling.

Swapping the roles — fix m_n , minimize v_n , search over m — is not equivalent: $U_T(m, v)$ is not always decreasing in v , so the outer problem loses its clean equivalence to the joint optimum.

4.2 Inner Sub-Problem: Efficient-Frontier Diagnostics

To illustrate what the inner sub-problem produces on real data (calibration details in Section 5), Figure 1 shows one panel per regime — expansion (top), neutral (middle), crisis (bottom) — each plotting two curves in the (variance, expected return) plane:

- the *soft-risk* frontier $v_n \mapsto \bar{\rho}_n(v_n)$ (orange) — the concave envelope of the inequality-constrained inner problem (22) (variance $\leq v_n$); this is the object the outer search actually consumes;
- the *hard-risk* frontier $v_n \mapsto \rho_n(v_n)$ (blue) — the equality-constrained inner problem (23) obtained by disciplined convex-concave programming (DCCP), shown for reference.

The two coincide where the inequality binds at the optimum — expansion and neutral regimes throughout, crisis regime up to the interior peak; past the peak, ρ_n slides onto the lower branch of the mean-variance parabola while $\bar{\rho}_n$ stays on the upper. Vertical dotted guides bracket the admissible variance interval $[v_n^{\min}, v_n^{\max}]$ ($v_n^{\min}$ from a convex QP; $v_n^{\max}$ from a DCCP whose optimum sits at a corner of Ω_n); horizontal dashed guides bracket the expected-return envelope $[m_n^{\min}, m_n^{\max}]$ (trivial LPs on Ω_n). The product $Q^{(v)} = \prod_n [v_n^{\min}, v_n^{\max}]$ is the hyper-box over which the outer generalized-simulated-annealing loop searches.

Denote by $\bar{m}_n^{\max} \equiv \max_{v \in [v_n^{\min}, v_n^{\max}]} \bar{\rho}_n(v)$ the highest drift attainable on the soft envelope over the admissible variance range. Below this level — at variances where the soft envelope is still climbing — the two frontiers *coincide*, $\bar{\rho}_n(v) = \rho_n(v)$. Once the soft envelope reaches $\bar{m}_n^{\max}$, it extrapolates flatly at this value; the hard equality frontier can then diverge, staying constant or decreasing.

Across the three panels one sees three typical topologies. In *expansion* (top) the frontier rises monotonically to $v_n^{\max}$: $\bar{m}_n^{\max}$ is attained at the right edge and coincides with the horizontal ceiling $m_n^{\max}$. In *neutral* (middle) $\bar{m}_n^{\max}$ is again attained at $v_n^{\max}$ but sits strictly below the ceiling. In *crisis* (bottom) $\bar{m}_n^{\max}$ is reached at an *interior* variance $\bar{v}_n^{\max}$; past this peak the soft envelope stays flat at $\bar{m}_n^{\max}$, while the hard equality frontier decreases (blue below orange). Strictly speaking, past $\bar{v}_n^{\max}$ the plot is no longer the efficient frontier: a rational allocation does not take on extra variance once $\bar{m}_n^{\max}$ is attained.




Figure 1: Per-regime efficient frontier of the inner sub-problem on the calibrated sector universe under long-only / full-load constraints. Orange: soft-risk envelope $\bar{\rho}_n(v_n)$ (inequality, (22)). Blue: hard-risk frontier $\rho_n(v_n)$ (equality, (23)) obtained by DCCP. Vertical dotted guides mark the admissible variance interval $[v_n^{\min}, v_n^{\max}]$; horizontal dashed guides mark the expected-return envelope $[m_n^{\min}, m_n^{\max}]$.

4.3 Soft vs. hard forms: single-regime analysis

Each panel of Figure 1 is structurally a $K = 1$ instance: the inner sub-problem of a single regime, with the outer local-time law ν_T collapsed to $\mathcal{T}_1 = T$ deterministically. We use it to determine *when* the soft form (22) suffices in place of the hard equality.

In the Gaussian limit of Section 3.3 the goal-based utility factors as

$$(26) \quad U_T[u] = \mathcal{N}\left(\sqrt{T} \frac{u^\top \mu - r}{\sqrt{u^\top C u}}\right),$$

so maximising U_T reduces to maximising the excess Sharpe ratio, in either the simple or the log-return convention:

$$(27) \quad S(v) = \frac{\rho(v) - r}{\sqrt{v}}, \quad S_L(v) = \frac{\rho(v) - \frac{1}{2}v - r}{\sqrt{v}}.$$

Whether the soft form suffices comes down to *where* $S(v)$ (or $S_L(v)$) is maximised. Recall from Section 4.2 that for $v \leq \bar{v}^{\max}$ the two frontiers coincide, $\rho(v) = \bar{\rho}(v)$, while past $\bar{v}^{\max}$ the soft envelope plateaus at $\bar{m}^{\max}$ and the hard frontier stays the same or decreases. If the Sharpe maximum lies to the left of $\bar{v}^{\max}$ the two forms give the same outer optimum; if it lies to the right, they can differ.

Differentiating,

$$(28) \quad \frac{dS}{dv} = \frac{\rho'(v) - \frac{1}{2}\frac{\rho(v)-r}{v}}{\sqrt{v}}, \quad \frac{dS_L}{dv} = \frac{\rho'(v) - \frac{1}{2}\frac{\rho(v)-r}{v} - \frac{1}{4}}{\sqrt{v}},$$

and evaluating at the soft peak $v = \bar{v}^{\max}$ (where $\rho'(\bar{v}^{\max}) = 0$),

$$(29) \quad \left. \frac{dS}{dv} \right|_{\bar{v}^{\max}} = -\frac{\bar{m}^{\max} - r}{2(\bar{v}^{\max})^{3/2}}, \quad \left. \frac{dS_L}{dv} \right|_{\bar{v}^{\max}} = -\frac{\bar{m}^{\max} - r}{2(\bar{v}^{\max})^{3/2}} - \frac{1}{4\sqrt{\bar{v}^{\max}}}.$$

The sign splits the analysis in two.

Common case ($r < \bar{m}^{\max}$). Both derivatives at $\bar{v}^{\max}$ are strictly negative (log case automatically, from the extra $-\frac{1}{4\sqrt{\bar{v}^{\max}}}$), so the Sharpe maximum lies strictly to the left of $\bar{v}^{\max}$; and to the right, on the plateau $v \in [\bar{v}^{\max}, v^{\max}]$, the soft Sharpe $(\bar{m}^{\max} - r)/\sqrt{v}$ is strictly decreasing in v while the hard Sharpe is pointwise smaller since $\rho(v) \leq \bar{m}^{\max}$. Both maxima therefore sit inside $[v^{\min}, \bar{v}^{\max}]$, where $\rho \equiv \bar{\rho}$: the soft form gives the same outer optimum as the hard one, and a plain convex QP suffices. On the calibrated data this covers the *expansion* and *neutral* panels of Figure 1, and every (T, r) cell in Section 5.

Aspirational targets ($r > \bar{m}^{\max}$). The derivatives at $\bar{v}^{\max}$ can flip sign (in particular $dS/dv|_{\bar{v}^{\max}} > 0$), so the Sharpe maximum can migrate past $\bar{v}^{\max}$ into the range where $\rho(v) < \bar{\rho}(v) = \bar{m}^{\max}$. The two outer optima can then differ, and DCCP at the inner level may become necessary. Such (T, r) cells are rare in institutional practice — a legitimate “reach for the moon” request — and the hard-soft gap they produce is genuine.

Concrete illustration: tightening the long-only cap. Without weight bounds ($u_i \in \mathbb{R}$) the frontier reduces to the standard Markowitz efficient frontier: $\rho(v)$ is a monotonically increasing $\sqrt{v}$ -type curve with $\bar{\rho} = \rho$ everywhere. The interior peak in the crisis panel of Figure 1 is entirely a consequence of the long-only box $u_i \in [0, 1]$: past the peak, the equality $u^\top C u = v$ can only be met by tilting into lower-drift corner sectors, and ρ falls. Tightening the box to $u_i \in [0, 0.75]$ sharpens the effect and pushes the setup into the aspirational regime (Figures 2 and 3): $S_L(v)$ is maximised strictly past $\bar{v}^{\max}$, where $\rho < \bar{\rho}$, and the hard-form outer optimum beats the soft-form one — confirming that the $r < \bar{m}^{\max}$ restriction of the common case is not vacuous.




Figure 2: Crisis-regime efficient frontier under a tighter per-sector cap $u_i \in [0, 0.75]$. Orange: soft envelope $\bar{\rho}(v)$; blue: hard-equality $\rho(v)$; guides as in Figure 1. The interior peak $\bar{v}^{\max}$ is sharper than under the $[0, 1]$ cap.




Figure 3: Excess Sharpe $S_L(v)$ under the $[0, 0.75]$ cap of Figure 2. The maximum sits past $\bar{v}^{\max}$, where $\rho(v) < \bar{\rho}(v)$: the hard-form outer optimum beats the soft-form one.

We accordingly adopt the soft form as the default inner solve (one convex QP per outer v -evaluation) whenever $r < \bar{m}^{\max}$, and fall back to hard-form DCCP only for the rare aspirational cells with $r \geq \bar{m}^{\max}$. Every calibrated cell in Section 5 sits in the common regime, so the fast path is used throughout the empirical results.

The resulting procedure is Algorithm 2 unfolded at $K = 1$.

Algorithm 1 Two-level scheme, Gaussian limit ($K = 1$).

1. Compute the admissible variance range $v^{\min} = \min_{u \in \Omega, \mathbb{1}^\top u = 1} u^\top C u$ (a convex QP) and $v^{\max} = \max_{u \in \Omega, \mathbb{1}^\top u = 1} u^\top C u$ (a DCCP).
2. **Target barrier (QP $\leftrightarrow$ DCCP switch).** Compute the soft-frontier peak $\bar{m}^{\max} = \max_{v \in [v^{\min}, v^{\max}]} \bar{\rho}(v)$ by solving (22) on a coarse variance grid. If $r < \bar{m}^{\max}$, use the soft-form inner problem below (fast convex QP path); otherwise, replace $\bar{\rho}(v)$ by the hard-form $\rho(v)$ solved by DCCP for $v \in [\bar{v}^{\max}, v^{\max}]$ (slow aspirational path).
3. For $v \in [v^{\min}, v^{\max}]$, evaluate the soft-risk inner problem

$$\bar{\rho}(v) = \max_u u^\top \mu \quad \text{s.t.} \quad u^\top C u \leq v, \quad \mathbb{1}^\top u = 1, \quad u \in \Omega,$$

as a convex QP.

4. Form the scalar excess Sharpe $S(v) = (\bar{\rho}(v) - r)/\sqrt{v}$ and solve $v^* = \arg \max_{v \in [v^{\min}, v^{\max}]} S(v)$ by a one-dimensional global search (e.g. Brent with restarts, or grid + local polish).
 5. Return $u^* = \text{optimizer of the inner QP at } v = v^*$; the corresponding utility is $U_T[u^*] = \mathcal{N}(\sqrt{T} S(v^*))$.
-

The Bromwich contour of Section 4.5.2 likewise degenerates in this limit and recovers (26) in closed form.

4.4 The General K Case: Sign Analysis of the Soft Form

The $K = 1$ analysis lifts to $K > 1$ by differentiating the outer utility with respect to each per-regime variance v_k . Introduce the τ -weighted per-path drift and variance $\bar{M}(\tau, v) \equiv \sum_n \tau_n \bar{\rho}_n(v_n)$, $\bar{V}(\tau, v) \equiv \sum_n \tau_n v_n$, and the log-return analogue $\bar{M}^L(\tau, v) \equiv \bar{M}(\tau, v) - \frac{1}{2} \bar{V}(\tau, v)$. A direct differentiation (Appendix B) shows that, evaluated at the per-regime soft peak $v_k = \bar{v}_k^{\max}$ where $\bar{\rho}'_k(\bar{v}_k^{\max}) = 0$, the outer utility partial reads

$$(30) \quad \left. \frac{\partial U_T(\bar{\rho}(v), v)}{\partial v_k} \right|_{v_k = \bar{v}_k^{\max}} = -\frac{\sqrt{T}}{2} \mathbb{E}_\infty \left[\phi(\bar{S}_T) \frac{\tau_k (\bar{M} - r)}{\bar{V}^{3/2}} \right],$$

with ϕ the standard normal density and the log analogue carrying $\bar{M}^L - r + \bar{V}$ in place of $\bar{M} - r$. The sign of (30), set by $\phi \geq 0$, $\tau_k \geq 0$, and the τ -mean of $\bar{M} - r$ (or $\bar{M}^L + \bar{V} - r$ for logs), splits the analysis into two regimes.

Common case (r below the π_∞ -weighted drift, heuristically). A strict pointwise sufficient condition for the soft form to work is $r < \min_n \bar{m}_n^{\max}$: it guarantees $\bar{\rho}_n(v_n) > r$ for every n , hence $\bar{M}(\tau, v) > r$ for every τ , and hence (30) strictly negative for every k . This bound is usually too conservative — a single low-peak regime (a deeply negative crisis regime, say) drags $\min_n \bar{m}_n^{\max}$ down even when that regime carries little stationary weight $\pi_\infty^{(n)}$. In practice we use either

- the π_∞ -weighted heuristic threshold, $r < \mathbb{E}_\infty [\bar{M}(\tau, v)] \approx \sum_n \pi_\infty^{(n)} \bar{\rho}_n(v_n)$; or, more sharply,
- a direct Monte Carlo estimate of (30) at $v_k = \bar{v}_k^{\max}$ for every k (with τ -samples drawn from ν_T), retaining the r -range on which every regime's derivative comes out negative.

On the retained range the outer optimum v^* satisfies $v_k^* < \bar{v}_k^{\max}$ regime by regime and the soft form suffices; the log case is even more easily negative thanks to the extra $+\bar{V}$ shift in the τ -mean (see Appendix B).

Aspirational targets (r above the π_∞ -weighted drift). When r exceeds the heuristic threshold — or, more precisely, when the MC estimate of $\partial U_T / \partial v_k$ at $v_k = \bar{v}_k^{\max}$ comes out non-negative for some regime k — $\bar{M} - r$ becomes negative on the τ -realisations concentrated on regime k , and the outer optimum pushes

v_k^* past $\bar{v}_k^{\max}$ into the range where $\rho_k < \bar{\rho}_k = \bar{m}_k^{\max}$. The hard-form DCCP inner solve is then required in regime k ; in every other regime the soft form still suffices, so the fallback is per-regime.

On the calibrated data of Section 5 every (T, r) cell we consider satisfies the heuristic: the outer GSA loop lands on the orange soft envelope in every panel of Figure 1, and the per-regime hard-form DCCP fallback is not exercised.

We can now collect the resulting two-level procedure. Algorithm 2 states the common case in which the sign-analysis criterion holds so the inner primitive is the fast soft QP (22); the DCCP fallback (23) replaces it in any regime for which the criterion fails.

Algorithm 2 Two-level scheme, general regime-switching case.

1. Compute the admissible variance hyper-box $Q^{(v)} = \prod_{n=1}^K [v_n^{\min}, v_n^{\max}]$: for each regime n , $v_n^{\min} = \min u^{(n)\top} C^{(n)} u^{(n)}$ over $\{u^{(n)} \in \Omega_n : \mathbb{1}^\top u^{(n)} = 1\}$ is a convex QP, and $v_n^{\max}$ over the same feasible set is a DCCP whose optimum sits at a corner of Ω_n .

2. **Target barrier (per-regime QP $\leftrightarrow$ DCCP switch).** Compute the per-regime soft-frontier peaks $\bar{m}_n^{\max} = \max_{v_n \in [v_n^{\min}, v_n^{\max}]} \bar{\rho}_n(v_n)$ and mark each regime n as *soft-safe* if the target lies below its per-regime barrier. Two threshold levels are available (cf. Section 4.4):

- (a) *strict pointwise*: $r < \min_n \bar{m}_n^{\max}$ marks every regime soft-safe;
- (b) *sharp per-regime*: check the sign of the MC estimate of $\partial U_T / \partial v_k \big|_{v_k = \bar{v}_k^{\max}}$ from (30); regime k is soft-safe when this comes out strictly negative.

Set the inner primitive for each regime to the soft QP (22) on the soft-safe regimes and to the hard-form DCCP (23) on the remaining (aspirational) regimes.

3. For each candidate variance vector $v = (v_1, \dots, v_K) \in Q^{(v)}$, solve the K per-regime inner problems from Step 2 to obtain the frontier drifts $\bar{\rho}_n(v_n)$ (or $\rho_n(v_n)$ for any aspirational regime).
 4. Evaluate the outer utility $U_T(\bar{\rho}(v), v)$ from (21) — by Monte Carlo over regime local times (Section 4.5.1) or by the Bromwich formula (Section 4.5.2).
 5. Solve $v^* = \arg \max_{v \in Q^{(v)}} U_T(\bar{\rho}(v), v)$ on the K -dimensional non-convex outer objective by generalized simulated annealing [Tsallis and Stariolo, 1996, Xiang et al., 1997].
 6. Return $u^* = \{u^{(n)*}\}_{n=1}^K$, the K inner-QP optimizers evaluated at $v = v^*$; the corresponding utility is $U_T[u^*] = U_T(\bar{\rho}(v^*), v^*)$.
-

4.5 Numerical Evaluation of the Utility

With the outer v fixed and the inner QPs (22) providing $\bar{\rho}(v)$, the remaining task is to evaluate $U_T(\bar{\rho}(v), v)$. We give two routes: unbiased Monte Carlo over regime local times (Section 4.5.1) and a semi-analytical Bromwich inversion of a matrix-exponential moment generating function (Section 4.5.2) — orders of magnitude faster, and used by default in the numerical experiments.

4.5.1 Monte Carlo Over Regime Local Times

Simulating the local-time vector $\mathcal{T}$ requires only exact-jump simulation of the CTMC z_t with generator Λ starting from π_∞ : draw exponential holding times in each state, accumulate them into the corresponding $\mathcal{T}_n$, and stop at horizon T . No discretization of W is needed, and no time step enters. Averaging $\mathcal{N}(\mathcal{S}_T[u; \mathcal{T}])$ over L such draws yields an unbiased estimator of $U_T[u]$ with $\mathcal{O}(L^{-1/2})$ error, tight enough for practice at $L \sim 10^4$ – 10^5 .

4.5.2 Semi-Analytical Evaluation via Moment Generating Function

For faster evaluation we bypass Monte Carlo entirely by giving a closed-form expression for the moment generating function of X_T . Integrating out the Brownian motion in $q_T(\xi|k) \equiv \mathbb{E} [e^{\xi X_T} | z_0 = k]$ gives

$$(31) \quad q_T(\xi|k) = \mathbb{E} \left[\exp \left(\xi \int_0^T u^\top \mu(z_t) dt + \frac{\xi^2}{2} \int_0^T u^\top C(z_t) u dt \right) \middle| z_0 = k \right].$$

The remaining expectation is over a functional of regime local times and can be evaluated in closed form on the finite state space [Çınlar, 1972]:

$$(32) \quad q_T(\xi|k) = \left\{ e^{T \left(\Lambda + \xi D_m + \frac{\xi^2}{2} D_v \right)} \mathbb{1} \right\}_k, \quad D_m = \text{diag} \left(u^{(n)\top} \mu^{(n)} \right), \quad D_v = \text{diag} \left(u^{(n)\top} C^{(n)} u^{(n)} \right).$$

A derivation sketch is given in the Appendix. Averaging over the stationary distribution π_∞ ,

$$(33) \quad q_T(\xi) = \pi_\infty^\top e^{T \left(\Lambda + \xi D_m + \frac{\xi^2}{2} D_v \right)} \mathbb{1}.$$

Combining with the Bromwich representation $\mathbb{1}_{x>0} = \frac{1}{2\pi i} \int_{C_+} \frac{d\xi}{\xi} e^{\xi x}$ (a vertical contour to the right of zero) yields the success probability as a single one-dimensional inverse Laplace integral,

$$(34) \quad U_T[u] = \frac{1}{2\pi i} \int_{C_+} \frac{d\xi}{\xi} q_T(\xi) e^{-\xi r T},$$

which is orders of magnitude faster to evaluate than the Monte-Carlo form (11), at the cost of a matrix exponential per node on the contour.

5 Numerical Experiments

We validate the framework on $T_{\text{tot}} = 25$ years of daily total returns for the eleven GICS sectors (SPDR sector ETF universe), together with the CBOE VIX index over the same window. The experiments are a proof of concept in the long-only class: they isolate the value of regime conditioning under the stationary framework and quantify its statistical significance, but do not attempt to price transaction costs or to promote a specific weight vector for live use. Institutional deployment would run through the tracking-and-turnover-bounded *Hedged* variant defined below, whose numbers appear alongside those of the raw *Optimal* throughout.

5.1 Data, regimes, and strategies

Regime identification. We take $K = 3$ regimes labelled *expansion*, *neutral*, and *crisis*, and identify the current label from the VIX with the standard practitioner thresholds $[0, 0.20, 0.30]$ (VIX expressed as an annualized decimal, so that 0.20 and 0.30 correspond to the quoted VIX values of 20 and 30): $z_t = 0$ when $\text{VIX}_t < 0.20$, $z_t = 1$ for $\text{VIX}_t \in [0.20, 0.30]$, and $z_t = 2$ for $\text{VIX}_t \geq 0.30$. Empirical occupation frequencies over the sample are approximately 60%, 25%, and 15%, consistent with the intuition that expansion is modal while crisis is a fat but rare tail. Modest changes to the thresholds shift the local times $\mathcal{T}_n$ but leave the ranking of state-conditional Sharpe ratios intact, so the qualitative backtest conclusions are robust to the specific partition.

Strategies. Four portfolios enter the comparison:

- **Optimal** — regime-dependent weights $u_{T,r}^{(\text{opt})}(z)$ maximizing the goal-reaching probability under normalization and long-only constraints.
- **Hedged** — $u_{T,r}^{(\text{opt})}$ projected onto a benchmark-tracking and turnover-constrained feasible set (defined below).

- **Gaussian** — goal-based optimum under a single-regime Gaussian model (regime label ignored; drift and covariance averaged over the whole window), isolating the value added by regime conditioning.
- **S&P 500** — market-capitalization-weighted sector index, as passive baseline.

The Hedged variant $u_{T,r}^{(\text{hedge})}(t)$ maximises overlap with the Optimal under long-only, full-load, tracking, and turnover constraints:

$$\begin{aligned}
& u_{T,r}^{(\text{hedge})}(t) \cdot u_{T,r}^{(\text{opt})}(t) \rightarrow \max, \\
& 0 \leq u_{T,r}^{(\text{hedge})}(t)_i \leq 1, \quad i = 1, \dots, N, \quad \sum_i u_{T,r}^{(\text{hedge})}(t)_i = 1, \\
& |u_{T,r}^{(\text{hedge})}(t) - u^{(\text{SP})}(t)| \leq q_{\text{tilt}}, \\
& |u_{T,r}^{(\text{hedge})}(t) - u_{T,r}^{(\text{hedge})}(t-1)| \leq q_{\text{to}}.
\end{aligned}$$

We take $q_{\text{tilt}} = 40\%$ and $q_{\text{to}} = 20\%$, values representative of the tracking-error budgets and rebalancing intensities used by sovereign wealth and pension funds. The problem is a standard convex QP; tighter q_{tilt} pulls performance towards the S&P 500, looser bounds recover the unconstrained optimum.

Performance metric. The primary metric is the historical rolling-window probability of reaching the target return:

$$(35) \quad P_{T,r}[\Pi] = \frac{1}{N_t} \sum_{t < T_{\text{tot}} - T} \mathbb{1}_{R(t-T,t)[\Pi] > r}, \quad R(t-T,t)[\Pi] = \frac{1}{T} \sum_{t-T < s < t} \frac{\Delta \Pi(s)}{\Pi(s)}.$$

Consistently with the two return conventions of Section 2.3, we also report the log-return version

$$(36) \quad R^{\log}(t-T,t)[\Pi] = \frac{1}{T} \log \frac{\Pi(t)}{\Pi(t-T)} = \frac{1}{T} \sum_{t-T < s < t} \log \left(1 + \frac{\Delta \Pi(s)}{\Pi(s)} \right),$$

the discrete-time counterpart of R_T/T in (7). The identity between the along-path frequency $P_{T,r}$ and the model-implied utility $U_T[u]$ holds because $u(z)$ depends only on the current regime label; this alignment is what makes $P_{T,r}$ a direct empirical estimator of $U_T[u]$ (see Introduction, §1, ¶2).

5.2 Headline result

The Optimal strategy exceeds the passive S&P 500 baseline by 15–40% of goal-reaching probability at horizons $T \geq 3$ years and target rates $r \geq 10\%$, peaking at $\approx 40\%$ for $r = 15\%$, $T \in \{3, 5\}$ years. The Hedged variant captures typically half of that gain under tracking-error and turnover constraints, and the single-regime Gaussian collapse trails the Optimal by 5–12% at the same aggressive cells, isolating regime conditioning as the source of the lift. Figure 4a–4b visualise the picture; Table 1 reports the numerical detail for simple returns, and its log-return companion (Table 3, Appendix C.4) agrees within 0.01–0.02 on Sharpe and $\sim 1\%$ on return/vol.

Target rate (%)	Horizon	Strategy	Prob (%)	Sharpe	Return (%)	Vol (%)	Daily (Max) turnover (%)
5	1M	optimal	63	0.75	18	24	0.8 (200.0)
5	1M	gaussian	62	0.68	17	25	1.3 (200.0)
5	1M	hedged	64	0.75	15	20	4.6 (50.0)
5	1M	sp500	64	0.66	12	19	0.5 (16.4)
5	1Y	optimal	85	0.76	18	24	0.7 (200.0)
5	1Y	gaussian	79	0.68	17	25	1.3 (200.0)
5	1Y	hedged	83	0.75	15	20	4.8 (50.0)
5	1Y	sp500	79	0.66	12	19	0.5 (16.4)
5	3Y	optimal	97	0.76	18	24	0.7 (200.0)
5	3Y	gaussian	97	0.68	17	25	1.3 (200.0)
5	3Y	hedged	92	0.76	15	20	4.8 (50.0)
5	3Y	sp500	87	0.66	12	19	0.5 (16.4)
5	5Y	optimal	100	0.76	18	24	0.7 (200.0)
5	5Y	gaussian	100	0.68	17	25	1.3 (200.0)
5	5Y	hedged	97	0.75	15	20	4.7 (50.0)
5	5Y	sp500	84	0.66	12	19	0.5 (16.4)
10	1M	optimal	60	0.76	18	24	0.8 (200.0)
10	1M	gaussian	59	0.7	18	25	1.1 (200.0)
10	1M	hedged	61	0.75	15	20	4.5 (50.0)
10	1M	sp500	60	0.66	12	19	0.5 (16.4)
10	1Y	optimal	74	0.76	18	24	0.7 (200.0)
10	1Y	gaussian	68	0.7	18	25	1.1 (200.0)
10	1Y	hedged	69	0.75	15	20	4.7 (50.0)
10	1Y	sp500	66	0.66	12	19	0.5 (16.4)
10	3Y	optimal	89	0.76	18	24	0.7 (200.0)
10	3Y	gaussian	82	0.7	18	25	1.1 (200.0)
10	3Y	hedged	84	0.75	15	20	4.7 (50.0)
10	3Y	sp500	72	0.66	12	19	0.5 (16.4)
10	5Y	optimal	91	0.76	18	24	0.7 (200.0)
10	5Y	gaussian	90	0.7	18	25	1.1 (200.0)
10	5Y	hedged	80	0.76	15	20	4.7 (50.0)
10	5Y	sp500	69	0.66	12	19	0.5 (16.4)
15	1M	optimal	57	0.77	19	24	0.7 (200.0)
15	1M	gaussian	56	0.7	18	25	1.1 (200.0)
15	1M	hedged	56	0.75	15	20	4.9 (50.0)
15	1M	sp500	55	0.66	12	19	0.5 (16.4)
15	1Y	optimal	61	0.76	18	24	0.7 (200.0)
15	1Y	gaussian	58	0.7	18	25	1.1 (200.0)
15	1Y	hedged	53	0.77	15	20	4.9 (50.0)
15	1Y	sp500	46	0.66	12	19	0.5 (16.4)
15	3Y	optimal	66	0.76	18	24	0.7 (200.0)
15	3Y	gaussian	54	0.7	18	25	1.1 (200.0)
15	3Y	hedged	47	0.76	15	20	4.8 (50.0)
15	3Y	sp500	27	0.66	12	19	0.5 (16.4)
15	5Y	optimal	64	0.76	18	24	0.7 (200.0)
15	5Y	gaussian	60	0.7	18	25	1.1 (200.0)
15	5Y	hedged	41	0.76	15	20	4.8 (50.0)
15	5Y	sp500	24	0.66	12	19	0.5 (16.4)

Table 1: Historical goal-reaching probability, Sharpe ratio, annualized return and volatility, and daily/max turnover for the four strategies across the target-rate/horizon grid; simple-return convention. Log-return counterpart in Table 3 (Appendix C.4).





Figure 4: Historical goal-reaching probabilities across target rate and investment horizon, visualizing Table 1 (simple returns; log-return counterpart in Table 3, Appendix C.4). *Optimal* is the regime-based, horizon-specific portfolio; *Gaussian* is the single-regime benchmark portfolio; *Hedged* is the turnover- and tilt-constrained portfolio; *S&P 500* is the historical benchmark. Uncertainty bars are $\pm$ two Monte-Carlo standard deviations.

Calibration diagnostics. The per-regime drift and volatility parameters $(\mu^{(n)}, \sigma^{(n)})$ are estimated from sample moments on trailing windows whose length grows with regime rarity: 5, 7, and 12 years for expansion,

neutral, and crisis. This rarity-scaled design lets the crisis calibration accumulate enough tail observations for a stable estimate while keeping the expansion regime responsive to recent conditions. The transition generator Λ is estimated by matching empirical transition frequencies over 12 years. The complete model — $(\mu^{(n)}, \sigma^{(n)})$, Λ , and the resulting regime-conditional Optimal weights $u_{T,r}^{(\text{opt})}(n)$ — is refreshed on a rolling quarterly schedule; between refreshes the policy is $u_{T,r}^{(\text{opt})}(z_t)$ evaluated at the observed VIX regime label. The single-regime Gaussian benchmark (the practitioner *constant-risk* model) is recalibrated monthly on a 3-year trailing window, reflecting its lower parameter count.


Table 2 tracks the annualized drift and volatility of the realized strategy across the 93 quarterly refreshes. Volatility is stable everywhere ($\text{CV}(\sigma) \in [0.20, 0.31]$). Drift is tightest in the modal $K = 3$ neutral cell ($\text{CV}(\mu) \approx 0.27$) and noisier in the rare regimes (0.81/1.21), motivating the rarity-scaled windows; the $K = 1$ Gaussian sits at 0.83, a $\sim 3\times$ penalty relative to $K = 3$ neutral. The VIX regime label acts as an informative side variable in the calibration: conditioning on z_t lets the estimator recover the state-dependent drift and volatility of the underlying data-generating process, whereas the $K = 1$ Gaussian averages across the same latent states and absorbs the between-state dispersion of μ into the sampling noise of a single moment. Drift signs are stable in every $K = 3$ row, and the residual parameter noise is far smaller than the 30–40% probability edge reported in Table 1.

Model	Regime	Drift μ (%)	CV(μ)	Vol σ (%)	CV(σ)
Single-regime ($K = 1$)	—	19.6 (16.2)	0.83	22.9 (7.1)	0.31
Regime-switching ($K = 3$)	Expansion	33.9 (27.6)	0.81	18.7 (5.8)	0.31
	Neutral	14.6 (3.9)	0.27	19.3 (5.4)	0.28
	Crisis	−20.4 (24.6)	1.21	32.0 (6.3)	0.20


Table 2: Empirical stability of the realized-strategy calibration across the 93 quarterly refreshes of the backtest. Annualized drift and volatility (in %) are reported as *mean (robust std)*, where the robust standard deviation is the median absolute deviation rescaled to a Gaussian equivalent. CV denotes the coefficient of variation, robust std divided by |mean|. Strategy moments are essentially horizon-independent across $T \in \{1M, 1Y, 3Y, 5Y\}$ (differences of at most 0.02%), so a single row per (model, regime) is reported.

The two model classes handle non-stationarity by opposite mechanisms. The single-regime Gaussian model has no state variable, so its only way to adapt is to recalibrate frequently on a short trailing window; each calibration is imprecise (few observations per moment), and the policy jumps every time the window rolls forward. The regime-switching model captures state-dependence explicitly through the label z_t and does not need to chase conditions via cadence: per-regime fitting requires longer windows — 12 years for crisis — but each calibration is then a more precise estimator of the conditional distribution, so refreshes can be less frequent (quarterly) and the policy remains stable between them.

Cumulative view. For a concrete cumulative view of the same (T, r) cell, Figure 5 tracks realized wealth and drawdown at a 10% target and a one-year optimization horizon, under both return conventions.



(a) Simple returns.



(b) Log returns.

Figure 5: Cumulative wealth (top) and drawdown (bottom) over the backtest window at a 10% target rate and a one-year optimization horizon. Strategy definitions as in Figure 4a.

Computational cost. The two-level scheme of Section 4.1 is fast enough for a rolling backtest. A single (T, r) optimization runs $\mathcal{O}(10^3)$ inner QPs inside the generalized-simulated-annealing outer loop and completes in a few seconds on a laptop. The full backtest grid — three target rates $\times$ four horizons $\times$ four strategies $\times$ 25 years of daily rebalancing — takes minutes end-to-end, of which the vast majority is spent

on statistical-diagnostic bookkeeping rather than on the optimization itself.

5.3 Statistical inference: the probability channel


The point-estimate probabilities of Table 1 are more informative than the paired Sharpe difference, for a specific reason. A goal-based objective deforms the return distribution rightward past rT ; Sharpe is a two-moment summary and averages away exactly that shape. Concretely, the Sharpe difference between the Optimal strategy and the S&P 500 baseline is ≈ 0.10 , but the paired-Sharpe standard error (45) at $\rho \approx 0.9$ and $T \approx 20$ years gives $SE \approx 0.12$, placing the observed difference at $\approx 0.8\sigma$ — suggestive but well short of the 2σ threshold. The strong evidence lives in the probability-gain channel: paired-difference MC noise on the goal-probability difference is ~ 5 – 10% at long horizons, making the 30–40% gains 2σ -significant. Appendix C makes precise why: the paired probability difference scales as $(1 - \rho)^{1/4}$ where ρ is the strategy–benchmark return correlation (≈ 0.9 in our sample), which puts the probability channel well into the noise-cancelling regime. For reference, the paired Sharpe difference scales as $\sqrt{1 - \rho}$; both channels benefit from pairing, but the probability channel with a slightly weaker ρ -driven suppression — the price of being a nonlinear functional of the return.

Paired Monte-Carlo protocol. Because both the strategy and the benchmark induce stationary return processes by construction, the goal-probability statistic has a well-defined ensemble expectation, and its MC standard error is estimated consistently by simulating jointly from the calibrated regime model. The paired protocol at each grid point (T, r) is:


1. Compute the realized P&L trajectories Π^{strat} and Π^{SP} of the strategy and the S&P 500 on the historical data, and their along-path goal-reaching probabilities $\hat{P}_{T,r}^{\text{strat}}$ and $\hat{P}_{T,r}^{\text{SP}}$ from (35).
2. Jointly calibrate the regime-switching model of Section 2 to those two return series.
3. Simulate L synthetic joint trajectories of length T_{tot} under the calibrated model and compute along-path probabilities on a grid of candidate target rates r' , separately for the strategy and the benchmark.
4. Pick the effective target rate r^* whose average simulated probability matches the empirical historical one; because of finite-sample deviations, r^* is typically close to r but not exactly equal.
5. Report the MC standard deviations of P_{T,r^*}^{strat} , of P_{T,r^*}^{SP} , and of their difference across the L simulated paths — the “Strategy MC std”, “Benchmark MC std”, and “Prob gain MC std” columns of Tables 4–5 in Appendix C.4.

A natural alternative would be to obtain confidence intervals by (block) bootstrapping the historical P&Ls directly. We deliberately avoid that route: any bootstrap that resamples returns — pointwise or in blocks — shuffles chunks of the trajectory relative to the VIX and regime path, breaking the mean-reverting, state-dependent structure that the optimisation is designed to exploit. Simulating jointly from the calibrated regime-switching model preserves that structure by construction and is therefore the natural noise generator for statistics whose signal comes from state-dependent drifts and transitions.

Figure 6a–6b visualises the probability gain with its MC uncertainty envelope; the tabular detail sits in Appendix C.4. At long horizons and aggressive target rates the paired-difference SE sits at ~ 5 – 10% , so the goal-reaching gains of 30–40% clear the 2σ bar comfortably; at $T = 1$ month all noise columns collapse into the single-digit floor and the point-estimate gains are correspondingly small, consistent with $\sqrt{T}$ scaling.



(a) Simple returns.



(b) Log returns.

Figure 6: Probability gain (in %, relative to the S&P 500) across target rate and investment horizon, with paired Monte-Carlo uncertainty bands. Strategy definitions as in Figure 4a; tabular detail in Appendix C.4 (Tables 4–5).

5.4 Concentration and the diversification trade-off

Figure 7 shows the unconstrained regime-conditional Optimal weights across the backtest window: at each quarterly refresh, close to 100% of the portfolio sits in a single sector per regime.




Figure 7: Optimal regime-conditional weights over the backtest window, at target rate $r = 10\%$, horizon $T = 5Y$, log returns, and default long-only bounds $u_i^{(n)} \in [0, 1]$. The top three panels show the sector weights $u_{T,r}^{(\text{opt})}(n)$ fitted for each regime ($n = 0$ expansion, $n = 1$ neutral, $n = 2$ crisis) at every quarterly calibration refresh; the bottom panel overlays the realized policy $u_{T,r}^{(\text{opt})}(z_t)$ (coloured areas) with the observed regime label z_t (black step function).

The winning sector rotates over time — Energy pre-2013, Utilities and Consumer Staples in the mid-2010s, Information Technology and Industrials more recently — but at any given calibration date only one or two sectors carry material weight. The pattern is not particular to the plotted ($r = 10\%$, $T = 5Y$) cell: across every (r, T) combination in the grid the per-regime Optimal weight histograms nearly coincide. The mechanism is a corner effect on the efficient frontier under long-only bounds and moderate cross-sectional Sharpe dispersion: the inner QP (22) saturates on the top-Sharpe single-sector portfolio over a wide range of the variance parameter v_n , so different (r, T) cells select different outer v_n^* but land on nearly the same weight vector. This concentration is intrinsic to the probability-of-reaching-a-target utility itself, not to the stationary reduction: Browne [Browne, 1999] documented the same all-in one-asset behaviour in the dynamic HJB setting. The practical consequence is that the unconstrained Optimal is not tradable at institutional scale — which is what the Hedged variant and the bounded-optimization variants below address.

Diversification sensitivity. To measure the probability cost of a hard diversification cap, we re-run the Optimal strategy under progressively tighter per-sector long-only bounds $u_i^{(n)} \in [0, 1]$, $[0, 0.5]$, and $[0, 0.25]$, forcing at least 1, 2, and 4 sectors respectively to receive positive weight in each regime. Figure 8 plots the resulting goal-reaching probability gain over the S&P 500 baseline. The gain erodes as the cap tightens, most sharply at long horizons and aggressive targets where the unconstrained edge is largest: at $r = 15\%$, $T = 5Y$ it falls from $\approx 34\%$ under $[0, 1]$ to $\approx 19\%$ under $[0, 0.5]$ to $\approx 12\%$ under $[0, 0.25]$. The gain nonetheless remains materially positive even at the tightest cap: the regime-conditional edge survives a four-sector-minimum constraint at the aggressive cells, confirming that the mechanism is a genuine regime-driven

allocation signal rather than an artifact of unconstrained concentration. Short-horizon panels ($T = 1M$) are near-flat and near-zero, consistent with the $\sqrt{T}$ scaling of the Sharpe objective. A practitioner can therefore read off, from a single row of Figure 8, the probability cost of their own governance-driven diversification bound.




Figure 8: Probability gain of the regime-conditional Optimal strategy over the S&P 500 baseline, as a function of the per-sector long-only weight cap $u_i^{(n)} \in [0, b]$ with $b \in \{1, 0.5, 0.25\}$. Three panels correspond to the target rates $r \in \{5\%, 10\%, 15\%\}$; each curve corresponds to one investment horizon $T \in \{1M, 1Y, 3Y, 5Y\}$ (log returns). The horizontal reference at zero marks parity with the passive benchmark.

5.5 Deployment considerations

Return, volatility, Sharpe. The four portfolios split into a low- and a high-risk group. The S&P 500 sits at return 12% / vol 19% and Hedged at 15%/20% (Hedged is anchored to the benchmark by construction); Gaussian delivers 17–18% at $\approx 25\%$ vol, and Optimal reaches 18–19% at $\approx 24\%$ — roughly 6% of extra return for $\approx 5\%$ of extra vol above the benchmark. Sharpe ratios: 0.75–0.77 for Optimal and Hedged, 0.68–0.70 for Gaussian, 0.66 for the S&P 500. Regime conditioning delivers most of the Sharpe lift: the Gaussian variant, which discards the regime label but is otherwise the same optimization, sits at 0.68–0.70 — above the benchmark but well below the regime-aware strategies.

Horizon dependence and turnover. At $T \lesssim 1$ year probability differences between strategies are small, consistent with the $\sqrt{T}$ scaling of the Sharpe-based objective. From $T = 3$ –5 years the Optimal delivers material uplift across all target rates, exceeding the noise envelope; the uplift widens with r : at $r = 15\%$, $P_{5Y}^{\text{opt}} = 64\%$ vs. $P_{5Y}^{\text{SP500}} = 24\%$. Average daily turnover is ~ 70 bps for the Optimal and ~ 40 bps for the Hedged; peak turnover reaches $\sim 200\%$ for the Optimal at regime transitions — a direct consequence of concentration — whereas the Hedged is capped at $q_{\text{to}} = 20\%$ by construction.

Extensions. The framework carries over unchanged under the following deployment contexts.

- *Larger universes.* The three drivers of concentration (small N , ranking bias of the probability objective, regime conditioning) all soften as N grows. On single-name equities, sector-neutral factor sleeves, or cross-asset instruments, the Optimal spreads across a wider set of near-optimal names — the mechanism generalises even though the specific sector histogram of Figure 7 would not. Computationally, the outer generalized-SA search remains K -dimensional regardless of N ; only the inner per-regime QP scales with N , and at $N = 500$ it is still trivial for modern solvers.
- *Long-short and leverage.* The framework accepts general linear and box constraints without modification: swapping $u_i^{(n)} \in [0, 1]$ for $u_i^{(n)} \in [-a, b]$ with a gross-exposure cap $\sum_i |u_i^{(n)}| \leq L$ (or a net-exposure constraint $\sum_i u_i^{(n)} = c$) drops in as an update to Ω_n , and Algorithm 2 carries over unchanged.

- *Custom regime labels.* The VIX-based partition is a demonstration choice, not a framework requirement. Any exogenous discrete state process z_t — macro-cycle indicator, credit state, central-bank regime, sentiment classifier — plugs in, provided its generator Λ can be estimated and the irreducibility assumption of Section 2.1 holds. Expert views can be encoded directly as prior-weighted transition rates.
- *Additional constraints via Ω_n .* Herfindahl caps, minimum-number-of-names constraints, factor-exposure limits, sector-neutrality relative to a benchmark, and drawdown-linked penalties all fit as linear or convex constraints in the per-regime feasible set. The diversification sensitivity of Figure 8 generalises to any such family, giving a governance-vs-probability trade-off the practitioner can price against their own risk budget.
- *Transaction costs.* Turnover is reported in Table 1 but not priced explicitly. On liquid ETF-scale instruments the round-trip cost is $\sim 5\text{--}10$ bps, so Optimal’s average daily one-sided turnover of $\sim 0.7\%$ translates into an annualised drag of only $\sim 5\text{--}10$ bps — well below the multi-percentage-point probability gains reported here. Peak-day cost is more delicate: at regime transitions Optimal can rebalance up to $\sim 200\%$ of the portfolio in a single day, carrying $10\text{--}20$ bps of round-trip cost concentrated on that day. The Hedged variant is bounded by construction — its per-sector turnover cap $q_{to} = 20\%$ per day removes the peak-day risk at the price of some probability gain, as visible in the diversification-sensitivity results. Beyond the ETF universe, single-name and less-liquid markets carry larger round-trip costs and non-trivial market impact for institutional-scale trades. Where t-costs bite harder, they can be absorbed into Ω_n as an L_1 turnover penalty $\lambda \|u^{(n)} - u_{\text{prev}}^{(n)}\|_1$: this enters as a set of linear-plus-box constraints, keeps the inner QP convex, and does not disturb the outer generalized-simulated-annealing loop.
- *Robustness.* First-order analytical noise-budget estimates in Appendix D argue that the headline probability gains at long horizons and aggressive targets survive natural perturbations of the VIX thresholds, the calibration windows, and the sample length; a numerical sensitivity sweep is left for future work.

6 Conclusion

On a proof-of-concept universe of eleven GICS sectors over 25 years, in the long-only class, the stationary regime-conditional goal-based allocation exceeds a passive S&P 500 baseline by 15–40% of goal-reaching probability at multi-year horizons and aggressive target rates, 2σ -significant under the paired Monte-Carlo standard errors of Appendix C. A state-blind single-regime collapse of the same optimization captures only part of the lift, isolating regime conditioning — rather than the optimization apparatus itself — as the source of the gain.

The result is not a numerical accident. Because the policy $u(z)$ depends only on the current regime label, both the weight process and the induced return process are strictly stationary, and ergodicity of the regime chain aligns the along-path frequency of goal attainment with the model-implied utility. This identity between historical backtest and model expectation is a structural feature of the stationary formulation that time-dependent HJB frameworks do not possess, and it is what makes the proof-of-concept magnitude informative: the mechanism carries over to larger universes and different partitions, even where specific numbers do not.

Two lessons about goal-based portfolios generalise beyond the specific setup. First, a probability-of-target objective deforms the return distribution rightward past rT ; Sharpe, being a two-moment summary, averages away exactly the shape being optimised, so the statistical evidence for such a strategy lives in the probability channel rather than the Sharpe channel. The paired-difference standard errors of Appendix C are the inference machinery that makes this signal legible. Second, the same objective is intrinsically concentrating: under long-only constraints with moderate cross-sectional Sharpe dispersion, the unconstrained optimum places nearly 100% of the portfolio on the top-Sharpe sector per regime. This is Browne’s phenomenon in a stationary setting, and it motivates the tracking-and-turnover-bounded *Hedged* variant which recovers most of the probability gain in an institutionally tradable form.

A natural open direction is the aspirational-target regime, $r > \bar{m}^{\max}$, where the sign-analysis criterion of Section 4 fails and the soft- and hard-form outer optima genuinely diverge. Every (T, r) cell in our experiments lies safely inside the fast convex range; the reach-for-the-moon targets — and the per-regime DCCP fallbacks that unlock them — warrant a dedicated study, both in terms of the additional probability gain the hard form can capture and in terms of the deployment cost of the local-only convergence of the non-convex inner solve. On the empirical side, the analytical robustness bounds of Appendix D bracket the exposure of the headline gains to natural perturbations of the setup; a full numerical sensitivity sweep across VIX thresholds and calibration windows is a complementary follow-up.

Taken together, the results argue that the pairing of a stationary state-only policy with a probability-of-target objective is a distinctive and defensible way to think about long-horizon goal-based allocation: it earns its keep both statistically — through the paired probability channel that makes small regime-driven distributional shifts legible — and operationally — through backtestability that other frameworks give up.

Acknowledgements

The authors thank their colleagues at ADIA for valuable discussions. Any remaining errors are our own.

References

- [Alqubaisi et al., 2026] Alqubaisi, A., Antonov, A., and Muravey, D. (2026). Stationary portfolio goal-based optimization. SSRN working paper.
- [Bauman et al., 2024] Bauman, T., Goluža, S., Gasperov, B., and Kostanjcar, Z. (2024). Deep reinforcement learning for goal-based investing under regime-switching. In Lutchyn, T., Ramírez Rivera, A., and Ricaud, B., editors, *Proceedings of the 5th Northern Lights Deep Learning Conference (NLDL)*, volume 233 of *Proceedings of Machine Learning Research*, pages 13–19. PMLR.
- [Bielecki and Pliska, 1999] Bielecki, T. R. and Pliska, S. R. (1999). Risk-sensitive dynamic asset management. *Applied Mathematics & Optimization*, 39(3):337–360.
- [Browne, 1999] Browne, S. (1999). Reaching goals by a deadline: Digital options and continuous-time active portfolio management. *Advances in Applied Probability*, 31:551–577. Reprinted in *Foundations of Financial Economics*, World Scientific, 2013.
- [Çınlar, 1972] Çınlar, E. (1972). Markov additive processes. II. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 24(2):95–121.
- [Nystrup et al., 2016] Nystrup, P., Madsen, H., and Lindström, E. (2016). Dynamic portfolio optimization across hidden market regimes. *Quantitative Finance*, 18:83 – 95.
- [Sass and Haussmann, 2004] Sass, J. and Haussmann, U. G. (2004). Optimizing the terminal wealth under partial information: The case of bounded portfolio positions. *Integral Transforms and Special Functions*, 15(5):445–452.
- [Tsallis and Stariolo, 1996] Tsallis, C. and Stariolo, D. A. (1996). Generalized simulated annealing. *Physica A*, 233(1–2):395–406.
- [Xiang et al., 1997] Xiang, Y., Sun, D. Y., Fan, W., and Gong, X. G. (1997). Generalized simulated annealing algorithm and its application to the Thomson model. *Physics Letters A*, 233(3):216–220.
- [Zhou and Yin, 2003] Zhou, X. Y. and Yin, G. (2003). Markowitz’s mean-variance portfolio selection with regime switching: A continuous-time model. *SIAM Journal on Control and Optimization*, 42(4):1466–1491.

A Derivation of the Moment Generating Function

Recall that the cumulative simple return of the portfolio at T reads

$$X_T = \int_0^T u(z_t)^\top (\mu(z_t) dt + \sigma(z_t) dW_t).$$

Its moment generating function (MGF) $\mathbb{E}[\exp(\xi X_T) \mid z_0 = k]$ can be simplified by integrating out dW (which is independent of the regime process):

$$q_T(\xi|k) = \mathbb{E}[e^{\xi X_T} | z_0 = k] = \mathbb{E}\left[e^{\xi \int_0^T u(z_t)^\top \mu(z_t) dt + \frac{1}{2} \xi^2 \int_0^T u(z_t)^\top C(z_t) u(z_t) dt} \mid z_0 = k\right]$$

where cov matrix $C(z) = \sigma(z) \sigma(z)^\top$.

The MGF can be presented as

$$q_T(\xi|k) = \mathbb{E}\left[e^{\sum_n Q_n(\xi) \int_0^T \mathbb{1}_{z_s=n} ds} \mid z_0 = k\right]$$

with quadratic functions

$$Q_n(\xi) = \xi u^{(n)\top} \mu^{(n)} + \frac{1}{2} \xi^2 u^{(n)\top} C^{(n)} u^{(n)}$$

The moment generation function can be calculated as follows. Denote

$$f(t) = \mathbb{E}\left[e^{\int_0^t \sum_n q_n \mathbb{1}_{z_s=n} ds} \mid F_t\right] = e^{\int_0^t \sum_n q_n \mathbb{1}_{z_s=n} ds} g(t, z_t)$$

The process $f(t)$ is martingale, thus,

$$0 = \mathbb{E}[df(t) \mid F_t] = e^{\int_0^t \sum_n q_n \mathbb{1}_{z_s=n} ds} \mathbb{E}[dg(t, z_t) \mid F_t] + \left(\sum_n Q_n \mathbb{1}_{z_t=n}\right) e^{\int_0^t \sum_n q_n \mathbb{1}_{z_s=n} ds} g(t, z_t)$$

The differential expectation can be easily calculated

$$\begin{aligned} \mathbb{E}[dg(t, z_t) \mid z_t = k] &= \mathbb{E}[g(t+dt, z_{t+dt}) - g(t, z_t) \mid z_t = k] = \mathbb{E}\left[\sum_n g^{(n)}(t+dt) \mathbb{1}_{z_{t+dt}=n} - g^{(k)}(t) \mid z_t = k\right] \\ &= \sum_n g^{(n)}(t+dt) \mathbb{E}[\mathbb{1}_{z_{t+dt}=n} \mid z_t = k] - g^{(k)}(t) = \sum_n g^{(n)}(t+dt) (\delta_{nk} - \Lambda_{kn} dt) - g^{(k)}(t) \\ &= \left(\partial_t g^{(k)}(t) + \sum_n \Lambda_{kn} g^{(n)}(t)\right) dt \end{aligned}$$

Also, we have

$$\mathbb{E}\left[\left(\sum_n q_n \mathbb{1}_{z_t=n}\right) g(t, z_t) \mid z_t = k\right] = q_k g^{(k)}(t)$$

Finally, the components $g^{(k)}(t)$ of the function $g(t, z)$ satisfy a linear vector equation

$$\partial_t g^{(k)}(t) + \sum_n \Lambda_{kn} g^{(n)}(t) + q_k g^{(k)}(t) = 0$$

Denoting matrix $A_{kn} = \Lambda_{kn} + \delta_{kn} q_k$, we obtain the following linear equation

$$(37) \quad \partial_t G(t) = -A G(t) \quad \text{for} \quad G(T) = \mathbb{1}$$

where column $G(t) = [g^{(1)}, \dots, g^{(K)}]^\top$. Clearly, its solution reads

$$(38) \quad G(t) = e^{(T-t)A(t)} \mathbb{1}$$

Coming back to the exponential expectation, we obtain

$$\mathbb{E} \left[e^{\sum_n q_n \int_0^T \mathbb{1}_{z_s=n} ds} | z_0 = k \right] = \left\{ e^{T(\Lambda + \text{diag}(\mathbf{q}))} \mathbb{1} \right\}_k$$

for the diagonal matrix $(\text{diag}(\mathbf{q}))_{nk} = \delta_{nk} q_n$. Then, the model moment generating function reads

$$q_T(\xi|k) = \mathbb{E} \left[e^{\sum_n Q_n(\xi) \int_0^T \mathbb{1}_{z_s=n} ds} | z_0 = k \right] = \left\{ e^{T(\Lambda + \xi D_m + \frac{1}{2} \xi^2 D_v)} \mathbb{1} \right\}_k$$

where the diagonal matrices are

$$(D_m)_{nk} = \delta_{nk} u^{(n)\top} \mu^{(n)} \quad \text{and} \quad (D_v)_{nk} = \delta_{nk} u^{(n)\top} C^{(n)} u^{(n)}.$$

B Sign Criterion for the Soft Form (General K)

This appendix collects the differentiation that produces the sign criterion of Section 4.4. Substituting $m_n = \bar{\rho}_n(v_n)$ into (21), the argument of $\mathcal{N}$ inside the expectation is

$$(39) \quad \bar{\mathcal{S}}_T(\tau; v) = \sqrt{T} \frac{\bar{M}(\tau, v) - r}{\sqrt{\bar{V}(\tau, v)}}, \quad \bar{M}(\tau, v) \equiv \sum_n \tau_n \bar{\rho}_n(v_n), \quad \bar{V}(\tau, v) \equiv \sum_n \tau_n v_n,$$

and differentiating at fixed τ ,

$$(40) \quad \frac{\partial \bar{\mathcal{S}}_T}{\partial v_k} = \frac{\tau_k \sqrt{T}}{\sqrt{\bar{V}}} \left[\bar{\rho}'_k(v_k) - \frac{1}{2} \cdot \frac{\bar{M} - r}{\bar{V}} \right].$$

This is the $K = 1$ bracket (29) with $\bar{\rho}'_k(v_k)$ replacing $\rho'(v)$ and the τ -weighted ratio $(\bar{M} - r)/\bar{V}$ replacing $(\rho - r)/v$. The log-return objective adds an extra $-\frac{1}{2}$ inside the bracket, with

$$(41) \quad \bar{M}^L(\tau, v) \equiv \sum_n \tau_n (\bar{\rho}_n(v_n) - \frac{1}{2} v_n)$$

in place of $\bar{M}$. The outer utility partial then reads

$$(42) \quad \frac{\partial U_T(\bar{\rho}(v), v)}{\partial v_k} = \mathbb{E}_\infty \left[\phi(\bar{\mathcal{S}}_T) \frac{\partial \bar{\mathcal{S}}_T}{\partial v_k} \right]$$

(ϕ the standard normal density; log analogue with $\bar{\mathcal{S}}_T^{(L)}$). Evaluating at $v_k = \bar{v}_k^{\max}$ (where $\bar{\rho}'_k(\bar{v}_k^{\max}) = 0$),

$$(43) \quad \left. \frac{\partial \bar{\mathcal{S}}_T}{\partial v_k} \right|_{v_k = \bar{v}_k^{\max}} = -\frac{\tau_k \sqrt{T}}{2 \bar{V}^{3/2}} (\bar{M} - r), \quad \left. \frac{\partial \bar{\mathcal{S}}_T^{(L)}}{\partial v_k} \right|_{v_k = \bar{v}_k^{\max}} = -\frac{\tau_k \sqrt{T}}{2 \bar{V}^{3/2}} (\bar{M}^L - r + \bar{V}),$$

and substituting into (42) yields the outer utility partial at the peak, (30), which drives the criterion of Section 4.4. The log-case bracket carries an extra $+\bar{V}$ shift, which is why the log case is uniformly more easily negative than the simple case.

C Standard Errors for Paired Sharpe and Probability Differences

This appendix collects the closed-form standard-error formulas that underlie the paired Monte-Carlo protocol described in Section 5 under *Confidence intervals*. Throughout, $R^{(1)}, R^{(2)}$ are two return series with correlation $\rho = \text{Corr}(R^{(1)}, R^{(2)})$ observed over T years; we derive the standard error of the paired Sharpe-ratio difference and then of the paired goal-probability difference.

C.1 Paired Sharpe-ratio difference

Let S_1, S_2 be the annualised Sharpe ratios of the two series. For reference, the standard error of an *individual* annualised Sharpe estimator $\hat{S}$ from T years of i.i.d. returns is (Lo 2002)

$$(44) \quad \text{SE}(\hat{S}) = \sqrt{\frac{1 + \frac{1}{2}S^2}{T}},$$

the noise a reader would face for either strategy in isolation. Under i.i.d. returns, Jobson–Korkie (1981) and Memmel (2003) give the standard error of the paired difference $\hat{S}_1 - \hat{S}_2$ as

$$(45) \quad \text{SE}(\hat{S}_1 - \hat{S}_2) = \sqrt{\frac{1}{T} \left[\underbrace{2(1-\rho)}_{\text{shared-noise cancellation}} + \underbrace{\frac{1}{2}(S_1^2 + S_2^2 - 2\rho^2 S_1 S_2)}_{\text{drift-magnitude correction}} \right]}.$$

Two limiting cases make the ρ -dependence transparent. If the two strategies are perfectly correlated ($\rho \rightarrow 1$), both bracketed terms vanish and the SE collapses to zero: whatever noise moves $\hat{S}_1$ moves $\hat{S}_2$ by the same amount, and it drops out on subtraction. If instead they are uncorrelated ($\rho = 0$), the shared-noise piece reaches its maximum value $2/T$ and the SE approaches what one gets by pooling the two individual Sharpe SEs $\sqrt{(1 + \frac{1}{2}S^2)/T}$ in quadrature. Realistic strategies live in between: they typically share a big chunk of calendar-time volatility with their benchmark, so pairing buys them a lot of noise reduction, but not all of it. In the practically relevant regime ρ close to 1, the shared-noise piece dominates and one has the useful leading-order approximation

$$(46) \quad \text{SE}(\hat{S}_1 - \hat{S}_2) \approx \sqrt{\frac{2(1-\rho)}{T}},$$

which makes the ρ -driven noise reduction transparent: halving $(1-\rho)$ shrinks the SE by a factor $\sqrt{2} \approx 1.4$.

For a numerical illustration matched to the main text, take Optimal vs. S&P 500 ($S_1 \approx 0.76$, $S_2 \approx 0.66$) with $\rho \approx 0.9$ over an effective $T \approx 20$ -year backtest window: the leading-order approximation (46) gives $\text{SE} \approx \sqrt{0.2/20} = 0.10$, and the full formula (45) evaluates to $\text{SE} \approx 0.12$ (the drift-magnitude correction adds only about 0.02). The observed difference $\hat{S}_1 - \hat{S}_2 \approx 0.10$ therefore lives at roughly 0.8σ — suggestive but well short of the usual 2σ significance bar. For scale, treating the two Sharpes as if they came from independent samples would give each individual estimator an SE of about 0.22–0.25; pairing at $\rho \approx 0.9$ roughly halves the noise, but half of “too large” is still “too large” at 20 years of data.

C.2 Paired goal-probability difference

Denote by $R^{(1)}, R^{(2)}$ the horizon- T returns, jointly distributed with correlation ρ and joint CDF

$$F(x_1, x_2) = \mathbb{P}[R^{(1)} \leq x_1, R^{(2)} \leq x_2].$$

The marginal and joint goal probabilities are

$$(47) \quad p_i = 1 - F_{R^{(i)}}(rT), \quad p_{12} = 1 - F_{R^{(1)}}(rT) - F_{R^{(2)}}(rT) + F(rT, rT).$$

For reference, the standard error of an *individual* probability estimator $\hat{p}_i$ from n i.i.d. Bernoulli trials is

$$(48) \quad \text{SE}(\hat{p}_i) = \sqrt{\frac{p_i(1-p_i)}{n}},$$

the standard binomial noise. The variance of the paired difference $\hat{p}_1 - \hat{p}_2$ from n i.i.d. paired trials is

$$(49) \quad \text{Var}(\hat{p}_1 - \hat{p}_2) = \frac{1}{n} \left[p_1(1-p_1) + p_2(1-p_2) - 2(p_{12} - p_1 p_2) \right].$$

The covariance term $p_{12} - p_1 p_2$ enters through the joint CDF and is a monotone function of ρ : perfect comovement pushes $p_{12} \rightarrow \min(p_1, p_2)$, maximises the covariance, and collapses the difference-noise; independence gives $p_{12} = p_1 p_2$, kills the covariance, and returns us to the pooled-independent-samples baseline. This is the same shared-calendar-time cancellation mechanism as in the Sharpe formula (45), just filtered through the joint CDF rather than through the first two moments.

When $p_1 \approx p_2 \equiv p$ and ρ is close to 1 — the regime relevant to a strategy compared to its own benchmark — we can approximate p_{12} by expanding the joint density near the diagonal threshold. Modelling $(R^{(1)}, R^{(2)})$ as jointly Gaussian with equal marginals and correlation ρ , a leading-order calculation in $1 - \rho$ gives

$$(50) \quad p - p_{12} \approx \phi(\Phi^{-1}(p)) \sqrt{\frac{1 - \rho}{\pi}},$$

where ϕ and Φ are the standard normal density and CDF, so $\Phi^{-1}(p)$ is the marginal p -quantile in standardised units and $\phi(\Phi^{-1}(p))$ is the standard normal density evaluated there.

Sketch of the derivation of (50). Introduce the standardised threshold $z = \Phi^{-1}(1 - p)$, so that $p = \Phi(-z)$ and $p_{12} = \Phi_2(-z, -z; \rho)$, where $\Phi_2(\cdot, \cdot; \rho)$ is the standard bivariate normal CDF with correlation ρ . Plackett's identity for its derivative in ρ ,

$$(51) \quad \frac{\partial}{\partial \rho} \Phi_2(a, b; \rho) = \frac{1}{2\pi\sqrt{1 - \rho^2}} \exp\left(-\frac{a^2 - 2\rho ab + b^2}{2(1 - \rho^2)}\right),$$

evaluated at $a = b = -z$, gives $a^2 - 2\rho ab + b^2 = 2z^2(1 - \rho)$ and $2(1 - \rho^2) = 2(1 - \rho)(1 + \rho)$, so

$$\frac{\partial p_{12}}{\partial \rho} = \frac{1}{2\pi\sqrt{1 - \rho^2}} \exp\left(-\frac{z^2}{1 + \rho}\right) \stackrel{\rho \rightarrow 1}{\sim} \frac{\phi(z)}{\sqrt{2\pi} \sqrt{2(1 - \rho)}}.$$

Integrating from ρ up to 1 (where $p_{12} \rightarrow p$) yields

$$p - p_{12}(\rho) = \int_{\rho}^1 \frac{\partial p_{12}}{\partial t} dt \sim \frac{\phi(z)}{\sqrt{2\pi}} \cdot \sqrt{2(1 - \rho)} = \phi(z) \sqrt{\frac{1 - \rho}{\pi}},$$

which is (50) after noting $\phi(z) = \phi(-z) = \phi(\Phi^{-1}(p))$.

Substituting into (49), which for $p_1 = p_2 = p$ reduces to $\text{Var}(\hat{p}_1 - \hat{p}_2) = 2(p - p_{12})/n$, yields the compact simplification

$$(52) \quad \text{SE}(\hat{p}_1 - \hat{p}_2) \approx \sqrt{\frac{2\phi(\Phi^{-1}(p))}{n}} \left(\frac{1 - \rho}{\pi}\right)^{1/4},$$

the direct probability-space analogue of the Sharpe leading-order formula (46). The $\rho \rightarrow 1$ cancellation is at work in both, with one notable difference: the paired-Sharpe SE vanishes as $\sqrt{1 - \rho}$, whereas the paired-probability SE vanishes only as $(1 - \rho)^{1/4}$. The slower rate reflects the fact that a probability is a nonlinear (indicator) functional of the return, so paired subtraction only picks up the “curvature” of the joint density along the diagonal threshold rather than a full cancellation of the shared-noise term.

C.3 Overlapping-window trials in the historical backtest

One complication keeps us from using (49) directly on the historical estimator. The frequency (35) sums indicators over N_t *overlapping* rolling windows on a single realised path, so successive trials share a substantial chunk of returns and are auto-correlated — the effective sample size is smaller than the raw window count. Rather than model this dependence in closed form, we estimate $\text{SE}(\hat{p}_1 - \hat{p}_2)$ directly by the paired Monte-Carlo protocol described under *Confidence intervals* in Section 5. That estimator is the operational analogue of (49) for dependent trials and delivers the difference-noise reported in the *Prob gain MC std* column of Tables 4–5.

C.4 Backtest tables: log-return convention and probability-noise envelopes

This subsection collects the three backtest tables deferred from Section 5: the log-return counterpart of the simple-return backtest table (Table 3), and the simple- and log-return probability-noise tables produced by the paired Monte-Carlo protocol of Section 5.3 (Tables 4–5).

Target rate (%)	Horizon	Strategy	Prob (%)	Sharpe	Return (%)	Vol (%)	Daily (Max) turnover (%)
5	1M	optimal	62	0.75	18	24	0.7 (200.0)
5	1M	gaussian	61	0.68	17	25	0.6 (200.0)
5	1M	hedged	63	0.74	15	20	4.6 (50.0)
5	1M	sp500	63	0.66	12	19	0.5 (16.4)
5	1Y	optimal	81	0.76	18	24	0.7 (200.0)
5	1Y	gaussian	73	0.68	17	25	0.6 (200.0)
5	1Y	hedged	79	0.75	15	20	4.6 (50.0)
5	1Y	sp500	76	0.66	12	19	0.5 (16.4)
5	3Y	optimal	89	0.76	18	24	0.7 (200.0)
5	3Y	gaussian	87	0.68	17	25	0.6 (200.0)
5	3Y	hedged	87	0.75	15	20	4.6 (50.0)
5	3Y	sp500	84	0.66	12	19	0.5 (16.4)
5	5Y	optimal	91	0.76	18	24	0.7 (200.0)
5	5Y	gaussian	92	0.68	17	25	0.6 (200.0)
5	5Y	hedged	85	0.75	15	20	4.7 (50.0)
5	5Y	sp500	78	0.66	12	19	0.5 (16.4)
10	1M	optimal	59	0.76	18	24	0.7 (200.0)
10	1M	gaussian	58	0.66	17	25	0.7 (200.0)
10	1M	hedged	60	0.75	15	20	4.6 (50.0)
10	1M	sp500	60	0.66	12	19	0.5 (16.4)
10	1Y	optimal	69	0.76	18	24	0.7 (200.0)
10	1Y	gaussian	63	0.66	17	25	0.7 (200.0)
10	1Y	hedged	66	0.75	15	20	4.8 (50.0)
10	1Y	sp500	63	0.66	12	19	0.5 (16.4)
10	3Y	optimal	79	0.76	18	24	0.7 (200.0)
10	3Y	gaussian	72	0.66	17	25	0.7 (200.0)
10	3Y	hedged	81	0.76	15	20	4.7 (50.0)
10	3Y	sp500	69	0.66	12	19	0.5 (16.4)
10	5Y	optimal	75	0.76	18	24	0.7 (200.0)
10	5Y	gaussian	73	0.66	17	25	0.7 (200.0)
10	5Y	hedged	71	0.76	15	20	4.7 (50.0)
10	5Y	sp500	65	0.66	12	19	0.5 (16.4)
15	1M	optimal	56	0.77	19	24	0.8 (200.0)
15	1M	gaussian	56	0.67	17	25	0.6 (200.0)
15	1M	hedged	56	0.76	15	20	4.8 (50.0)
15	1M	sp500	55	0.66	12	19	0.5 (16.4)
15	1Y	optimal	57	0.76	18	24	0.7 (200.0)
15	1Y	gaussian	54	0.67	17	25	0.6 (200.0)
15	1Y	hedged	50	0.77	15	20	4.7 (50.0)
15	1Y	sp500	46	0.66	12	19	0.5 (16.4)
15	3Y	optimal	59	0.76	18	24	0.7 (200.0)
15	3Y	gaussian	48	0.67	17	25	0.6 (200.0)
15	3Y	hedged	42	0.76	15	20	4.7 (50.0)
15	3Y	sp500	24	0.66	12	19	0.5 (16.4)
15	5Y	optimal	57	0.76	18	24	0.7 (200.0)
15	5Y	gaussian	46	0.67	17	25	0.6 (200.0)
15	5Y	hedged	38	0.76	15	20	4.8 (50.0)
15	5Y	sp500	23	0.66	12	19	0.5 (16.4)

Table 3: Regime-based goal backtest (log returns). Log-return counterpart of Table 1; values agree with the simple-return table within 0.01–0.02 on Sharpe and $\sim 1\%$ on return/vol.

Target rate (%)	Horizon	Strategy	Backtest prob (%)	Strategy MC std (%)	Backtest benchmark prob (%)	Benchmark MC std (%)	Backtest prob gain (%)	Prob gain MC std (%)
5	1M	optimal	63	3	64	3	-1	2
5	1M	gaussian	62	3	64	3	-2	2
5	1M	hedged	64	3	64	3	0	2
5	1Y	optimal	85	7	79	8	6	5
5	1Y	gaussian	79	7	79	8	-1	4
5	1Y	hedged	83	7	79	8	4	3
5	3Y	optimal	97	9	87	10	10	6
5	3Y	gaussian	97	10	87	10	10	6
5	3Y	hedged	92	9	87	10	4	4
5	5Y	optimal	100	9	84	11	15	7
5	5Y	gaussian	100	11	84	11	16	6
5	5Y	hedged	97	9	84	11	13	4
10	1M	optimal	60	3	60	3	0	3
10	1M	gaussian	59	3	60	3	-1	2
10	1M	hedged	61	3	60	3	1	2
10	1Y	optimal	74	7	66	8	9	5
10	1Y	gaussian	68	8	66	8	3	5
10	1Y	hedged	69	8	66	8	3	3
10	3Y	optimal	89	9	72	10	17	6
10	3Y	gaussian	82	11	72	12	9	6
10	3Y	hedged	84	10	72	11	12	4
10	5Y	optimal	91	10	69	12	22	8
10	5Y	gaussian	90	10	69	11	21	6
10	5Y	hedged	80	15	69	17	11	6
15	1M	optimal	57	3	55	3	2	3
15	1M	gaussian	56	3	55	3	1	2
15	1M	hedged	56	3	55	3	1	2
15	1Y	optimal	61	8	46	9	16	6
15	1Y	gaussian	58	8	46	9	12	5
15	1Y	hedged	53	9	46	9	7	3
15	3Y	optimal	66	14	27	15	39	10
15	3Y	gaussian	54	14	27	14	28	9
15	3Y	hedged	47	14	27	14	21	6
15	5Y	optimal	64	18	24	18	39	13
15	5Y	gaussian	60	18	24	18	36	11
15	5Y	hedged	41	18	24	17	16	7

Table 4: Paired probability-noise envelope, simple returns. For each $(r, T, \text{strategy})$ cell the table reports the historical probabilities of strategy and benchmark and their paired Monte-Carlo standard deviations, together with the paired gain and its MC standard deviation. The gain-noise column is systematically several times smaller than either individual-strategy noise column at horizons $T \geq 1$ year, consistent with the ρ -driven cancellation predicted by (49).

Target rate (%)	Horizon	Strategy	Backtest prob (%)	Strategy MC std (%)	Backtest benchmark prob (%)	Benchmark MC std (%)	Backtest prob gain (%)	Prob gain MC std (%)
5	1M	optimal	62	3	64	3	-2	2
5	1M	gaussian	61	3	64	3	-2	2
5	1M	hedged	63	3	64	3	0	2
5	1Y	optimal	81	8	76	8	6	5
5	1Y	gaussian	73	8	76	8	-2	4
5	1Y	hedged	79	8	76	8	4	3
5	3Y	optimal	89	11	84	12	5	7
5	3Y	gaussian	87	12	84	12	3	6
5	3Y	hedged	87	11	84	12	2	4
5	5Y	optimal	91	12	78	14	13	8
5	5Y	gaussian	92	14	78	14	14	8
5	5Y	hedged	85	12	78	14	7	5
10	1M	optimal	59	3	60	3	0	3
10	1M	gaussian	58	3	60	3	-1	2
10	1M	hedged	60	3	60	3	0	2
10	1Y	optimal	69	8	63	8	6	5
10	1Y	gaussian	63	8	63	8	-1	5
10	1Y	hedged	66	8	63	8	2	3
10	3Y	optimal	79	12	69	12	10	7
10	3Y	gaussian	72	13	69	13	3	7
10	3Y	hedged	81	11	69	12	12	4
10	5Y	optimal	75	16	65	18	10	11
10	5Y	gaussian	73	16	65	16	9	9
10	5Y	hedged	71	16	65	18	7	6
15	1M	optimal	56	3	55	3	1	3
15	1M	gaussian	56	3	55	3	1	2
15	1M	hedged	56	3	55	3	1	2
15	1Y	optimal	57	9	46	9	12	6
15	1Y	gaussian	54	9	46	9	8	5
15	1Y	hedged	50	9	46	8	5	3
15	3Y	optimal	59	14	24	14	36	9
15	3Y	gaussian	48	15	24	14	25	8
15	3Y	hedged	42	14	24	14	18	5
15	5Y	optimal	57	18	23	18	34	12
15	5Y	gaussian	46	18	23	18	24	10
15	5Y	hedged	38	18	23	16	15	7

Table 5: Paired probability-noise envelope, log returns. As Table 4 but under the log-return convention.

D Analytical Robustness of the Empirical Findings

The empirical block of Section 5 rests on a specific set of choices: a VIX-based three-regime partition with thresholds $[0.20, 0.30]$, per-regime windows $\{5, 7, 12\}$ years, an eleven-sector universe, and 25 years of daily data. This appendix does not attempt a full numerical sensitivity sweep — that is left for future work — but gives first-order analytical estimates of how each natural perturbation to those choices would move the reported probability gains. The bounds use only quantities already tabulated in Table 2 and the paired-difference formulas of Appendix C, and are meant as order-of-magnitude brackets rather than tight numerical claims.

R.1 VIX threshold shifts. A shift of ± 5 points in the practitioner thresholds $[0.20, 0.30]$ (moving to $[0.15, 0.25]$ or $[0.25, 0.35]$) relabels approximately 5–10% of daily observations between adjacent regimes,

since the VIX distribution is dense near the modal value of 20. On the relabelled slice, the effective per-regime drift shifts by $\Delta_n \equiv \bar{\rho}_n(\text{new}) - \bar{\rho}_n(\text{baseline})$, bounded above by the between-regime drift gap in Table 2 ($\sim 20\text{--}50\%$ in absolute terms) times the fraction relabelled. Propagating through the first-order Taylor expansion of U_T ,

$$|\Delta U_T| \lesssim \sqrt{T} \phi(\bar{\mathcal{S}}) \cdot |\Delta_n| \cdot \tau_n,$$

where τ_n is the occupation fraction of the affected regime, gives an upper bound of a few percentage points on the aggressive-cell gain. The residual gain therefore remains 2σ -significant.

R.2 Calibration windows. Halving each per-regime window (from $\{5, 7, 12\}$ years to $\{2.5, 3.5, 6\}$) doubles the sampling noise of the moment estimator, so $\text{CV}(\mu)$ in Table 2 rises by a factor of $\sqrt{2} \approx 1.4$. Propagating this through the same first-order argument as R.1 widens the probability-gain noise band by a comparable factor. Since the baseline $\text{CV}(\mu)$ values are already small relative to the 30–40% probability edge, the widened band does not consume the gain at the aggressive cells. In the other direction, longer windows reduce responsiveness — the trade-off flagged in Section 5.2 that motivates the rarity-scaled window design in the first place.

R.3 Sample length. The paired-difference standard errors of Appendix C scale as $1/\sqrt{N_{\text{eff}}}$, where N_{eff} is the effective sample size after overlapping-window autocorrelation. Halving the total sample from 25 to ~ 12 years raises the paired SE by $\sqrt{2}$, from $\sim 5\text{--}10\%$ at long horizons to $\sim 7\text{--}14\%$. The aggressive (T, r) cell’s 39% gain still clears the 2σ threshold; moderate-cell gains near the noise floor lose significance, consistent with the paper’s caveat that short-horizon gains are small.

Summary. Across all three perturbations, the probability gain at long horizons and aggressive target rates remains positive and materially non-zero. The strong-headline cell ($r = 15\%, T = 5Y$) with 39% gain retains 2σ -significance under all worst-case combinations of the above bounds. What deteriorates first is the significance of the moderate cells ($r = 5\%, T = 1Y$), which sit close to the noise floor in the baseline and become indistinguishable from zero under aggressive perturbations. This is consistent with the $\sqrt{T}$ -scaling of the Sharpe channel and does not affect the paper’s core claim.

E Supplementary Figures

This appendix collects additional diagnostics for the backtest of Section 5. Figures 9 and 10 show the risk-return trade-off across strategies and horizons; Figures 11 and 12 display mean sector allocations per regime for the 10% target rate. Strategy definitions are as in Figure 4a. In the scatter plots, *colors identify strategies* while *marker shapes identify horizons*.




Figure 9: Risk-return scatter across strategies and horizons, simple returns. Colors identify strategies; marker shapes identify horizons.




Figure 10: Risk-return scatter across strategies and horizons, log returns. Colors identify strategies; marker shapes identify horizons.

Mean sector allocations by regime: target 10% (simple returns)




Figure 11: Mean sector allocations per regime at a 10% target rate, simple returns.

Mean sector allocations by regime: target 10% (log returns)




Figure 12: Mean sector allocations per regime at a 10% target rate, log returns.