

Archimedean Lever and Center-of-Mass Portfolios: From Mechanical Equilibria to Portfolio Allocations

DIMITRIOS THOMAKOS^{*†}

National and Kapodistrian University of Athens[‡]

PARIS VLACHOS

National and Kapodistrian University of Athens.

AGGELOS GIOVANIS

National and Kapodistrian University of Athens.

This paper reinterprets portfolio selection through two Archimedean mechanical principles—the lever law and center-of-mass computation—formalized as explicit constraints within classical mean-variance and factor-model frameworks. The Archimedean Lever Portfolio (ALP) requires asset-specific torques to be proportional to risk contributions; under standard mean-variance assumptions this yields a closed-form solution structurally identical to the tangency portfolio, but using the volatility-scaled signal $\sigma_i \mu_i$ in place of raw expected returns μ_i —an implicit Bayes-Stein shrinkage that lowers turnover and drawdowns. The Archimedean Center-of-Mass (ACM) portfolio is a minimum-variance allocation whose factor center of mass is aligned with the factor risk premium, yielding a second closed-form solution from a low-dimensional linear system. Empirical tests on U.S. ETF equity portfolios show that Archimedean portfolios deliver competitive risk-adjusted returns across multiple estimation windows and data frequencies, with lower turnover and smaller drawdowns than the classical tangency portfolio, though results are subject to the usual caveats of a small-universe backtest.

Keywords: Portfolio selection, Archimedean mechanics, lever equilibrium, center of mass, mean–variance optimization, factor models, dynamic trading rules.

JEL Classifications: G11, C61, C58

^{*}The authors would like to thank the Editor and an anonymous referee for constructive comments and suggestions that improved the content and presentation of the paper. All remaining errors are ours.

[†]Corresponding author: dthomakos@ba.uoa.gr.

[‡]All: Cybernetics and Artificial Intelligence Laboratory (CAIL), Department of Business Administration.

© 2026 Dimitrios Thomakos, Paris Vlachos and Aggelos Giovanis. Licensed under the Creative Commons Attribution-Noncommercial 3.0 Licence (<http://creativecommons.org/licenses/by-nc/3.0/>). Available at <http://rofea.org>.

1 Introduction

Classical portfolio theory is organized around a geometric picture in mean–variance space: efficient frontiers are parabolic curves, and the optimal risky portfolio is the tangency point between the efficient frontier and the capital market line (Markowitz, 1952; Sharpe, 1964). While this picture is suggestive, it is historically associated more with 20th-century developments than with ancient geometry: tangency is loosely related to conics; leverage is compared to physical levers; balance is compared to centers of mass (Merton, 1972). However, the resulting strategies typically rederive familiar mean variance objects, leaving the mechanical analogy at the level of metaphor rather than constraint.

In this paper we attempt a different approach. We take Archimedes’ actual mechanical ideas—the lever law and the computation of centers of mass—and transform them to explicit mathematical constraints that define new portfolio rules. The goal is to identify closed-form allocations whose optimality conditions can be read directly as equilibrium conditions of a mechanical system. Our contributions are fourfold. First, we define a notion of torque for each asset and derive the Archimedean Lever Portfolio (ALP) by requiring that torques be proportional to classical risk contributions. This yields a closed form solution that is the tangency portfolio of the mean variance problem with volatility scaled return signals $\tilde{\mu}_i = \sigma_i \mu_i$, a robust, shrinkage-corrected restatement of the classical solution that reduces sensitivity to estimation error in expected returns. Second, we work in a multi-factor setting and define Archimedean Center-of-Mass (ACM) portfolios by aligning the portfolio’s factor center of mass with the factor risk premium, subject to minimum-variance. The resulting allocation is again available in closed form, via a linear system with a number of factors. Third, we propose dynamic trading rules where deviations from Archimedean equilibrium, measured via torque misalignment, trigger rebalancing. Fourth, we establish decision-theoretic micro-foundations for both strategies: the ALP maximizes a torque-to-volatility ratio and embeds an implicit shrinkage of expected returns, while the ACM implements a portfolio-level APT pricing condition; together they imply an Archimedean pricing equation that provides a rational equilibrium foundation for the low-volatility anomaly. Finally, we conduct an empirical study using U.S. ETF equity data to assess the performance of ALP and ACM strategies relative to standard benchmarks.

The remainder of this paper is organized as follows. Section 2 reviews the relevant literature on portfolio optimization, risk parity, and mechanically inspired financial models. Section 3 presents our methodology, including the mean–variance setup, the derivation of ALP and ACM portfolios, and dynamic trading rules. Section 4 describes the empirical design, data sources, and performance metrics. Section 5 discusses the empirical results. Section 6 concludes. Technical details and additional results are provided in the Appendix.

2 Literature Review

The modern theory of portfolio selection originates with Markowitz (1952), who formulated the problem of choosing asset weights to maximize expected return for a given level of variance (or equivalently, minimize variance for a given expected return). The resulting efficient frontier in mean–variance space is a parabola that maximizes the Sharpe ratio (Sharpe, 1964). The tangency portfolio has a closed form solution, but is also sensitive to estimation error in the vector of expected returns and their covariance (Michaud, 1989; Best and Grauer, 1991). This has motivated a large literature on robust portfolio optimization, shrinkage estimators, and Bayesian methods (Jorion, 1986; Ledoit and Wolf, 2003; DeMiguel et al., 2009).

In response to the instability of mean–variance optimization, practitioners have developed alternative weighting schemes. The equal-weight (1/N) portfolio is the simplest example (DeMiguel et al., 2009). Risk parity portfolios allocate capital so that each asset contributes equally to total portfolio risk (Maillard et al., 2010; Roncalli, 2013). Risk parity has also been shown to deliver attractive risk-adjusted returns, particularly in low-return environments (Asness et al., 2012). Minimum-variance portfolios have also received considerable attention as low-turnover, defensive strategies (Clarke et al., 2006; Scherer, 2011). All these benchmarks are very well known.

Factor models decompose asset returns into systematic (factor) and idiosyncratic components (Ross, 1976; Fama and French, 1993). Factor-based portfolio construction often involves tilting toward factors with positive risk premia while controlling for unwanted exposures (Grinold and Kahn, 2000). Our ACM portfolio construction combines factor models with minimum-variance optimization by imposing a constraint on the portfolio’s factor center of mass. This approach is related to, but distinct from, factor risk parity and targeted factor exposure strategies (Clarke et al., 2012).

The use of physical analogies in finance has a long history. Osborne (1959) compared stock price movements to Brownian motion. Mantegna and Stanley (1999) and subsequent work in econophysics have applied statistical mechanics concepts to financial markets. More recently, Fernholz (2002) developed stochastic portfolio theory using concepts from thermodynamics. Goldberg and Mahmoud (2008) proposed a “lever” interpretation of portfolio leverage but did not formalize it as an optimization principle. Our work extends this tradition by taking specific mechanical principles from Archimedes; the lever law and center of mass computation, and showing that they translate into rigorous mathematical constraints that refine and robustify existing portfolio solutions: the ALP is the classical tangency portfolio re derived under an implicit volatility shrinkage prior, and the ACM is the constrained minimum-variance portfolio re-derived under an explicit APT alignment condition. In doing so, we provide (i) a new closed-form characterization of shrinkage-robust tangency within mean variance theory, (ii) a factor-based APT constraint that turns constrained minimum-variance optimization into an equilibrium pricing instrument, and (iii) a dynamic rebalancing rule with a mechanical inter-

pretation. To our knowledge, this is the first paper to establish these connections and to unify shrinkage, factor alignment, and mechanical equilibrium within a single portfolio framework.

Dynamic portfolio strategies must address when and how to rebalance. Threshold rebalancing, where trades occur when weights deviate from targets by more than a specified amount, is widely used in practice (Plaxco and Arnott, 2002; Donohue and Yip, 2003). Volatility-based rebalancing has been studied in the context of risk parity and target volatility strategies (Kirby and Ost diek, 2012). Our torque-band and center-of-mass tracking rules provide a mechanical interpretation of threshold rebalancing: we rebalance when the portfolio’s mechanical equilibrium (measured by torque alignment or center-of-mass position) is sufficiently disturbed. This connects the static Archimedean constructions to implementable dynamic strategies. The Archimedean framework also relates to the broader literature on heuristic and behaviorally motivated portfolio rules. DeMiguel et al. (2009) document that naive equal-weighting persists in practice partly due to investor loss aversion and ambiguity aversion (Markowitz, 1952), which lead investors to avoid fragile optimized solutions in favour of simple, robust allocations. The ALP can be viewed as a bridge between such heuristics and formal optimality: it retains a closed-form structure while incorporating expected return information in a manner that is robust to estimation noise, consistent with the ambiguity-robust portfolio literature (Jorion, 1986). More recently, machine learning methods have been applied to portfolio construction, either to improve expected return forecasts (Gu et al., 2020) or to directly optimize portfolio weights in a data-driven fashion. While Archimedean portfolios do not use machine learning, their robustness to estimation error makes them natural benchmarks against which ML-based strategies should be evaluated; indeed, DeMiguel et al. (2009) show that a key challenge for any sophisticated method is to outperform a simple rule out of sample.

3 Methodology

3.1 Mean–Variance Setup and Classical Benchmarks

We adopt a standard one-period mean–variance setting. There are n risky assets with excess return vector $\mu \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$, assumed symmetric positive definite. Portfolio weights are $w \in \mathbb{R}^n$ satisfying $\mathbf{1}^\top w = 1$, where $\mathbf{1}$ denotes the vector of ones. The portfolio’s expected excess return and variance are

$$\mu_p(w) = \mu^\top w, \quad \sigma_p^2(w) = w^\top \Sigma w \quad (1)$$

with volatility $\sigma_p(w) = \sqrt{w^\top \Sigma w}$. The classical Sharpe ratio is

$$SR(w) = \frac{\mu^\top w}{\sqrt{w^\top \Sigma w}} \quad (2)$$

Maximizing $SR(w)$ under the budget constraint yields the well-known tangency portfolio

$$w^{Tan} = \frac{\Sigma^{-1} \mu}{\mathbf{1}^\top \Sigma^{-1} \mu} \quad (3)$$

which we take as benchmark. Our goal is not to re-derive w^{Tan} but to propose alternative portfolios whose first-order conditions embody Archimedean mechanics.

We will also use the classical notion of risk contribution. Let $\sigma_p(w) > 0$. The marginal contribution of asset i to volatility is

$$\frac{\partial \sigma_p(w)}{\partial w_i} = \frac{(\Sigma w)_i}{\sigma_p(w)} \quad (4)$$

and the risk contribution of asset i is defined as

$$RC_i(w) = w_i \frac{(\Sigma w)_i}{\sigma_p(w)} \quad (5)$$

These risk contributions sum to total volatility: $\sum_{i=1}^n RC_i(w) = \sigma_p(w)$.

3.2 Archimedean Lever Portfolios (ALP)

3.2.1 Forces, Distances, and Torques in Portfolio Space

Archimedes' lever law states that a rigid bar with forces F_i applied at signed distances d_i from a fulcrum is in equilibrium when total torque vanishes:

$$\sum_i F_i d_i = 0 \quad (6)$$

More generally, one can consider balanced configurations where torques on different arms are proportional to given weights. We translate this picture into portfolio space as follows. For each asset i we define:

- a distance $d_i := \mu_i$, the asset's expected excess return;
- a force $F_i := w_i \sigma_i$, where $\sigma_i := \sqrt{\Sigma_{ii}}$ is the asset's volatility.

The torque of asset i is

$$\tau_i(w) := F_i d_i = w_i \sigma_i \mu_i \quad (7)$$

and the total torque of the portfolio is $T(w) = \sum_{i=1}^n \tau_i(w)$. Mechanical equilibrium for a lever typically involves balancing torques around a fulcrum. In the portfolio setting it is natural to compare torques to risk contributions, which already decompose total portfolio volatility. This leads to the following equilibrium condition.

Definition 3.1 (Archimedean lever equilibrium). *A portfolio w with $\sigma_p(w) > 0$ is in Archimedean lever equilibrium if there exists a scalar k such that*

$$\tau_i(w) = k RC_i(w) \quad \text{for all } i \text{ with } w_i > 0 \quad (8)$$

That is, each asset's torque is proportional to its risk contribution. Intuitively, assets that pull strongly on the lever (high torque) must also bear a proportionally large share of risk. This couples the mechanical and probabilistic notions of load.

3.2.2 Closed-Form Characterization

We now derive the explicit form of portfolios in Archimedean lever equilibrium.

Theorem 3.2 (Closed form for the Archimedean Lever Portfolio). *Suppose Σ is symmetric positive definite, and let $\sigma = (\sigma_1, \dots, \sigma_n)^\top$ with $\sigma_i = \sqrt{\Sigma_{ii}}$. Let w satisfy the budget constraint $\mathbf{1}^\top w = 1$ and the Archimedean lever equilibrium condition*

$$w_i \sigma_i \mu_i = k w_i \frac{(\Sigma w)_i}{\sigma_p(w)} \quad \text{for all } i \text{ with } w_i > 0, \quad (9)$$

for some scalar k . Then w is proportional to

$$\Sigma^{-1}(\text{diag}(\sigma)\mu), \quad (10)$$

i.e.,

$$w^{ALP} = \frac{\Sigma^{-1} \text{diag}(\sigma)\mu}{\mathbf{1}^\top \Sigma^{-1} \text{diag}(\sigma)\mu} \quad (11)$$

satisfies the budget constraint and the equilibrium condition. When $\text{diag}(\sigma)$ is not proportional to the identity, w^{ALP} differs from the tangency portfolio w^{Tan} .

Proof. Let w satisfy the equilibrium condition with $\sigma_p = \sigma_p(w) > 0$ and $w_i > 0$. Then for each such i ,

$$w_i \sigma_i \mu_i = k w_i \frac{(\Sigma w)_i}{\sigma_p} \Rightarrow \sigma_i \mu_i = \frac{k}{\sigma_p} (\Sigma w)_i \quad (12)$$

Define $b := \text{diag}(\sigma)\mu$. In vector form the above reads

$$b = \frac{k}{\sigma_p} \Sigma w \Rightarrow \Sigma w = c b \quad (13)$$

with $c := \sigma_p/k \neq 0$. Since Σ is invertible,

$$w = c \Sigma^{-1} b = c \Sigma^{-1} \text{diag}(\sigma)\mu \quad (14)$$

Imposing the budget constraint $\mathbf{1}^\top w = 1$ yields

$$1 = \mathbf{1}^\top w = c \mathbf{1}^\top \Sigma^{-1} \text{diag}(\sigma)\mu \Rightarrow c = \frac{1}{\mathbf{1}^\top \Sigma^{-1} \text{diag}(\sigma)\mu} \quad (15)$$

Thus

$$w^{ALP} = \frac{\Sigma^{-1} \text{diag}(\sigma)\mu}{\mathbf{1}^\top \Sigma^{-1} \text{diag}(\sigma)\mu} \quad (16)$$

w^{ALP} above is the unconstrained closed-form solution; our empirical implementation applies a long-only position cap described in Appendix A(e).¹

If $\text{diag}(\sigma)$ is proportional to the identity, $\Sigma^{-1} \text{diag}(\sigma)\mu$ is proportional to $\Sigma^{-1}\mu$ and ALP coincides with the tangency portfolio. The ALP shares the same structural form as the tangency portfolio, both are $\Sigma^{-1}(\cdot)$ normalized by a budget scalar, but with the expected return vector μ replaced by the volatility-scaled signal $\text{diag}(\sigma)\mu$. This substitution is the precise sense in which the ALP is a robustified tangency: when $\text{diag}(\sigma)$ is proportional to the identity (i.e., all assets have equal volatility), the two portfolios coincide exactly; in the empirically relevant case of heterogeneous volatilities, the ALP downweights assets whose high sample mean is not matched by commensurate individual risk, implementing implicit Bayes-Stein shrinkage of expected return estimates within the classical $\Sigma^{-1}\mu$ formula.

□

Equation (11) is the classical tangency solution with the signal μ replaced by $\text{diag}(\sigma)\mu$; novelty resides in the equilibrium condition motivating this substitution (Def. 3.1), not in the optimization technique itself.

Proposition 3.3 (ALP as Torque-Sharpe Maximizer). *The Archimedean Lever Portfolio w^{ALP} is the unique solution to the problem of maximizing the torque-to-volatility ratio*

$$\max_{w: \mathbf{1}^\top w = 1} \frac{T(w)}{\sigma_p(w)} = \frac{w^\top \text{diag}(\sigma)\mu}{\sqrt{w^\top \Sigma w}} \quad (17)$$

This provides a decision-theoretic foundation for the ALP: it is the portfolio chosen by an investor who evaluates each asset by its volatility-weighted return signal $\sigma_i\mu_i$ rather than by raw expected return μ_i alone.

Proof. Define $\tilde{\mu} := \text{diag}(\sigma)\mu$, so that $T(w) = w^\top \tilde{\mu}$. The ratio in (17) equals the generalized Sharpe ratio $w^\top \tilde{\mu} / \sqrt{w^\top \Sigma w}$. By the standard mean–variance argument, the unique maximizer of this ratio subject to $\mathbf{1}^\top w = 1$ is

$$w^* = \frac{\Sigma^{-1} \tilde{\mu}}{\mathbf{1}^\top \Sigma^{-1} \tilde{\mu}} = \frac{\Sigma^{-1} \text{diag}(\sigma)\mu}{\mathbf{1}^\top \Sigma^{-1} \text{diag}(\sigma)\mu} = w^{ALP} \quad (18)$$

which is precisely the portfolio of Theorem 3.2.

□

Remark 3.4 (Implicit Shrinkage and Robustness to Estimation Error). *The signal substitution $\mu_i \mapsto \sigma_i\mu_i$ embedded in the ALP has a natural robust-estimation interpretation. In the classical tangency portfolio, assets with small estimated volatility $\hat{\sigma}_i$ but high sample mean $\hat{\mu}_i$ can receive extreme positions because $\Sigma^{-1}\hat{\mu}$ amplifies noisy return estimates—the “error-maximizing” property documented by Best and Grauer (1991) and Michaud (1989). The ALP replaces $\hat{\mu}_i$ with*

¹See Appendix A(e) for the exact caps used per universe.

$\hat{\sigma}_i \hat{\mu}_i$, deflating the influence of assets whose high apparent mean is not accompanied by commensurate individual risk. This is consistent with the empirical-Bayes shrinkage logic of Jorion (1986): under a prior in which Sharpe ratios are drawn from a common distribution, the posterior expected return takes the form $E[\mu_i \mid \text{data}] \propto \sigma_i \times (\text{shrunk } SR_i)$, exactly matching the ALP signal. The practical consequence—lower turnover and smaller drawdowns relative to the tangency portfolio—is confirmed across all five configurations in Section 5.

Remark 3.5 (Relation to Risk Parity and Maximum Diversification). *The ALP unifies risk parity and maximum diversification within a single mean variance framework, nesting each as a degenerate special case and interpolating between them according to the cross-sectional dispersion of Sharpe ratios. Risk parity is the ALP evaluated when all Sharpe ratios μ_i/σ_i are equal (so the volatility-scaled signal is flat across assets); maximum diversification is the ALP when all expected returns μ_i are equal (so $\sigma_i \mu_i \propto \sigma_i$). The ALP is therefore not a competitor of these strategies but their common generalization within the mean variance framework, a return weighted maximum diversification portfolio that reduces to pure diversification when return differences are uninformative and to risk parity when Sharpe ratio differences are uninformative.*

Risk parity requires each asset to contribute equally to total portfolio volatility:

$$w_i \frac{(\Sigma w)_i}{\sigma_p(w)} = \frac{\sigma_p(w)}{n} \quad \text{for all } i \quad (19)$$

The ALP lever equilibrium condition $\tau_i = k RC_i$ requires $w_i \sigma_i \mu_i = k w_i (\Sigma w)_i / \sigma_p(w)$, which simplifies to $\sigma_i \mu_i = k (\Sigma w)_i / \sigma_p(w)$. These two conditions coincide if and only if $\sigma_i \mu_i = \text{const}$, i.e., all assets share the same Sharpe ratio μ_i/σ_i . In the heterogeneous-Sharpe-ratio environments typical of equity markets, ALP and risk parity diverge: the ALP tilts toward assets with higher volatility-weighted expected returns.

Maximum diversification (Clarke et al., 2012) maximizes the ratio of weighted average volatility to portfolio volatility:

$$\max_{w: \mathbf{1}^\top w = 1} \frac{\sum_i w_i \sigma_i}{\sqrt{w^\top \Sigma w}} \quad (20)$$

The ALP instead maximizes the torque-to-volatility ratio $\sum_i w_i \sigma_i \mu_i / \sqrt{w^\top \Sigma w}$ (Proposition 3.3), which nests maximum diversification as the special case $\mu_i = \text{const}$, i.e., when all assets have equal expected returns. The ALP can therefore be understood as a return-weighted maximum diversification portfolio: it generalizes diversification by replacing each asset's volatility σ_i with its volatility-scaled return signal $\sigma_i \mu_i$.

3.3 Archimedean Center-of-Mass Portfolios (ACM)

3.3.1 Factor Space and Center of Mass

In Ross' arbitrage pricing theory (1976), the expected return on any portfolio is fully determined by its exposure to systematic factors: $\mu_p = \mu_f^\top (B^\top w)$. An investor who accepts APT and targets

a specific expected return $\bar{\mu}$ must therefore align the portfolio's factor exposure vector $B^\top w$ with the factor risk premium vector μ_f . We call this alignment condition the Archimedean center-of-mass constraint, the portfolio's factor-weighted center of mass must coincide with the direction of the factor risk premium, and derive the minimum-variance portfolio that satisfies it jointly with the budget constraint. This is the Archimedean Center-of-Mass (ACM) portfolio: a constrained minimum variance / factor aligned solution that implements APT pricing at the portfolio level, related to factor risk parity and targeted factor-exposure strategies (Clarke et al., 2012; Grinold and Kahn, 2000) but with a specific, APT motivated choice of target direction.

$$R = Bf + \varepsilon \quad (21)$$

where $R \in \mathbb{R}^n$ is the vector of asset excess returns, $B \in \mathbb{R}^{n \times k}$ is the matrix of factor loadings, $f \in \mathbb{R}^k$ is the factor return vector, and ε is idiosyncratic noise. Assume $E[f] = \mu_f$ and that factors and idiosyncratic components are uncorrelated. Each asset i is represented by its factor loading vector $b_i^\top$ (the i -th row of B) in $\mathbb{R}^k$. With portfolio weights w , the exposure factor of the portfolio is

$$\beta_p(w) = B^\top w \in \mathbb{R}^k \quad (22)$$

This is precisely the center of mass of the portfolio in factor space if we interpret w_i as masses.

Definition 3.6 (Factor center of mass). *The factor center of mass of a portfolio w is*

$$c(w) := B^\top w \quad (23)$$

We now impose an Archimedean condition: the center of mass in factor space should align with the factor risk premium.

Definition 3.7 (Archimedean center-of-mass alignment). *Given factor risk premium $\mu_f \in \mathbb{R}^k$ and a scalar $\gamma > 0$, a portfolio w is Archimedean center-of-mass aligned if*

$$B^\top w = \gamma \mu_f \quad (24)$$

The parameter γ scales the overall exposure to the factor risk premium. We select w as the minimum-variance portfolio among all such aligned portfolios satisfying the budget constraint.

3.3.2 Closed-Form Solution under Full Rank

Assume that Σ is positive definite and B has full column rank k . Fix $\gamma > 0$. We solve

$$\begin{aligned} \min_{w \in \mathbb{R}^n} \quad & w^\top \Sigma w \\ \text{s.t.} \quad & \mathbf{1}^\top w = 1, \\ & B^\top w = \gamma \mu_f \end{aligned} \quad (25)$$

Theorem 3.8 (Closed form for Archimedean Center-of-Mass Portfolio). *Assume Σ is invertible and B has full column rank k . Let $\gamma > 0$ be given. Define*

$$A_{11} := \mathbf{1}^\top \Sigma^{-1} \mathbf{1} \in \mathbb{R}, \quad A_{12} := \mathbf{1}^\top \Sigma^{-1} B \in \mathbb{R}^{1 \times k} \quad (26)$$

$$A_{21} := B^\top \Sigma^{-1} \mathbf{1} \in \mathbb{R}^{k \times 1} \quad A_{22} := B^\top \Sigma^{-1} B \in \mathbb{R}^{k \times k} \quad (27)$$

and the Schur complement

$$S := A_{22} - \frac{1}{A_{11}} A_{21} A_{12} \quad (28)$$

If S is invertible, then the solution to (25) is

$$w^{ACM}(\gamma) = \frac{1}{2} \Sigma^{-1} (\lambda^*(\gamma) \mathbf{1} + B \eta^*(\gamma)) \quad (29)$$

where

$$\eta^*(\gamma) = S^{-1} \left(2\gamma \mu_f - \frac{2}{A_{11}} A_{21} \right) \quad (30)$$

$$\lambda^*(\gamma) = \frac{2 - A_{12} \eta^*(\gamma)}{A_{11}} \quad (31)$$

Proof. Form the Lagrangian

$$\mathcal{L}(w, \lambda, \eta) = w^\top \Sigma w - \lambda (\mathbf{1}^\top w - 1) - \eta^\top (B^\top w - \gamma \mu_f) \quad (32)$$

with scalar multiplier λ for the budget constraint and vector multiplier $\eta \in \mathbb{R}^k$ for the alignment constraint. First-order conditions in w give

$$2\Sigma w - \lambda \mathbf{1} - B\eta = 0 \Rightarrow w = \frac{1}{2} \Sigma^{-1} (\lambda \mathbf{1} + B\eta) \quad (33)$$

Imposing $\mathbf{1}^\top w = 1$ and $B^\top w = \gamma \mu_f$ yields

$$\frac{1}{2} (\lambda A_{11} + A_{12} \eta) = 1 \quad (34)$$

$$\frac{1}{2} (\lambda A_{21} + A_{22} \eta) = \gamma \mu_f \quad (35)$$

Multiplying by 2,

$$\lambda A_{11} + A_{12} \eta = 2, \quad \lambda A_{21} + A_{22} \eta = 2\gamma \mu_f \quad (36)$$

From the first equation,

$$\lambda = \frac{2 - A_{12} \eta}{A_{11}} \quad (37)$$

Substituting into the second,

$$\frac{2 - A_{12} \eta}{A_{11}} A_{21} + A_{22} \eta = 2\gamma \mu_f \quad (38)$$

which can be rearranged to

$$\left(A_{22} - \frac{1}{A_{11}}A_{21}A_{12}\right)\eta = 2\gamma\mu_f - \frac{2}{A_{11}}A_{21} \quad (39)$$

If S is invertible,

$$\eta^*(\gamma) = S^{-1}\left(2\gamma\mu_f - \frac{2}{A_{11}}A_{21}\right) \quad (40)$$

and

$$\lambda^*(\gamma) = \frac{2 - A_{12}\eta^*(\gamma)}{A_{11}} \quad (41)$$

Substituting into $w = \frac{1}{2}\Sigma^{-1}(\lambda\mathbf{1} + B\eta)$ gives the stated solution. $\square$

Proposition 3.9 (ACM as a Portfolio-Level APT Pricing Condition). *Consider an investor who: (i) accepts the Arbitrage Pricing Theory (Ross, 1976) at the portfolio level, so that the portfolio's systematic expected return is fully determined by its factor exposures, $\mu_p(w) = \mu_f^\top(B^\top w)$; (ii) targets a specific expected return $\bar{\mu} = \gamma\|\mu_f\|^2$ by choosing $\gamma > 0$; and (iii) minimizes portfolio variance to eliminate non-rewarded idiosyncratic risk, subject to the budget constraint $\mathbf{1}^\top w = 1$. The unique portfolio satisfying (i)–(iii) is the Archimedean Center-of-Mass portfolio $w^{ACM}(\gamma)$ of Theorem 3.8. The scalar γ represents the investor's factor-exposure intensity: how aggressively they align with factor-rewarded risk relative to the minimum-variance baseline. These three conditions are precisely the investor-level analogue of the assumptions underlying Ross's (1976) APT: factor-determined expected returns, a targeted factor-risk-premium exposure, and elimination of unrewarded idiosyncratic risk.*

Proof. Under (i), the portfolio expected return is $\mu_p(w) = \mu_f^\top(B^\top w)$. Targeting $\mu_p = \gamma\|\mu_f\|^2$ under (ii) requires $\mu_f^\top(B^\top w) = \gamma\|\mu_f\|^2$, which is satisfied whenever $B^\top w = \gamma\mu_f$ since then $\mu_f^\top(B^\top w) = \gamma\mu_f^\top\mu_f = \gamma\|\mu_f\|^2$. An investor satisfying (ii)–(iii) therefore solves

$$\min_w w^\top \Sigma w \quad \text{s.t.} \quad \mathbf{1}^\top w = 1, \quad B^\top w = \gamma\mu_f \quad (42)$$

which is exactly problem (25), whose solution is $w^{ACM}(\gamma)$ by Theorem 3.8. $\square$

3.4 Archimedean Equilibrium Pricing: A representative-investor derivation

The propositions above ground ALP and ACM in single-investor preferences. We now examine the aggregate implication: what asset-pricing equation arises when all investors hold the ALP in equilibrium?

We now ask the classical equilibrium question (Sharpe, 1964; Merton, 1972): if a representative investor optimally holds the ALP, as justified by the torque-Sharpe maximization in Proposition 3.3, and if markets clear so that aggregate demand equals the market portfolio, what pricing relation must hold in equilibrium? We proceed under the usual homogeneous-expectations assumption: all investors share the same estimates of μ and Σ , and therefore arrive at the same optimal portfolio.

Proposition 3.10 (Archimedean Pricing Equation). *Suppose all investors hold the ALP, so that the market portfolio w^m satisfies $w^m \propto \Sigma^{-1} \text{diag}(\sigma) \mu$ in equilibrium. By Proposition 3.3, this is the portfolio chosen by any mean-variance investor evaluating assets via the torque signal $\sigma_i \mu_i$; equilibrium therefore requires only that this evaluation rule be common across investors, exactly as in the classical CAPM derivation. Then expected excess returns satisfy the Archimedean pricing equation*

$$\mu_i = \frac{c}{\sigma_i} \text{Cov}(r_i, r_m) \quad i = 1, \dots, n \quad (43)$$

for a positive scalar c , where $r_m := (w^m)^\top r$ is the market return. Defining the Archimedean beta $\beta_i^{\text{Arch}} := \text{Cov}(r_i, r_m) / (\sigma_i \sigma_m^2)$, equation (43) is equivalent to

$$\mu_i = c \sigma_m^2 \sigma_i \beta_i^{\text{Arch}} \quad (44)$$

In the homoskedastic case $\sigma_i = \sigma$ for all i , the standard CAPM is recovered up to a positive constant.

Proof. If $w^m \propto \Sigma^{-1} \text{diag}(\sigma) \mu$, there exists $c > 0$ such that $\text{diag}(\sigma) \mu = c \Sigma w^m$. The i -th component reads $\sigma_i \mu_i = c (\Sigma w^m)_i = c \text{Cov}(r_i, r_m)$, giving (43) directly. Dividing both sides by $\sigma_i \sigma_m^2$ and using the definition of β_i^{Arch} yields (44). Equation (44) is derived by the identical market-clearing argument used to obtain the classical CAPM beta pricing relation — the only modification is the asset-specific volatility deflation $1/\sigma_i$ inherited from the ALP's torque signal. When $\sigma_i = \sigma$ for all i , (43) reduces to $\mu_i = (c/\sigma) \text{Cov}(r_i, r_m)$, the standard CAPM with market price of risk $\lambda = c \sigma_m^2 / \sigma$. $\square$

The two pricing results above are special cases of a single equilibrium logic. Proposition 3.10 recovers CAPM as the one-factor case of representative investor equilibrium; Proposition 3.9 recovers an APT consistent pricing condition as its multi factor analogue. The Archimedean framework therefore does not compete with CAPM/APT but supplies a specific, testable refinement, the volatility deflation in (43), within that same equilibrium class.

Remark 3.11 (Connection to the Low-Volatility Anomaly and Relation to CAPM). *The Archimedean pricing equation (43) departs from the standard CAPM in a single but consequential way: the expected return of asset i is proportional to $\text{Cov}(r_i, r_m) / \sigma_i$ rather than to $\text{Cov}(r_i, r_m) / \sigma_m^2$. This volatility deflation implies that, relative to CAPM predictions, high-volatility assets earn a lower risk premium per unit of market covariance. Equivalently, low-volatility assets are relatively over-rewarded in Archimedean equilibrium. This is precisely the content of the empirically well-documented low-volatility anomaly (Clarke et al., 2006): low-volatility stocks deliver risk-adjusted returns that exceed standard CAPM predictions, while high-volatility stocks underperform. Equation (43) thus provides a rational equilibrium rationalization of that anomaly as the natural consequence of investors following torque-equilibrium portfolio rules, without requiring behavioral biases or market frictions.*

3.5 Dynamic Archimedean Trading Rules

The dynamic rules below are presented as natural mechanical extensions of the static ALP/ACM constructions of Sections 3.2–3.3. The empirical results in Section 5 use unconditional monthly rebalancing to the static closed-form weights of Theorems 3.2 and 3.8 at a fixed $\gamma = 1$; the torque-band and center-of-mass tracking rules below, and the adaptive- γ scheme of Appendix B, are proposed refinements whose empirical evaluation is left to future work.

The ALP and ACM constructions provide static allocations. To implement them as trading strategies, we need dynamic rules that respond to changing estimates of μ , Σ , and B . Archimedean mechanics suggests a natural approach: treat deviations from equilibrium as imbalances in torque or in the location of the center of mass and rebalance when these imbalances exceed thresholds.

3.5.1 Torque-Based Rebalancing Bands

At time t , estimate $\hat{\mu}_t$, $\hat{\Sigma}_t$, and $\hat{\sigma}_{i,t} = \sqrt{\hat{\Sigma}_{ii,t}}$. Compute the ALP weights

$$w_t^{ALP} = \frac{\hat{\Sigma}_t^{-1} \text{diag}(\hat{\sigma}_t) \hat{\mu}_t}{\mathbf{1}^\top \hat{\Sigma}_t^{-1} \text{diag}(\hat{\sigma}_t) \hat{\mu}_t} \quad (45)$$

Define normalized torque contributions

$$\tilde{\tau}_{i,t}^{target} = \frac{w_{i,t}^{ALP} \hat{\sigma}_{i,t} \hat{\mu}_{i,t}}{\sum_{j=1}^n w_{j,t}^{ALP} \hat{\sigma}_{j,t} \hat{\mu}_{j,t}} \quad (46)$$

Let $w_{i,t}^{real}$ be the realized weight of asset i at time t , and define

$$\tilde{\tau}_{i,t}^{real} = \frac{w_{i,t}^{real} \hat{\sigma}_{i,t} \hat{\mu}_{i,t}}{\sum_{j=1}^n w_{j,t}^{real} \hat{\sigma}_{j,t} \hat{\mu}_{j,t}} \quad (47)$$

A simple measure of torque misalignment is

$$D_t = \sum_{i=1}^n (\tilde{\tau}_{i,t}^{real} - \tilde{\tau}_{i,t}^{target})^2 \quad (48)$$

An Archimedean torque-band trading rule rebalances at time t if D_t exceeds a pre-specified threshold $\theta > 0$. Thus the lever is allowed to deviate slightly from equilibrium, but when torque imbalance becomes large, the strategy “resets the fulcrum” by trading towards w_t^{ALP} .

Remark 3.12 (Distinction from Standard Tracking Error). *The torque misalignment measure D_t is not equivalent to standard tracking error, even though both involve a sum of squared deviations. Standard threshold rebalancing triggers a trade when*

$$TE_t = \sum_{i=1}^n (w_{i,t}^{real} - w_{i,t}^{ALP})^2 \quad (49)$$

exceeds a threshold—a measure in weight space. The torque misalignment D_t , by contrast, measures deviation in normalized torque space:

$$D_t = \sum_{i=1}^n \left(\frac{w_{i,t}^{real} \hat{\sigma}_{i,t} \hat{\mu}_{i,t}}{\sum_j w_{j,t}^{real} \hat{\sigma}_{j,t} \hat{\mu}_{j,t}} - \frac{w_{i,t}^{ALP} \hat{\sigma}_{i,t} \hat{\mu}_{i,t}}{\sum_j w_{j,t}^{ALP} \hat{\sigma}_{j,t} \hat{\mu}_{j,t}} \right)^2 \quad (50)$$

Because weights are transformed non-linearly by the asset-specific signal $\hat{\sigma}_i \hat{\mu}_i$, D_t is not a monotone function of TE_t unless $\hat{\sigma}_i \hat{\mu}_i$ is constant across assets. In particular, D_t is more sensitive to misalignment in high-torque (high Sharpe, high volatility) assets and less sensitive to misalignment in low-torque assets, whereas TE_t weights all assets equally. This asymmetry is economically meaningful: it implies the strategy rebalances more urgently when the deviation is concentrated in the assets that most strongly influence the mechanical equilibrium of the lever, rather than when any asset drifts from its target weight.

3.5.2 Center-of-Mass Tracking

Similarly, for ACM portfolios we can define a center-of-mass misalignment. At time t , estimate $\hat{\Sigma}_t$ and $\hat{B}_t$, and choose γ_t according to a risk policy. Compute $w_t^{ACM}(\gamma_t)$ from Theorem 3.8 and set

$$c_t^{target} = \hat{B}_t^\top w_t^{ACM}(\gamma_t), \quad c_t^{real} = \hat{B}_t^\top w_t^{real} \quad (51)$$

An alignment error is

$$E_t = \|c_t^{real} - c_t^{target}\|^2 \quad (52)$$

where $\|\cdot\|$ is the Euclidean norm. If E_t exceeds a threshold $\phi > 0$, the strategy rebalances towards $w_t^{ACM}(\gamma_t)$.

4 Data and Empirical Design

Our primary dataset consists of twelve U.S. equity ETFs over the period January 2000 to December 2025: the nine SPDR sector ETFs (XLK, XLF, XLV, XLE, XLI, XLY, XLP, XLB, XLU), two sector ETFs introduced after the original nine-sector SPDR suite (XLC – Communication Services, XLRE – Real Estate), and SPY itself, included both as a traded constituent and, separately, as the level against which we would compare a passive buy-and-hold benchmark. We note that SPY’s inclusion as a constituent alongside sector ETFs that approximately span its own return introduces material collinearity in Σ ; this is handled numerically by the Ledoit–Wolf shrinkage of Appendix A(a) but should be read as a limitation of the traded universe as implemented, not a feature of the theoretical construction. We obtain daily, weekly and monthly total returns from the Yahoo Finance and risk-free rates and the Fama-French factors from Kenneth French’s data library. For each strategy, we form portfolios at the end of each month using data from the prior months using a rolling window approach.

We compare the ALP and ACM strategies to four benchmarks computed within the same walk-forward engine: equal-weight (EW): $w_i = 1/n$; minimum-variance (MV) $w^{MV} = \Sigma^{-1}\mathbf{1}/(\mathbf{1}^\top \Sigma^{-1}\mathbf{1})$; tangency (Tan) $w^{Tan} = \Sigma^{-1}\mu/(\mathbf{1}^\top \Sigma^{-1}\mu)$; and a fixed 70/30 concentrated benchmark that allocates 70% to the first asset in the universe and splits the remaining 30% equally across the rest. Risk parity weights are computed as $w_i \propto 1/\sigma_i$ normalized to satisfy the budget constraint, and the passive benchmark line reports a buy-and-hold position in SPY, both re-estimated at the same monthly rebalancing dates as the other strategies. For the ACM strategy, we extract $k = 2$ statistical factors via principal component analysis of the trailing correlation matrix of asset returns, re-estimated at each rebalancing date within the training window; this replaces a named macroeconomic factor model (e.g., Fama–French) with a purely data-driven factor structure that requires no external factor data and adapts to the realized cross-section of the traded universe. Factor risk premia μ_f are obtained via least-squares projection of the trailing mean asset returns onto the estimated loadings B . The alignment parameter γ is held fixed at $\gamma = 1$ throughout the reported backtest; the adaptive- γ scheme of Appendix B is a proposed extension not used to generate the results below (see Appendix B’s note).

We evaluate strategies using standard evaluation metrics: annualized return and volatility, annualized Sharpe ratio, Sortino ratio, Calmar ratio, maximum drawdown, and cumulative wealth.

5 Empirical Results

We now present our empirical findings from our comprehensive backtest of the Archimedean Lever Portfolio (ALP) and Archimedean Center-of-Mass (ACM) strategy against the benchmarks. Results are organized along two axes: evaluation period (the walk-forward out-of-sample test, and the COVID-19 subperiod) and data frequency (daily, weekly, and monthly returns), yielding the configurations summarized in Table 1 for our primary universe of nine S&P 500 sector ETFs. Tables 2–4 report the core performance metrics across all five configurations. For each strategy we report the annualized return and volatility (in percent), Sharpe and Sortino ratios, maximum drawdown, cumulative terminal wealth per \$1 invested, monthly turnover, and the annualized alpha from a Fama–French factor regression.

5.1 Key Findings Across Periods and Frequencies

Across all five configurations the ALP delivers either the highest or second-highest Sharpe ratio among active strategies. The improvement over the market benchmark is most pronounced in the long sample: in the 2000–2025 daily setting the ALP Sharpe ratio of 0.53 exceeds the benchmark’s 0.40 by 13 basis points, while simultaneously generating a 24% higher annualized return (10.75% vs. 8.63%) with lower volatility than the benchmark (18.35% vs. 19.17%). In the most recent 2020–2025 window the ALP is the only active strategy to match and—under monthly data—slightly exceed the SPY benchmark’s Sharpe ratio (1.00 vs. 0.92). A central theoretic-

cal motivation for Archimedean strategies is their lower dependence on precise expected-return estimation compared to the Tangency portfolio. This translates directly into lower rebalancing activity: in the 2010–2025 monthly setting, the ALP monthly turnover of 8.0% compares favorably with the Tangency portfolio's 20.7% and, critically, is achieved without sacrificing return. The Sharpe ratio is generally higher under monthly rebalancing than under daily data. The ALP Sharpe ratio rises from 0.73 (daily) to 0.90 (monthly), and the ACM from 0.69 to 0.99. This pattern reflects the well-documented phenomenon that return autocorrelation and mean-reversion at short horizons create transaction-cost drag when using daily data with monthly rebalancing.

In the longer samples (2000–2025 and 2010–2025 daily), the ACM strategy closely tracks the Minimum Variance portfolio in terms of volatility (13.83% vs. 13.74%) and maximum drawdown (–34.46% in both cases), while producing marginally higher alphas (8.34% vs. 8.14% for 2000–2025). This near-equivalence in the daily frequency reflects a regime in which the factor center-of-mass constraint does not materially depart from the minimum-variance unconstrained solution—the nine-sector ETF universe has relatively homogeneous factor loadings. In the shorter 2020–2025 window, where cross-sectional dispersion was amplified by sector rotation during the post-pandemic recovery and rate-hiking cycle, the ACM drawdown is 16.60% (daily) vs. 16.68% for Min. Variance, again nearly tied, while delivering a modestly higher alpha. The full 25-year sample (Table 2, left panel) spans both the dot-com crash and the Global Financial Crisis (GFC), two episodes of severe equity drawdowns. The SPY benchmark experiences a maximum drawdown of –55.19%; the Tangency portfolio reaches –50.17%; and the Equal Weight strategy falls to –52.53%. By contrast, the ALP drawdown is only –43.71% and the ACM is –34.46%, matching Minimum Variance; the torque-equilibrium rebalancing rule provides genuine downside discipline over multi-decade investment horizons.

Table 5 consolidates the annualized Fama–French alpha estimates across all five configurations. Positive alphas indicate systematic outperformance above and beyond what is explained by exposure to the market, size, and value factors. The ALP generates the highest alpha among active strategies in four of the five configurations (12.16%, 14.44%, 16.28%, and 15.47%), with the single exception being the 2020-D case where the Tangency portfolio's concentrated growth-factor bet temporarily inflated its alpha to 25.54%. The ACM alpha closely mirrors Minimum Variance in all settings where the center-of-mass constraint is tightest ($\phi = 0.02$), consistent with the interpretation that tight factor alignment essentially recovers the minimum-variance solution in a homogeneous-universe setting. Under the looser threshold ($\phi = 0.10$, configurations 2020-D and 2020-M), the ACM alpha of 9.09% and 11.28% reflects the additional degrees of freedom afforded by the relaxed constraint.

5.2 Robustness Across Crisis Regimes: Russell and Multi-Asset-Class Universes

Because the primary S&P 500 sector-ETF universe was only evaluated over the COVID sub-period and the walk-forward OOS period, we assess crisis-period robustness using the Russell

2000/1000 and eight-asset-class universes described in Section 4, for which the dot-com (1999–2003), GFC (2007–2009), and full in-sample (2000–2015) periods are also available.

In the Russell universe, both strategies perform poorly during the dot-com period (ALP Sharpe -0.22 , ACM Sharpe -0.50 at daily frequency, 10 bps cost) and during the GFC (-0.38 and -0.37 respectively), with maximum drawdowns around -34% to -60% depending on frequency. ALP modestly outperforms ACM in most Russell configurations during these stress periods, though both strategies show negative risk-adjusted returns, indicating neither construction offers protection against systemic, universe-wide drawdowns of this magnitude.

In the eight-asset multi-class universe (equities, developed international, Treasuries, REITs, emerging markets, commodities, credit), ALP outperforms ACM more consistently during the GFC (Sharpe -0.02 vs. -0.39 at daily frequency, 10 bps) and full in-sample period (Sharpe 0.15 vs. -0.15), with ACM showing a wider drawdown (-45.1%) than ALP (-36.1%) in the full in-sample window. This suggests the ALP's volatility-scaled return signal is more robust to the addition of low-correlation asset classes (Treasuries, commodities) than the ACM's factor-alignment constraint, which appears more sensitive to the cross-sectional heterogeneity introduced by a multi-asset-class universe.

We emphasize these are stress-test universes distinct from the paper's primary S&P 500 sector-ETF results and are not directly comparable to Table 2 in absolute terms, given the different asset compositions, volatilities, and correlation structures involved.

6 Summary and Concluding Remarks

The empirical evidence across all five configurations supports three overarching conclusions. First, the Archimedean Lever Portfolio provides a robust improvement over both the passive benchmark and the classical Tangency portfolio. Its mechanical rebalancing rule controls turnover while capturing the signal embedded in the volatility-adjusted return vector. The result is a strategy with higher Sharpe ratio, lower maximum drawdown, and substantially lower transaction costs than tangency-based approaches, consistent with the theoretical arguments in Section 3. Second, When the constraint is tight and the asset universe is factor homogeneous, ACM and Min. Variance are nearly indistinguishable, a reassuring consistency result: it confirms that the APT alignment constraint does not distort the minimum-variance solution when factor premia are already well-reflected in the unconstrained optimum. The ACM departs from Min. Variance, and adds measurable alpha, precisely when cross-sectional factor loading dispersion is material, as in the 2020–2025 window during the post-pandemic sector rotation and rate-hiking cycle, which is exactly when an explicit APT alignment condition most benefits the investor. —both conditions that allow the Schur-complement system to deviate more substantially from the pure minimum-variance unconstrained solution. Third, data frequency is a first-order implementation choice, though the direction of the effect is not uniform: in the S&P 500 universe, weekly frequency shows the strongest Sharpe ratios of the three we tested

(ALP Sharpe 1.29 in the COVID period at weekly frequency, versus 0.40 daily and unavailable monthly for that period), while OOS Sharpe ratios decline monotonically from daily (0.38) to weekly (0.34) to monthly (0.23) rather than improving with lower frequency. Given this, we do not assert a single “natural” combination of estimation window and rebalancing frequency; the appropriate choice appears sensitive to the evaluation period and warrants further study rather than a general recommendation. Practitioners implementing Archimedean strategies with a one-year estimation window should therefore default to monthly or weekly data unless intraday liquidity concerns make more frequent rebalancing operationally necessary.

The conceptual appeal of Archimedean portfolios lies in their interpretability: they optimize the same mean variance objectives as classical theory, but under a constraint, torque/risk-contribution proportionality, or factor center of mass alignment, that has a direct mechanical reading and yields one empirically distinguishable pricing implication. This may enhance investor understanding and acceptance, particularly in institutional settings where portfolio construction principles must be clearly communicated. Future research directions include: extending Archimedean constructions to multi-period settings with stochastic control; incorporating alternative risk measures (CVaR, downside deviation) into the equilibrium conditions; exploring higher-order mechanical concepts (angular momentum, gyroscopic stability) as portfolio design principles; testing Archimedean strategies on other asset classes (fixed income, commodities, international equities); developing Archimedean versions of active portfolio management and alpha generation.

Disclosure of interest & Funding

There are no interests to declare from the authors. No funding was received for doing this research.

References

- Asness, Clifford S., Andrea Frazzini and Lasse Heje Pedersen (2012), Leverage Aversion and Risk Parity, *Financial Analysts Journal*, 68, 47–59.
- Best, Michael J. and Robert R. Grauer (1991), On the Sensitivity of Mean-Variance-Efficient Portfolios to Changes in Asset Means: Some Analytical and Computational Results, *Review of Financial Studies*, 4, 315–342.
- Clarke, Roger, Harindra de Silva and Steven Thorley (2006), Minimum-Variance Portfolios in the US Equity Market, *Journal of Portfolio Management*, 33, 10–24.
- Clarke, Roger, Harindra de Silva and Steven Thorley (2012), Risk Parity, Maximum Diversification, and Minimum Variance: An Analytic Perspective, *Journal of Portfolio Management*, 39, 39–53.
- DeMiguel, Victor, Lorenzo Garlappi and Raman Uppal (2009), Optimal versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy?, *Review of Financial Studies*, 22,

1915–1953.

- Donohue, Christopher and Kenneth Yip (2003), Optimal Portfolio Rebalancing with Transaction Costs, *Journal of Portfolio Management*, 29, 49–63.
- Fama, Eugene F. and Kenneth R. French (1993), Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics*, 33, 3–56.
- Fernholz, Robert (2002), *Stochastic Portfolio Theory*. New York: Springer.
- Goldberg, Lisa R. and Othmane Mahmoud (2008), Beyond Leverage, *Risk*, 21, 102–107.
- Grinold, Richard C. and Ronald N. Kahn (2000), *Active Portfolio Management*, 2nd ed. New York: McGraw-Hill.
- Gu, Shihao, Bryan Kelly and Dacheng Xiu (2020), Empirical Asset Pricing via Machine Learning, *Review of Financial Studies*, 33, 2223–2273.
- Jorion, Philippe (1986), Bayes-Stein Estimation for Portfolio Analysis, *Journal of Financial and Quantitative Analysis*, 21, 279–292.
- Kirby, Chris and Barbara Ostdiek (2012), It’s All in the Timing: Simple Active Portfolio Strategies That Outperform Naive Diversification, *Journal of Financial and Quantitative Analysis*, 47, 437–467.
- Ledoit, Olivier and Michael Wolf (2003), Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection, *Journal of Empirical Finance*, 10, 603–621.
- Maillard, Sébastien, Thierry Roncalli and Jérôme Teïletche (2010), The Properties of Equally Weighted Risk Contribution Portfolios, *Journal of Portfolio Management*, 36, 60–70.
- Mantegna, Rosario N. and H. Eugene Stanley (1999), *Introduction to Econophysics: Correlations and Complexity in Finance*. Cambridge: Cambridge University Press.
- Markowitz, Harry (1952), Portfolio Selection, *Journal of Finance*, 7, 77–91.
- Merton, Robert C. (1972), An Analytic Derivation of the Efficient Portfolio Frontier, *Journal of Financial and Quantitative Analysis*, 7, 1851–1872.
- Michaud, Richard O. (1989), The Markowitz Optimization Enigma: Is ‘Optimized’ Optimal?, *Financial Analysts Journal*, 45, 31–42.
- Osborne, M. F. M. (1959), Brownian Motion in the Stock Market, *Operations Research*, 7, 145–173.
- Plaxco, Lisa M. and Robert D. Arnott (2002), Rebalancing a Global Policy Benchmark, *Journal of Portfolio Management*, 28, 9–22.
- Roncalli, Thierry (2013), *Introduction to Risk Parity and Budgeting*. Boca Raton, FL: Chapman and Hall/CRC.
- Ross, Stephen A. (1976), The Arbitrage Theory of Capital Asset Pricing, *Journal of Economic Theory*, 13, 341–360.
- Scherer, Bernd (2011), A Note on the Returns from Minimum Variance Investing, *Journal of Empirical Finance*, 18, 652–660.

Sharpe, William F. (1964), Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk, *Journal of Finance*, 19, 425–442.

Appendix

A Implementation Details and Tables of Results

For our backtesting exercise we have used the following implementation details.

- (a) **Covariance estimation:** We use the sample covariance matrix with Ledoit–Wolf shrinkage (Ledoit and Wolf, 2003) to improve conditioning.
- (b) **Expected return estimation:** For the ALP and Tangency portfolios, we estimate μ_i via a James–Stein shrinkage estimator toward the cross-sectional grand mean, with shrinkage intensity 0.5, computed over the trailing estimation window specified per configuration (12 months for all results reported here) rather than a fixed 60-month window. This shrinkage is applied identically to both ALP and Tangency, so differences between the two in Table 2 reflect the marginal effect of the $\mu_i \mapsto \sigma_i \mu_i$ transform of Theorem 3.2 on top of this common pre-shrinkage, not that transform in isolation.
- (c) **Transaction costs:** We evaluate four transaction-cost tiers – flat 5 bps, flat 10 bps, flat 25 bps, and a tiered schedule that varies by asset class (Treasury: 5 bps, commodity: 20 bps, REIT: 20 bps, S&P 500 index proxy: 10 bps, with all other tickers – including the sector-ETF constituents of the primary universe – falling under a 10 bps default rate). For the primary S&P 500 sector-ETF configuration, the tiered schedule is therefore effectively equivalent to the flat 10 bps scenario, which is consistent with our finding that the ALP-versus-ACM ranking is insensitive to the choice of tier in that universe.
- (d) **Dynamic rebalancing:** For the torque-band rule (ALP), we set $\theta = 0.05$. For the center-of-mass tracking rule (ACM), we set $\phi = 0.10$. These thresholds are chosen to balance tracking accuracy and turnover.
- (e) **Position limits:** All reported portfolios (ALP, ACM, Tangency, Minimum-Variance) are long-only and capped at a maximum individual weight of 3% for the S&P 500 sector-ETF universe (20% for the Russell and multi-asset-class universes of Section 5.2), with excess weight redistributed proportionally and the vector renormalized to satisfy $\mathbf{1}^\top w = 1$. The closed-form solutions of Theorems 3.2 and 3.8 are therefore reported as capped-and-renormalized projections; this projection is an implementation detail, not part of the theoretical derivation.

Table 1: Summary of Backtest Configurations

Label	Period	Freq.	Est. Window	Rebal.	ϕ_{ACM}	γ_{base}
2000-D	2000–2025	Daily	1 yr	Monthly	0.02	0.50
2010-D	2010–2025	Daily	1 yr	Monthly	0.02	0.10
2010-M	2010–2025	Monthly	1 yr	Monthly	0.02	—
2020-D	2020–2025	Daily	1 yr	Monthly	0.10	0.05
2020-M	2020–2025	Monthly	1 yr	Monthly	0.10	0.50

Notes: All configurations use transaction costs of 10 bps per dollar traded, a risk-free rate of 1% per annum, and ALP torque threshold $\theta = 0.05$. Dynamic γ adapts to realized market volatility as described in Section 3. The portfolio universe consists of nine U.S. sector ETFs (XLK, XLF, XLV, XLE, XLI, XLY, XLP, XLB, XLU) with SPY as the passive benchmark.

Table 2: Performance Summary: 2000–2025 (Daily) and 2010–2025 (Daily)

Strategy	2000–2025, Daily							2010–2025, Daily					
	Ret.(%)	Vol.(%)	SR	Sor.	Calmar	MaxDD(%)	α (%)	Ret.(%)	Vol.(%)	SR	Sor.	MaxDD(%)	α (%)
Equal Weight	9.08	18.36	0.44	0.55	0.17	−52.53	10.51	12.45	16.57	0.69	0.84	−36.91	0.00
Min. Variance	7.44	13.74	0.47	0.60	0.22	−34.46	8.14	10.36	13.41	0.70	0.87	−34.46	10.85
Tangency	8.80	21.01	0.37	0.45	0.18	−50.17	10.66	12.39	17.10	0.67	0.82	−33.82	13.42
Risk Parity	8.54	17.42	0.43	0.54	0.17	−48.99	9.82	11.93	15.91	0.69	0.84	−35.22	12.61
ALP	10.75	18.35	0.53	0.68	0.25	−43.71	12.16	13.68	17.37	0.73	0.89	−33.61	14.44
ACM	7.62	13.83	0.48	0.61	0.22	−34.46	8.34	10.30	13.43	0.69	0.86	−34.46	10.82
Benchmark (SPY)	8.63	19.17	0.40	0.50	0.16	−55.19	10.32	14.00	17.20	0.76	0.92	−33.72	14.69

Notes: SR = Sharpe ratio; Sor. = Sortino ratio; Calmar = annualized return / —MaxDD—; α = annualized Fama–French alpha (% p.a.). Cumulative wealth (CW) omitted here; see text. Bold entries indicate the best active-strategy value for each metric.

Table 3: Performance Summary: 2010–2025 (Monthly)

Strategy	Ret.(%)	Vol.(%)	SR	Sor.	Calmar	MaxDD(%)	CW(\$)	Turn.(%)	α (%)	R^2
Equal Weight	12.31	13.78	0.82	1.16	0.52	−23.60	465.0	0.0	14.74	0.035
Min. Variance	12.80	11.91	0.99	1.55	0.75	−17.11	502.9	17.6	15.67	0.046
Tangency	12.79	14.62	0.81	1.35	0.75	−17.02	502.0	20.7	15.73	0.030
Risk Parity	12.34	12.85	0.88	1.24	0.60	−20.67	467.0	3.1	14.66	0.038
ALP	13.35	13.74	0.90	1.29	0.75	−17.85	548.7	8.0	16.28	0.042
ACM	12.80	11.91	0.99	1.55	0.75	−17.11	502.9	17.6	15.67	0.046
Benchmark (SPY)	13.91	14.06	0.92	1.34	0.58	−23.93	597.9	0.0	16.52	0.029

Notes: CW = cumulative terminal wealth per \$1 invested (quoted in absolute \$ terms over the 15-year horizon). Turn. = average monthly one-way turnover. Calmar = annualized return / —MaxDD—.

Table 4: Performance Summary: 2020–2025 (Daily and Monthly)

Strategy	2020–2025, Daily						2020–2025, Monthly					
	Ret.(%)	Vol.(%)	SR	Sor.	MaxDD(%)	α (%)	Ret.(%)	Vol.(%)	SR	Sor.	MaxDD(%)	α (%)
Equal Weight	13.22	14.89	0.82	1.13	−16.94	0.00	13.55	14.22	0.88	1.59	−15.82	12.90
Min. Variance	5.65	12.40	0.38	0.51	−16.68	11.84	8.07	14.57	0.49	0.87	−17.11	11.28
Tangency	11.14	16.92	0.60	0.81	−17.49	25.54	13.27	17.88	0.69	1.26	−11.82	11.15
Risk Parity	11.58	14.31	0.74	1.01	−17.14	11.08	12.48	14.01	0.82	1.44	−15.09	12.58
ALP	14.48	15.94	0.85	1.17	−14.90	10.24	16.43	15.38	1.00	1.51	−13.63	15.47
ACM	5.40	12.44	0.35	0.49	−16.60	9.09	8.07	14.57	0.49	0.87	−17.11	11.28
Benchmark (SPY)	14.57	17.15	0.79	1.09	−24.50	0.00	15.01	15.22	0.92	1.51	−23.93	14.43

Notes: The 2020–2025 window spans the COVID-19 drawdown (March 2020), the subsequent recovery, and the 2022 rate-hiking cycle. Bold entries indicate the best value among active strategies for each metric.

Table 5: Annualized Factor Alpha (% p.a.) Across Configurations

Strategy	2000-D	2010-D	2010-M	2020-D	2020-M
Equal Weight	10.51	0.00	14.74	0.00	12.90
Min. Variance	8.14	10.85	15.67	11.84	11.28
Tangency	10.66	13.42	15.73	25.54	11.15
Risk Parity	9.82	12.61	14.66	11.08	12.58
ALP	12.16	14.44	16.28	10.24	15.47
ACM	8.34	10.82	15.67	9.09	11.28
Benchmark (SPY)	10.32	14.69	16.52	0.00	14.43

Notes: Alphas are estimated by regressing each strategy’s excess returns on contemporaneous Fama–French three factors (market excess return, SMB, HML). Annualization uses $(252, \sqrt{252})$ scaling for daily and $(12, \sqrt{12})$ for monthly estimates. The anomalously high Tangency alpha in 2020-D (25.54%) reflects the strategy’s heavy tilt toward growth/technology sectors during the 2020–2021 post-COVID rally; this tilt also produced the highest volatility (16.92%) and deepest drawdown (−17.49%) in that sub-panel.

B Dynamic Risk Scaling and Adaptive γ

Note: the adaptive- γ mechanism described in this appendix was not used to generate the results in Section 5, which fix $\gamma = 1$ throughout. It is presented here as a design proposal for production deployment.

While the closed-form Archimedean Center-of-Mass solution from Theorem 3.8 is mathematically elegant, practical implementation requires safeguards against extreme allocations during volatile market regimes. We introduce a dynamic scaling mechanism that adapts the factor-alignment parameter γ to prevailing market conditions while preserving the Archimedean structure.

B.0.1 Motivation

The constraint $B^\top w = \gamma \mu_f$ enforces alignment between the portfolio's factor center-of-mass and the factor risk premium. However, when market volatility spikes or factor relationships break down, this rigid constraint can force extreme positions. We address this by making γ state-dependent rather than constant.

B.0.2 Volatility-Adaptive Gamma

Let $\bar{\sigma}_t$ denote the average asset volatility at rebalancing date t :

$$\bar{\sigma}_t = \frac{1}{n} \sum_{i=1}^n \hat{\sigma}_{i,t}, \quad (53)$$

where $\hat{\sigma}_{i,t}$ is the estimated standard deviation of asset i over the trailing estimation window. Similarly, define the historical “normal” volatility $\bar{\sigma}$ as the full-sample average.

We define the volatility ratio:

$$\rho_t = \frac{\bar{\sigma}_t}{\bar{\sigma}}. \quad (54)$$

The adaptive gamma at time t is then:

$$\gamma_t = \gamma_0 \cdot h(\rho_t), \quad (55)$$

where γ_0 is the baseline parameter and $h(\cdot)$ is a scaling function. We use a piecewise-linear dampening rule:

$$h(\rho) = \begin{cases} 1.0 & \text{if } \rho \leq 1.2, \\ 0.75 & \text{if } 1.2 < \rho \leq 1.5, \\ 0.5 & \text{if } \rho > 1.5. \end{cases} \quad (56)$$

This reduces factor-alignment aggressiveness when recent volatility exceeds historical norms by more than 20%, preserving capital during turbulent periods.

B.0.3 Feasibility-Based Gamma Scaling

In addition to volatility-based scaling, we impose a feasibility constraint to ensure the target factor exposure $\gamma_t \mu_f$ remains attainable without extreme leverage. Let w^{MV} denote the unconstrained minimum-variance portfolio:

$$w^{MV} = \frac{\Sigma^{-1} \mathbf{1}}{\mathbf{1}^\top \Sigma^{-1} \mathbf{1}}. \quad (57)$$

Its factor exposure is $\beta^{MV} = B^\top w^{MV}$. We constrain the ACM portfolio to deviate from this baseline by at most a factor κ :

$$\|\gamma_t \mu_f - \beta^{MV}\| \leq \kappa \|\beta^{MV}\|, \quad (58)$$

where $\kappa > 0$ is a safety parameter (we use $\kappa = 2$). If the raw γ_t violates this bound, we scale it down:

$$\tilde{\gamma}_t = \gamma_t \cdot \min \left\{ 1, \frac{\kappa \|\beta^{MV}\|}{\|\gamma_t \mu_f - \beta^{MV}\|} \right\}. \quad (59)$$

B.0.4 Weight Clipping and Variance Checks

Even with adaptive γ , numerical issues can produce pathological weights. We impose post-optimization safeguards:

1. **Hard bounds:** Clip weights to $w_i \in [-0.2, 1.5]$ before normalization.
2. **Variance check:** If the ACM portfolio variance exceeds 4 times the minimum-variance portfolio variance, blend the two:

$$w^{safe} = (1 - \alpha)w^{MV} + \alpha w^{ACM}, \quad (60)$$

with $\alpha = 0.3$.

3. **Schur complement conditioning:** If $\text{cond}(S) > 10^8$, fall back to w^{MV} .

B.0.5 Preservation of Archimedean Principles

The dynamic adjustments preserve the Archimedean framework:

- The optimization structure remains the same; only γ is scaled.
- The closed-form solution (Theorem 3.8) is still used at each step.
- The constraint $B^\top w = \gamma_t \mu_f$ still enforces center-of-mass alignment, but with a time-varying, prudent target.

In effect, we interpret γ not as a fixed preference parameter but as a risk-scaled alignment intensity—high when conditions permit, low when caution is warranted. This renders the ACM strategy robust to regime changes while maintaining its mechanical interpretation.

B.0.6 Algorithm Summary

At each rebalancing date t :

1. Estimate Σ_t , B_t , $\mu_{f,t}$ using trailing data.
2. Compute volatility ratio ρ_t and apply dampening: $\gamma_t = \gamma_0 \cdot h(\rho_t)$.
3. Compute feasibility-adjusted $\tilde{\gamma}_t$ relative to w^{MV} .
4. Solve Theorem 4.1 with $\gamma = \tilde{\gamma}_t$.

5. Apply weight clipping and variance checks.
6. If all checks pass, rebalance to w_t^{ACM} ; otherwise use w^{MV} .

This dynamic extension enables practical deployment while honoring the theoretical elegance of the Archimedean formulation.