

Stochastic Portfolio Theory, Machine-Checked: What Is Proven, What Is Sealed Against Mathlib, and What Stays a Modeling Assumption

A verified trust-map for functionally generated portfolios

Dr. Tamás Nagy

ORCID: 0009-0004-8079-4679

Independent Researcher

tnagyphd@gmail.com

<https://thalens.org>

Working Paper

Build-std v2.0+3e9dce8e72 | 2026-09-01 01:26 | 0d8fc1a1

Practitioner's Summary

Stochastic Portfolio Theory (SPT), developed by Robert Fernholz, makes a claim that sounds too good to be true: a portfolio built mechanically from a suitable function of the market weights — no forecasts, no estimated expected returns — can capture a structural growth advantage over the capitalization-weighted market index. In the trace-model formulation used in this paper, the sign certificate applies when the generator is **log-concave**. The textbook examples are the equal-weighted portfolio and the entropy-weighted portfolio. The advantage has a name, the *excess growth rate*.

This paper does not add a new strategy. It does something a practitioner should arguably want more: it takes the SPT machinery apart and labels every piece by how much you can trust it. Each result is run through a proof kernel that refuses to accept a step it cannot check, and the two inequalities the theory actually leans on are additionally re-proved inside Mathlib, the formal mathematics library, with no gaps.

The result is a trust-map with three honest tiers. **First**, the elementary engine of diversification — the statement that a component-weighted average variance dominates the corresponding diagonal basket variance — is machine-derived from scratch as exact algebra. **Second**, the two analytic facts the theory cannot do without — that the entropy of the market weights is bounded by the log of the number of stocks, and that the explicit trace-model excess growth is non-negative for a log-concave generator and a valid covariance — are sealed against Mathlib's own theorems on convexity and positive-semidefinite matrices. These are checked conditionally on the displayed trace definition. **Third**, and stated just as plainly: the central identity of the theory, Fernholz's *master formula*, is an Itô-calculus decomposition that the kernel takes as a *premise*, not as something it proves. Mathlib v4.28 itself does not contain the general multidimensional, random-adapted-integrand machinery needed to discharge the full master formula, and we say so rather than dress an assumption as a theorem.

There is one more thing a practitioner must hear, because it is the most common way SPT is oversold. The excess growth rate being non-negative is a statement about a *rate*, not a guarantee

of *realized* outperformance over any finite horizon. We include a direct Monte-Carlo check of this. In a simulated market, the diversity- and entropy-weighted portfolios beat the index on only a **modest majority** of sample paths over five years, and a naive “small-cap premium” came out **negative** in realized terms. Neither portfolio wins path by path. The growth advantage is structural, but it competes against boundary leakage and finite-horizon dispersion. The honest takeaway is that SPT gives you a favorable component, not a free lunch.

What this buys the reader is confidence calibrated to evidence. When we say the diagonal rebalancing premium is positive, that is algebra you can re-check in seconds. When we say the trace-model sign follows from log-concavity and valid covariance, that is a Lean proof you can run. When we discuss the Atlas small-cap effect, that is a *conditional* statement resting on model assumptions we name explicitly. The contribution is the separation itself — and the discipline of never inviting you to mistake one tier for another.

Abstract

Stochastic Portfolio Theory provides a model-free framework in which functionally generated portfolios track the market index up to two terms: a boundary term in the generating function and the time-integral of an *excess growth rate*. We formalize the logical skeleton of this framework in a proof kernel and report, with deliberate precision, which parts are derived and which are assumed.

Three tiers result. (i) The elementary algebra of diversification — that a component-weighted average variance dominates the corresponding diagonal basket variance — is derived from first principles as a sum-of-squares certificate, generalized term-by-term beyond the two-asset case; with full correlation, the excess growth of an arbitrary portfolio is a non-negatively weighted sum of pairwise return-difference variances, $\gamma_\pi^* = \frac{1}{2} \sum_{i < j} w_i w_j \text{Var}(R_i - R_j) \geq 0$, provided those quantities come from a valid covariance matrix. The two-asset equal-weight specialization is derived explicitly as $\gamma_{\text{EW}}^* = \text{Var}(R_1 - R_2)/8$. (ii) The two analytic inequalities the theory cannot avoid are discharged against Mathlib: the diversity bound $0 \leq H(\mu) \leq \log n$ on the Shannon entropy of market weights (via Jensen and the concavity of $-x \log x$), and the sign of the explicitly defined trace-model rate $\gamma^* \geq 0$ for log-concave G and positive-semidefinite weighted covariance. Both Lean proofs are sorry-free and depend only on Mathlib’s standard axioms. (iii) The Fernholz master formula itself is retained as an explicit modeling hypothesis. Separate external Lean developments verify scalar Itô decompositions of $\log G(\mu)$ in two concrete models, but those decompositions do not contain the self-financing portfolio-wealth identity and therefore do not discharge the master formula.

We complement the formal layer with two empirical checks: a real-data validation on the Ken French 49-industry cross-section (2000–2023) confirming that every load-bearing inequality holds with economically large margins — a 1.5%/yr equal-weight excess growth rate, the trace/pairwise and diversification-return identities exact to machine precision — and a Monte-Carlo study showing that a non-negative excess growth *rate* does not imply realized outperformance over finite horizons. We then map the discharge path for the remaining assumption. The contribution is not a new portfolio but a verified, honestly-tiered account of an influential theory.

Epistemic status (read this first)

This paper mixes three kinds of statement. They are labeled throughout so the reader is never invited to mistake one for another.

Tier	What it means here	Representative results
D — Derived from primitives	Proved inside the kernel as exact real algebra (a Z3-backed sum-of-squares / Positivstellensatz certificate), under the variables and sign premises stated for each result.	$\bar{\sigma}_w^2 \geq \sigma_w^2$ (diagonal rebalancing premium ≥ 0); AM–GM core $uv \leq (u^2 + v^2)/2$; $0 < G_{EW} \leq 1$; diagonal excess growth ≥ 0 ; the concentration sandwich $\frac{1}{2}\sigma_{\min}^2(1 - H) \leq \gamma_{w,\text{diag}}^* \leq \frac{1}{2}\sigma_{\max}^2(1 - H)$; the correlated form $\gamma_{\pi}^* = \frac{1}{2} \sum_{i < j} w_i w_j \text{Var}(R_i - R_j) \geq 0$ for valid covariance input; the SPT–MPT bridge $2\gamma_{\pi}^* - [(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w] = \sum_{i < j} w_i w_j (\sigma_i - \sigma_j)^2$.
S — Sealed against Mathlib	The kernel states it as a hypothesis (a transcendental or higher-dimensional fact the scalar real-arithmetic kernel cannot close), but it is independently machine-proved in Lean 4 / Mathlib, sorry-free, standard axioms only.	Diversity bound $0 \leq H(\mu) \leq \log n$; trace-model excess-growth sign for log-concave G and positive-semidefinite covariance.
A — Assumed (modeling premise)	Posited in the kernel as an explicit, named hypothesis/fact and <i>restated</i> as a theorem. Not derived.	Fernholz master formula (Itô decomposition); portfolio representation $\pi_i = \mu_i D_i \log G$; the Atlas rank-model premia (small-cap, rebalancing-as-ergodic, Pareto tail).

“Machine-checked” without further qualification means tier **D** or **S**. Whenever a tier-**A** statement is used, it is flagged. The point of the paper is the boundary between these tiers, drawn exactly where the mathematics — not the marketing — puts it.

1. Introduction

1.1 The problem

A portfolio manager who has lived through a backtest knows the gap between a *rate* and a *result*. A model says the strategy earns 2% a year over its benchmark; the realized five-year track record is a coin flip. Stochastic Portfolio Theory sits precisely on this fault line. It proves — cleanly, and

without any forecast of expected returns — that a large class of mechanical portfolios carries a non-negative *excess growth rate* relative to the market. The theory is mathematically beautiful and has launched real funds. It is also routinely oversold, because the step from “non-negative growth rate” to “beats the market” is exactly the step the theory does *not* take.

So there are two questions a careful reader should ask of SPT, and they are different questions. First: *is the mathematics correct* — not “has it been peer-reviewed”, but “does every inequality actually follow, with no quietly-assumed step?” Second: *what does the correct mathematics actually promise* — a drift, or an outcome? This paper answers both, and keeps them apart.

1.2 Existing approaches

Fernholz’s original development (Fernholz, 2002; Fernholz and Karatzas, 2009) is rigorous in the ordinary mathematical sense: the master formula is an application of Itô’s lemma to $\log G(\mu(t))$, the excess growth rate emerges as the Itô correction term, and concavity of G delivers its sign. The arguments are standard stochastic analysis. What they are not — what no pencil-and-paper treatment is — is *mechanically checked*. The sign of the excess growth rate, in particular, is a positive-semidefinite quadratic-form statement that is “clearly true” in the way that invites a reader to nod past it. Our interest is in turning each such nod into either a certificate or an explicit, labeled assumption.

A second strand is empirical: diversity-weighted indexing and the documented small-cap and re-balancing premia (Fernholz, Garvy and Hannon, 1998). Here the literature is honest that realized performance is regime-dependent. We connect to this by reproducing, in a controlled simulation, the precise sense in which a positive growth rate fails to guarantee positive realized return.

1.3 This paper’s contribution

One claim, stated once. *We give a machine-checked account of the SPT framework that draws an exact line between what is proven, what is sealed against an external formal library, and what remains a modeling assumption — and we seal the two analytic inequalities the theory cannot do without.* The deliverable is the trust-map, not a portfolio.

Two consequences of that claim are worth stating because they are, to our knowledge, new at the level of *formalization*:

1. Under the paper’s explicit trace-model definition, non-negativity of the excess growth rate for a log-concave generator in any dimension is reduced to, and machine-proved as, the linear-algebra fact that **the trace of a product of two positive-semidefinite matrices is non-negative** (Section 5.2). The matrix theorem is sorry-free; the identification of that trace with the SPT drift remains part of the named model premise.
2. The diversity bound $0 \leq H(\mu) \leq \log n$ is sealed against Mathlib’s Jensen inequality and the concavity of $-x \log x$ (Section 5.1).

Beyond the formalization, the derived algebra collects a set of *quantitative* relations under one discipline. We make no strong novelty claim for the underlying mathematics: several ingredients are classical finance facts, including the Booth–Fama diversification return, the Elton–Gruber systematic-risk floor, and the excess-growth-rate optimization studied by Maeso and Martellini (2020). The contribution is to state the relations explicitly, certify the algebra, and lift the central sign to all n .

First (Section 4.4), the diagonal excess growth rate of an arbitrary weight vector is sandwiched, above and below, by the variance level times its Simpson diversity $1 - H$. Second (Section 4.6), the correlated excess growth rate of an arbitrary portfolio is shown to equal a non-negatively weighted sum of pairwise return-difference variances, with an exact equal-weight closed form only in the diagonal-covariance specialization. The same algebra certifies four portfolio-design facts in Section 4.7. Section 4.8 then derives the exact identity $2\gamma_\pi^* - [(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w] = \sum_{i < j} w_i w_j (\sigma_i - \sigma_j)^2$, together with a two-sided envelope and observable-input bounds. Section 4.9 gives hierarchical and asset-level decompositions. Finally, Section 7.1 tests the load-bearing identities on twenty-four years of Ken French industry returns and finds a 1.5%/yr equal-weight excess growth rate with the identities holding to machine precision.

1.4 Non-claims

We are explicit about what this paper does **not** do.

- We do **not** prove the Fernholz master formula, including in the two concrete models of Section 8. The external Lean developments prove the corresponding generator-side Itô decompositions; the missing bridge is the self-financing relative-wealth identity that turns a generator decomposition into a portfolio master formula.
- We do **not** claim functionally generated portfolios beat the market in realized terms. Section 7 shows directly that they often do not over five-year horizons.
- We do **not** claim the Atlas-model premia (small-cap, rebalancing-as-ergodic) as derived. They are conditional on rank-model hypotheses we state as such.
- “Sealed against Mathlib” means the truth rests on Mathlib’s audited theorems and standard axioms, not that the SPT kernel proves it natively. We distinguish the two.

1.5 Organization

Sections 2 and 3 fix the notation, state the modeling premises, and present the formalized framework as tier **A**. Section 4 then develops the derived algebra of diversification as tier **D**, while Section 5 presents the two Mathlib seals as tier **S**.

Section 6 treats relative arbitrage and the Atlas premia as conditional results. Section 7 supplies the Monte-Carlo honesty check, and Section 8 isolates the remaining general assumption and its discharge path.

Section 9 discusses implications and limits. The appendices index the formal statements, Lean bridges, and reproduction script.

2. Setup and notation

We consider a market of n stocks with positive capitalizations. The **market weight** of a stock is its share of total capitalization:

$$\mu_i(t) = \frac{X_i(t)}{\sum_{j=1}^n X_j(t)}, \quad \mu_i(t) > 0, \quad \sum_{i=1}^n \mu_i(t) = 1,$$

Thus the weight vector lives on the open unit simplex. In the kernel, `thm_weight_range` records positivity and `thm_weights_sum_one` records that the weights sum to one.

A **generating function** $G : \Delta^{n-1} \rightarrow (0, \infty)$ is a positive, twice-differentiable function of the market weights. The portfolio it generates — the *functionally generated portfolio* (FGP) — holds, in stock i , the weight

$$\pi_i(t) = \mu_i(t) \left[D_i \log G(\mu(t)) + 1 - \sum_{j=1}^n \mu_j(t) D_j \log G(\mu(t)) \right],$$

which we abbreviate, in the kernel's representation fact (`thm_portfolio_gen`, tier **A**), to the defining relation $\pi_i \propto \mu_i D_i \log G$. Two running examples:

- **Equal-weighted:** $G(\mu) = (\prod_i \mu_i)^{1/n}$, the geometric mean. We write G_{EW} for its two-asset instance $\sqrt{uv}$.
- **Entropy-weighted:** $G(\mu) = -\sum_i \mu_i \log \mu_i = H(\mu)$, the Shannon entropy of the weights.

The **relative value** of the FGP against the market is $V_\pi(t)/V_\mu(t)$, and the central object is its logarithm. The **excess growth rate** $\gamma^*(t)$ is the quantity SPT identifies as the FGP's structural growth advantage; its model definition (an Itô correction term) is the content of Section 3.

We use σ for the covariance matrix of instantaneous relative returns (the market's local covariance structure) and $D^2 \log G$ for the Hessian of $\log G$. *Log-concavity of G* means exactly that $\log G$ is concave, or equivalently $-D^2 \log G \succeq 0$. This is the curvature condition used by the trace-model sign certificate below; it should not be confused with the distinct standard SPT formulation written using $D^2 G/G$.

Assumption block (made explicit, relaxed nowhere in this paper). We take as given that (a) the capitalization processes are continuous semimartingales with a well-defined instantaneous covariation σ , (b) G is C^2 with the regularity needed for Itô's lemma, and (c) the finite-variation term is identified with the paper's trace-model rate $\gamma^* = \frac{1}{2} \text{tr}((-D^2 \log G)M)$. Assumptions (a)–(b) support the standard functionally generated-portfolio decomposition, but they do not by themselves imply the additional trace identification in (c): the standard Hessian relation $D^2 G/G = D^2 \log G + (D \log G)(D \log G)^\top$ contains a gradient-square term. The master formula used below therefore includes (c) as a separate tier-**A** premise. The kernel does not derive these stochastic inputs; it sees their stated consequence.

3. The formalized framework (tier A — the premises)

This section states the two facts the kernel *posits*. They are the economic-mathematical content of SPT; everything machine-derived later is a consequence of these plus elementary algebra.

3.1 The master formula

The cornerstone of SPT is the decomposition of relative log-value into a boundary term and a growth integral.

Premise 1 (Fernholz master formula; tier A). For an FGP generated by G ,

$$\log \frac{V_\pi(T)}{V_\mu(T)} = \underbrace{\log G(\mu(T)) - \log G(\mu(0))}_{\text{boundary (leakage) term}} + \int_0^T \gamma^*(t) dt.$$

In the kernel this is `F_master_formula`, restated as the theorem `thm_master_formula`. **It is assumed, not derived.** The identity is Itô's lemma applied to $\log G(\mu(t))$: the boundary term is the difference of the generating function at the endpoints, and the integrand γ^* is the second-order Itô correction. The kernel has no stochastic calculus, so it cannot produce this decomposition; it takes it as input. Section 8 returns to exactly what it would take to remove this assumption.

The decomposition is the source of every honest insight in SPT. Two definitional consequences are recorded as kernel theorems and *are* derived from the premise by linear arithmetic: the leakage term is the boundary difference (`thm_leakage`), and relative value splits as leakage plus growth integral (`thm_decomposition`).

3.2 The excess growth rate and its model definition

The integrand γ^* is where the theory earns its keep. Its model definition is the Itô-correction quadratic form

$$\gamma^*(t) = -\frac{1}{2} \sum_{i,j} \mu_i \mu_j (D^2 \log G)_{ij} \sigma_{ij} = \frac{1}{2} \text{tr}((-D^2 \log G) M), \quad M := \text{diag}(\mu) \sigma \text{diag}(\mu),$$

where σ is the relative-return covariance and M is its weighting by the diagonal market-weight matrix. This displayed equation is the **explicit model definition used in this paper**; the formal seal does not derive it from the standard functionally generated portfolio formula. We keep the weighting explicit because M is a congruence of σ , so a valid covariance $\sigma \succeq 0$ gives $M \succeq 0$ — the property the sign argument of Section 5.2 uses. For the market portfolio itself ($G \equiv \text{const}$, no curvature) the rate vanishes: $\gamma_{\text{mkt}}^* = 0$ (`thm_market_zero`). Conditional on this definition, the sign claim is: *log-concave G and valid covariance imply $\gamma^* \geq 0$* . In the kernel this claim is the hypothesis `H_gamma_nn`, restated as `thm_concave_growth`. Section 5.2 seals the matrix implication while retaining the trace identification as tier **A**.

4. The derived algebra of diversification (tier D)

Here we leave the premises behind. Every result in this section is proved inside the kernel as a sum-of-squares certificate — a Positivstellensatz witness found by a Z3-backed non-linear arithmetic tactic — with no premise beyond the stated variables. These are the parts of SPT that are *true as algebra*, independent of any stochastic model.

4.1 Why rebalancing earns a premium: averaged variance dominates the variance of the average

The single most useful elementary fact in SPT is that a *diversified* basket has lower diagonal variance than its component-weighted average variance. Concretely, with arbitrary portfolio weights $w_i \in$

$(0, 1)$ summing to one and per-asset variances $s_i^2 \geq 0$, define $\bar{\sigma}_w^2 = \sum_i w_i s_i^2$ and $\sigma_{w,\text{diag}}^2 = \sum_i w_i^2 s_i^2$. Then:

Theorem 1 (diagonal rebalancing premium ≥ 0 ; tier D). $\bar{\sigma}_w^2 \geq \sigma_{w,\text{diag}}^2$.

Why it is true, in one line. The difference telescopes term-by-term into a manifestly non-negative sum:

$$\bar{\sigma}_w^2 - \sigma_{w,\text{diag}}^2 = \sum_i (w_i - w_i^2) s_i^2 = \sum_i w_i(1 - w_i) s_i^2 \geq 0,$$

since each $w_i \in (0, 1)$ gives $w_i(1 - w_i) > 0$ and each $s_i^2 \geq 0$. The half-difference $(\bar{\sigma}_w^2 - \sigma_{w,\text{diag}}^2)/2$ is the diagonal excess-growth rate of the portfolio with weights w . Equal weighting is the specialization $w_i = 1/n$, not a premise of this identity.

Two points of honesty about Theorem 1. First, the kernel proves the **three-asset** case ($n = 3$) explicitly, but the certificate is *per-summand* — it never couples the indices — so it carries verbatim to any n ; we present it as the diagonal-covariance representative of the general statement, not as a two-asset coincidence. Second, the proof needs only $w_i \in [0, 1]$ and $s_i^2 \geq 0$; the simplex constraint $\sum w_i = 1$ is retained for fidelity to the SPT setup but is not load-bearing for the inequality. The kernel discharges this through a non-linear arithmetic certificate with no human-supplied hint terms — the solver finds the squares.

4.2 The equal-weighted generating function: a premise-free AM–GM

The equal-weighted portfolio is generated by the geometric mean $G_{\text{EW}} = \sqrt{uv}$ (two-asset case). Its two defining properties — positivity and an upper bound of one — reduce to the Arithmetic–Mean–Geometric–Mean inequality, which we prove in its cleanest possible form.

Theorem 2 (AM–GM core; tier D). For reals u, v : $uv \leq \frac{u^2 + v^2}{2}$.

This is the bare square $(u - v)^2 \geq 0$ rearranged — proved with **no premises at all**, separating the genuine mathematical content from any normalization. The SPT-facing corollaries follow for the *squared* generating value $g := G_{\text{EW}}^2 = uv$:

Theorem 3 (generating-value bounds; tier D). For $u, v > 0$ under the normalization $(u^2 + v^2)/2 = 1$: $0 < g = uv \leq 1$, hence $0 < G_{\text{EW}} = \sqrt{g} \leq 1$.

Positivity ($0 < g$) is a product of positives; the upper bound ($g \leq 1$) is Theorem 2 composed with the normalization $(u^2 + v^2)/2 = 1$, and it transfers to $G_{\text{EW}} = \sqrt{g}$ because $t \mapsto \sqrt{t}$ is increasing. One honesty note on this normalization: $(u^2 + v^2)/2 = 1$ is an algebraic scaling that isolates the AM–GM mechanism — it is the form in which the kernel (`thm_gew_pos`, `thm_amgm`) carries the bound — and is *not* the market simplex constraint $u + v = 1$. The economic content we extract is purely the mechanism “geometric-mean generating value is bounded”, not a literal market-weight identity. Where an earlier version of the kernel posited these as bare hypotheses, they are now clean consequences of Theorem 2.

4.3 Diagonal-covariance excess growth is non-negative

Combining the diagonal excess-growth definition with Theorem 1 closes the loop for an arbitrary weight vector:

Theorem 4 (diagonal excess growth ≥ 0 ; tier D). $\gamma_{w,\text{diag}}^* = \frac{1}{2}(\bar{\sigma}_w^2 - \sigma_{w,\text{diag}}^2) \geq 0$.

The legacy kernel theorem names `thm_ew_gamma` and `thm_ew_gamma_nn` retain “ew”, but the statements use generic weights w_i ; the paper therefore reports their actual arbitrary-weight content. The equality is definitional and the sign rests on Theorem 1, not on a postulated inequality. Equal weighting follows by setting $w_i = 1/n$. The correlated arbitrary-portfolio case is treated in Section 4.6, and the conditional log-concave trace model in Section 5.2.

4.4 A quantitative sharpening: excess growth is sandwiched by concentration

The textbook statement stops at the sign $\gamma_{w,\text{diag}}^* \geq 0$. The same sum-of-squares structure yields more. Write the portfolio’s **Herfindahl concentration index** $H = \sum_i w_i^2 \in [1/n, 1]$, so that $1 - H$ is the Simpson diversity. Using $\sum_i w_i = 1$, the diagonal diversification gap factorizes as

$$\bar{\sigma}_w^2 - \sigma_{w,\text{diag}}^2 = \sum_i w_i(1 - w_i) s_i^2, \quad \sum_i w_i(1 - w_i) = 1 - H,$$

and bounding each variance s_i^2 between a floor $\sigma_{\min}^2$ and a ceiling $\sigma_{\max}^2$ — term by term, since every coefficient $w_i(1 - w_i) \geq 0$ — gives a two-sided bound:

Theorem 5 (concentration sandwich; tier D). With per-asset variances in $[\sigma_{\min}^2, \sigma_{\max}^2]$,

$$\frac{1}{2} \sigma_{\min}^2 (1 - H) \leq \gamma_{w,\text{diag}}^* \leq \frac{1}{2} \sigma_{\max}^2 (1 - H).$$

Both inequalities are derived in the kernel as cold sum-of-squares certificates (no postulate), at the same $n = 3$ representative level as Theorem 1 and by the same term-by-term argument for general n . They refine $\gamma_{w,\text{diag}}^* \geq 0$: the diagonal rebalancing premium is pinned above and below by the portfolio’s diversity $1 - H$ and the cross-sectional variance band. The lower bound is strictly positive whenever $H < 1$ and $\sigma_{\min}^2 > 0$. When the variance cross-section is tight, $\gamma_{w,\text{diag}}^* \approx \frac{1}{2} \sigma^2 (1 - H)$. Equal weighting is again only the specialization $w_i = 1/n$.

4.5 The diversity range is derived; the diversification return is named

Two further facts fall out of the same elementary algebra, with no stochastic input. First, the gap of Theorem 1 has an exact name. Written out,

$$\bar{\sigma}_w^2 - \sigma_{w,\text{diag}}^2 = \sum_i w_i(1 - w_i) s_i^2,$$

which is the **Booth–Fama diversification return** for weights w , equal to $2\gamma_{w,\text{diag}}^*$ — now a kernel identity, not a slogan.

Second, the most-used concentration measure has a *fully derived* range. For the Herfindahl index $H = \sum_i w_i^2$,

$$\frac{1}{n} \leq H \leq 1, \quad \text{equivalently} \quad 1 \leq N_{\text{eff}} = \frac{1}{H} \leq n,$$

where N_{eff} is the effective number of holdings. The upper bound is $w_i^2 \leq w_i$; the lower bound is the Cauchy–Schwarz inequality $n \sum_i w_i^2 \geq (\sum_i w_i)^2$ specialized by $\sum_i w_i = 1$. Both are cold sum-of-squares certificates.

Where SPT carries the general one-parameter diversity $D_p \in [1, n]$ as a modeling range (Section 2), the canonical Herfindahl / effective-number diversity is here a **theorem**: uniform weights maximize it, a single dominant name minimizes it, and the bounds are proven rather than assumed.

4.6 The correlated general case: excess growth is a sum of pairwise relative variances

Everything in Sections 4.1–4.5 is the *diagonal*-covariance case. The practically central object keeps the full covariance matrix Σ . For any long-only portfolio with weights w on the simplex, the excess growth rate is $\gamma_\pi^* = \frac{1}{2}(\sum_i w_i \sigma_{ii} - w^\top \Sigma w)$, and the same elementary algebra that produced Theorem 1 produces its correlated form. Using $\sum_i w_i = 1$,

$$\sum_i w_i \sigma_{ii} - w^\top \Sigma w = \sum_{i < j} w_i w_j (\sigma_{ii} + \sigma_{jj} - 2\sigma_{ij}) = \sum_{i < j} w_i w_j \text{Var}(R_i - R_j).$$

Theorem 6 (correlated excess growth; tier D). For weights w on the simplex and a valid positive-semidefinite covariance matrix Σ ,

$$\gamma_\pi^* = \frac{1}{2} \left(\sum_i w_i \sigma_{ii} - w^\top \Sigma w \right) = \frac{1}{2} \sum_{i < j} w_i w_j (\sigma_{ii} + \sigma_{jj} - 2\sigma_{ij}) \geq 0,$$

since each pairwise term $\sigma_{ii} + \sigma_{jj} - 2\sigma_{ij} = \text{Var}(R_i - R_j) \geq 0$.

This is the cleanest algebraic statement that diversity generates growth: the excess growth of a portfolio is a non-negatively weighted sum of pairwise return-difference variances. It subsumes the diagonal Theorems 1 and 4. The formal source takes the three inequalities $r_{12}, r_{13}, r_{23} \geq 0$ as named hypotheses; it does not derive them from matrix positive-semidefiniteness. Semantically, those inequalities follow from the stated premise that Σ is a valid covariance matrix because $r_{ij} = (e_i - e_j)^\top \Sigma (e_i - e_j) \geq 0$. The tier-**D** certificate therefore proves the identity and the conditional sign from explicit pairwise non-negativity premises; no premise-free covariance-validity theorem is being claimed. Section 5.2 separately seals the conditional log-concave trace-model sign.

The sign holds for all n — not just the representative. The tier-**D** statements elsewhere in this section are certified at the $n = 3$ representative level with a per-summand argument that carries to general n by inspection. For the *central* claim — the sign of the correlated excess growth rate — we discharge the “by inspection” step and prove the general- n statement outright, over a finite-sum object with N free, using only the kernel’s finite-sum axioms (the kernel counterparts of Mathlib’s `Finset.sum_nonneg`, `sum_add_distrib`, `mul_sum`, and `sum_congr`). Two machine-checked $\forall N$ theorems do this. First, the sign:

Theorem (general- n excess-growth sign; tier D). For every dimension, let the weights be non-negative and let the symmetric pairwise relative variances be non-negative. Then

$$0 \leq \sum_i \sum_j w_i w_j r_{ij}.$$

Substituting return-difference variances into the excess-growth representation gives

$$r_{ij} = \text{Var}(R_i - R_j), \quad \gamma_\pi^* = \frac{1}{4} \sum_i \sum_j w_i w_j r_{ij} \geq 0.$$

Second, the identity that ties this double-sum representation to the trace form of the excess growth rate, again for arbitrary N :

Theorem (general- n pairwise—trace identity; tier D). For every dimension, weight vector, and covariance matrix,

$$\sum_i \sum_j w_i w_j (\sigma_{ii} + \sigma_{jj} - 2\sigma_{ij}) = 2 \left(\sum_k w_k \right) \left(\sum_k w_k \sigma_{kk} \right) - 2 \sum_i w_i \sum_j w_j \sigma_{ij}.$$

On the simplex, the right-hand side is four times the excess growth rate. Together with the preceding theorem, this makes non-negativity a genuine all-dimensional result rather than a representative calculation.

These two are the only statements in the derived layer proved over a symbolic, unbounded index set; the rest remain $n = 3$ representatives. We single out the sign because it is the theory’s load-bearing qualitative claim — the one a referee would most want free of any “representative” caveat.

A closed form follows for the canonical equal-weighted portfolio in the diagonal-covariance specialization. With $w_i = 1/n$ and $\sigma_{ij} = 0$ for $i \neq j$, the gap collapses to a single number:

Corollary (equal-weighted diversification return; tier D). $\gamma_{EW}^* = \frac{1}{2} (1 - 1/n) \overline{\sigma^2}$, where $\overline{\sigma^2} = \frac{1}{n} \sum_i \sigma_{ii}$ is the mean individual variance.

Within that diagonal model, the rebalancing premium grows linearly in the average variance level $\overline{\sigma^2}$ and in the diversity fraction $1 - 1/n$. With nonzero covariance the general pairwise formula above must be used; the displayed closed form is not a full-covariance identity.

4.7 Portfolio-design corollaries

The same drift-free algebra that underlies the excess growth rate also settles four facts that are the everyday vocabulary of portfolio construction. None needs stochastic calculus; each is a cold kernel certificate over weights w , per-asset volatilities σ_i , and covariances σ_{ij} (shown at $n = 3$, per-pair, hence general in n). We collect them because the trust-map discipline of this paper applies verbatim: these are tier **D**, not folklore.

Corollary 1 (diversification ratio ≥ 1 ; tier D). $w^\top \Sigma w \leq (\sum_i w_i \sigma_i)^2$, i.e. portfolio volatility never exceeds the weighted average of component volatilities. The gap is $\sum_{i < j} 2w_i w_j \sigma_i \sigma_j (1 - \rho_{ij}) \geq 0$.

The ratio $DR = (\sum_i w_i \sigma_i) / \sigma_p \geq 1$ is the diversification a portfolio is receiving: it equals 1 only under perfect correlation and grows as correlations fall. Maximizing it is Choueifaty’s most-diversified portfolio. The certificate makes the source of the benefit explicit — it is the correlation deficit $1 - \rho_{ij}$, weighted by the volatility-scaled co-holdings.

Corollary 2 (risk-contribution identity; tier D). $w^\top \Sigma w = \sum_i w_i (\Sigma w)_i$. Total portfolio variance decomposes *exactly* into per-asset risk contributions $RC_i = w_i (\Sigma w)_i$, with no residual.

This is Euler’s theorem for the (homogeneous) variance, and it is the entire basis of risk budgeting: “asset i carries RC_i / σ_p^2 of the portfolio’s risk” is a statement with no slack. Equalizing the RC_i is the risk-parity portfolio.

Corollary 3 (decorrelation monotonicity; tier D). Lowering every pairwise covariance lowers portfolio variance: $\sigma'_{ij} \leq \sigma_{ij}$ for all $i < j \Rightarrow w^\top \Sigma' w \leq w^\top \Sigma w$.

The formal justification for hunting low- or negatively-correlated diversifiers: at fixed weights and variances, risk is monotone in the off-diagonal covariances.

Corollary 4 (Elton–Gruber decomposition; tier D). For the equal-weighted portfolio, $\sigma_p^2 = \frac{1}{n} \overline{\sigma^2} + (1 - \frac{1}{n}) \overline{\text{cov}}$, where $\overline{\sigma^2}$ is the mean variance and $\overline{\text{cov}}$ the mean pairwise covariance.

As the number of names grows, the idiosyncratic term $\overline{\sigma^2}/n$ vanishes but the average-covariance term does not: $\sigma_p^2 \rightarrow \overline{\text{cov}}$. This is the exact, machine-checked statement of *why diversification cannot remove systematic risk* — the floor is the average covariance — and it quantifies how many names are needed to approach it.

4.8 Bridges and bounds, made explicit and machine-checked

The derived layer also yields exact relations among the excess growth rate, the diversification ratio, and the Herfindahl index. Maeso and Martellini (2020) study excess-growth-rate maximization as a smart-beta portfolio objective. Here we derive an explicit identity connecting that objective to the diversification-ratio gap and give the relation a machine certificate under the same honesty discipline: tier D, $n = 3$ representative, with per-pair certificates general in n .

A bridge between Fernholz and Choueifaty. Two diversification measures grew up in different literatures: the SPT *excess growth rate* $\gamma_\pi^* = \frac{1}{2}(\sum_i w_i \sigma_{ii} - w^\top \Sigma w)$, and the *diversification ratio* of the most-diversified-portfolio program (Choueifaty and Coignard, 2008; Choueifaty, Froidure and Reynier, 2013), whose squared gap is $(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w$. Maeso and Martellini (2020) study maximizing the excess growth rate as a smart-beta objective. We record the exact algebraic relation between the two objectives, which has a clean cancellation worth stating outright.

Theorem 7 (excess-growth / diversification-ratio identity; tier D). $2\gamma_\pi^* - [(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w] = \sum_{i < j} w_i w_j (\sigma_i - \sigma_j)^2 \geq 0$. Hence $\gamma_\pi^* \geq \frac{1}{2}[(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w]$, with equality iff all per-asset volatilities are equal.

The striking feature is that the *covariances cancel* in the difference: although both quantities depend on Σ , their gap is purely the weighted dispersion of volatilities $\sum_{i < j} w_i w_j (\sigma_i - \sigma_j)^2$. So twice the excess growth always dominates the diversification-ratio gap, and the two coincide exactly when the assets are equal-volatility. This places the SPT and most-diversified-portfolio notions of “diversification benefit” on one ruler.

The bridge is *two-sided*. The dispersion term is itself controlled — each $(\sigma_i - \sigma_j)^2 \leq (\sigma_{\max} - \sigma_{\min})^2$ — so the excess of γ_π^* over half the diversification-ratio gap is bounded by the volatility *range* times the concentration: $\gamma_\pi^* \leq \frac{1}{2}[(\sum_i w_i \sigma_i)^2 - w^\top \Sigma w] + \frac{1}{4}(\sigma_{\max} - \sigma_{\min})^2(1 - H)$. Together with Theorem 7 this brackets the excess growth between half the diversification-ratio gap and that same quantity plus a vol-range-and-concentration correction — the two measures cannot drift apart by more than $\frac{1}{4}(\sigma_{\max} - \sigma_{\min})^2(1 - H)$.

A guaranteed premium from observable inputs. The correlated excess growth admits a floor that a practitioner can evaluate from three numbers — the smallest asset volatility, the largest pairwise correlation, and the portfolio’s concentration.

Proposition 1 (guaranteed-premium floor; tier D). If $\sigma_i \geq \sigma_{\min}$ for all i and $\rho_{ij} \leq \rho_{\max} \leq 1$ for all pairs, then $\gamma_\pi^* \geq \frac{1}{2} \sigma_{\min}^2 (1 - \rho_{\max})(1 - H)$, where $H = \sum_i w_i^2$.

Unless the market is degenerate ($\sigma_{\min} = 0$) or perfectly correlated ($\rho_{\max} = 1$) or fully concentrated ($H = 1$), the rebalancing premium is *strictly* positive and bounded below by a quantity one can read off a covariance estimate.

The correlated concentration sandwich. The diagonal sandwich of Section 4.4 extends to arbitrary covariance once excess growth is written through the pairwise return-difference variances $r_{ij} = \text{Var}(R_i - R_j) = \sigma_{ii} + \sigma_{jj} - 2\sigma_{ij}$. Using $\sum_{i < j} w_i w_j = (1 - H)/2$,

Theorem 8 (correlated concentration sandwich; tier D). With $r_{ij} \in [r_{\min}, r_{\max}]$,
 $\frac{1}{4} r_{\min}(1 - H) \leq \gamma_{\pi}^* \leq \frac{1}{4} r_{\max}(1 - H)$.

The same concentration factor $1 - H$ that governs the equal-weighted premium (Theorem 5) governs the *correlated* premium of an arbitrary portfolio, now scaled by the band of pairwise relative variances rather than by a single variance level.

An accounting identity and a ceiling. Two transparent companions complete the picture. First, excess growth splits exactly into its uncorrelated part minus a covariance drag: $\gamma_{\pi}^* = \gamma_{\text{diag}}^* - \sum_{i < j} w_i w_j \sigma_{ij}$ — positive correlations *eat* precisely the weighted covariance sum out of the premium one would earn if holdings were independent. Second, the premium is capped: $\gamma_{\pi}^* \leq \frac{1}{2} \overline{\sigma^2}$, since portfolio variance is non-negative — excess growth can never exceed half the average individual variance.

Diversification of strategies, not just assets. The diagonal excess-growth functional is concave in the portfolio weights. Blending two portfolios therefore earns at least the average of their premia. The precise gain is the variance-weighted squared dispersion of their weights:

$$\gamma_{\text{diag}}^*\left(\frac{u+v}{2}\right) - \frac{1}{2}[\gamma_{\text{diag}}^*(u) + \gamma_{\text{diag}}^*(v)] = \frac{1}{8} \sum_i \sigma_{ii} (u_i - v_i)^2 \geq 0.$$

This is the rebalancing premium applied one level up: averaging two allocations never sacrifices premium, and strictly improves it whenever the two disagree.

When the assets share a variance level, a direct corollary is that the equal-weighted portfolio maximizes the premium:

$$\gamma_{\text{diag}}^* = \frac{1}{2} s(1 - H) \leq \frac{1}{2} s(1 - 1/n), \quad H \geq 1/n.$$

This gives a clean formal case for equal weighting when nothing distinguishes the assets' variances.

4.9 Decomposition and attribution

The bounds of the previous section say *how large* the premium is. Two further identities say *where it comes from* — one splitting it across a hierarchy of asset groups, the other splitting it across the individual assets. Both are exact, drift-free, and cold-certified (tier **D**, $n = 3$ representative).

Law of total excess growth. Group the assets, e.g. $A = \{1, 2\}$ and $B = \{3\}$, with group weights $W_g = \sum_{i \in g} w_i$. The total premium decomposes into a *between-group* term plus the within-group premia, each carried by its group weight:

$$\gamma_{\text{total}}^* = \gamma_{\text{between}}^* + \sum_g W_g \gamma_{\text{within},g}^*$$

This is the SPT analogue of the law of total variance: the rebalancing premium of a portfolio-of-portfolios is the premium *between* the sleeves plus the weighted premia *inside* them. A direct corollary — since each within-group relative variance is non-negative — is that **nesting never**

destroys premium: $\gamma_{\text{total}}^* \geq \sum_g W_g \gamma_{\text{within},g}^*$. Building a hierarchy of diversified sleeves can only add to the rebalancing premium, never subtract.

This is exactly the structure exploited by hierarchical risk-parity allocators: the identity guarantees the premium budget of a tree of groups is conserved and additive, so a top-down clustering carries an honest accounting of where the growth advantage is generated.

Excess-growth attribution. At the asset level, γ_π^* is a homogeneous-degree-2 form in the weights, so Euler’s theorem gives an exact additive attribution: writing each asset’s marginal contribution $\text{MC}_k = \frac{1}{2} w_k \sum_{j \neq k} w_j r_{kj}$ (its weight times its mean relative variance against the rest),

$$\sum_k \text{MC}_k = 2 \gamma_\pi^*, \quad \text{MC}_k \geq 0 \text{ for every } k.$$

Each asset’s share of the premium is non-negative — no holding drags the rebalancing premium down — and the contributions sum (up to the factor 2 from the degree-2 homogeneity) to the total. This is the premium-side counterpart of the familiar *risk* attribution $\sum_k w_k \partial_k \sigma_p = \sigma_p$: where risk-parity budgets *variance* across assets, this budgets *excess growth*, with the guarantee that the per-asset budget is everywhere non-negative.

5. The two Mathlib seals (tier S)

Two facts in SPT are analytic, not algebraic: they involve a transcendental function or an n -dimensional positive-semidefinite argument, and a scalar real-arithmetic kernel cannot close either. Rather than leave them as bare assertions, we discharge each against Mathlib — re-proving it as a stand-alone, sorry-free Lean 4 theorem that depends only on Mathlib’s three standard axioms (propext, Classical.choice, Quot.sound). The kernel still *states* them as hypotheses; the *truth* now rests on audited mathematics.

5.1 The diversity bound: $0 \leq H(\mu) \leq \log n$

The entropy-weighted portfolio is generated by the Shannon entropy $H(\mu) = -\sum_i \mu_i \log \mu_i$ of the market weights. SPT uses two facts about it: entropy is non-negative, and it is maximized by the uniform market at the value $\log n$. The upper bound is the *diversity bound* — it quantifies how far the market is from a monopoly. In the kernel these are the hypotheses `thm_entropy_nn` and `thm_entropy_max`; the upper bound needs the concavity of $-x \log x$, a transcendental fact no sum-of-squares certificate reaches.

Theorem 9 (diversity bound; tier S, Lean-sealed). For weights $\mu_i \geq 0$ with $\sum_i \mu_i = 1$,

$$0 \leq H(\mu) \leq \log n,$$

with equality at the uniform market $\mu_i = 1/n$.

The Lean proof (`SPT.entropy_bounds` in `lean_lab/LeanLab/SPTEntropyBound.lean`) has two moves. First, the simplex bounds make every entropy summand non-negative; this is Mathlib’s `Real.negMulLog_nonneg`.

Second, Jensen’s inequality is applied to the concave negative-log summand with uniform weights (`concaveOn_negMulLog`, `ConcaveOn.le_map_sum`):

$$\frac{1}{n} \sum_i (-\mu_i \log \mu_i) \leq -\left(\frac{1}{n}\right) \log \left(\frac{1}{n}\right) = \frac{1}{n} \log n,$$

Cancelling the positive uniform factor gives the stated upper bound. The proof is sorry-free; #print axioms reports the three standard axioms and nothing else.

5.2 The conditional trace-model sign: a positive-semidefinite trace inequality

The seal proves a precise conditional statement. Under the paper’s explicit model definition

$$\gamma^* = \frac{1}{2} \operatorname{tr}((-D^2 \log G) M), \quad M = \operatorname{diag}(\mu) \sigma \operatorname{diag}(\mu).$$

log-concavity of G means $-D^2 \log G \succeq 0$. The weighted covariance $M \succeq 0$, being a congruence of the valid relative covariance $\sigma \succeq 0$. Conditional on the displayed identification of γ^* , its sign is therefore the sign of **the trace of a product of two positive-semidefinite matrices**.

Theorem 10 (PSD-trace inequality; tier S, Lean-sealed). For real symmetric matrices $A, B \succeq 0$,

$$0 \leq \operatorname{tr}(AB).$$

Corollary (conditional trace-model sign). Log-concavity and positive-semidefiniteness of the weighted covariance give

$$-D^2 \log G \succeq 0, \quad M = \operatorname{diag}(\mu) \sigma \operatorname{diag}(\mu) \succeq 0,$$

and therefore

$$\gamma^* = \frac{1}{2} \operatorname{tr}((-D^2 \log G) M) \geq 0.$$

The Lean proof in `SPTEExcessGrowth.lean` mechanizes the textbook congruence argument. Let the positive-semidefinite square root of the first matrix be S ; Mathlib’s functional-calculus square root also supplies its Hermitian property. Appendix B records the exact theorem handles.

By cyclicity of the trace,

$$\operatorname{tr}(AB) = \operatorname{tr}(SSB) = \operatorname{tr}(SBS) = \operatorname{tr}(SBS^H).$$

The final matrix is a positive-semidefinite congruence of the second matrix, so its trace is non-negative by Mathlib’s trace-positivity theorem. The proof is sorry-free and depends only on the three standard axioms.

The matrix implication is checked in full dimension-free generality. Its scope is deliberately narrower than an unconditional theorem about every standard functionally generated portfolio: the seal covers the sign of γ^* **if** the paper’s trace-model definition $\gamma^* = \frac{1}{2} \operatorname{tr}((-D^2 \log G) M)$ holds, G is log-concave, and σ is a valid covariance. The trace identification remains tier **A**. What is beyond doubt is the conditional implication from these explicit premises.

6. Relative arbitrage and the Atlas premia (conditional results)

SPT’s most striking promises — *relative arbitrage* (beating the market with probability one over a long enough horizon) and the *small-cap* and *rebalancing* premia of the Atlas rank model — live one tier down in trustworthiness. They are real theorems *conditional on* model hypotheses, and we present them as such. Each is a kernel theorem whose proof is sound but whose premises are tier **A**.

6.1 Diversity and relative arbitrage

A market satisfies a *diversity* condition if no single stock comes to dominate: the maximum weight stays bounded below one, $\mu_{\max} \leq 1 - \delta$ for some $\delta > 0$ (thm_no_monopoly, thm_diverse_bounded). Under diversity the diversity-weighted portfolio has a strictly positive excess growth rate $\gamma_p^* > 0$ (thm_div_growth_pos), and over a horizon long enough that the integrated growth dominates the bounded leakage term, the relative value exceeds one:

Proposition 2 (relative arbitrage, conditional; tier A-premised). Under the diversity condition, there is a horizon T beyond which $V_\pi(T)/V_\mu(T) \geq 1$ — the portfolio beats the market.

The kernel derives this cleanly (thm_long_run_beats, thm_margin_pos) from the master formula plus a leakage bound — but both inputs are tier **A**. The honest statement is: *given* the master-formula decomposition and a bounded-leakage diversity regime, long-run outperformance follows. The “with probability one” flavor of the textbook theorem is a property of the continuous-time model we do not re-establish; the kernel works with the deterministic margin inequality.

6.2 The Atlas rank model: small-cap and rebalancing premia

The Atlas model is a rank-based market in which the smallest stocks have the highest growth rates, producing the empirically-motivated small-cap premium. The kernel encodes the rank structure (thm_rank_order, thm_rank_mean_revert, thm_growth_gap) and derives positivity of the premia by linear arithmetic from sign hypotheses:

Proposition 3 (Atlas premia, conditional; tier A-premised). In the rank model, the small-cap premium and the rebalancing premium are positive: thm_scp_pos, thm_rebal_pos.

These follow validly from the Atlas hypotheses (ergodicity with mixing-time bound thm_ergodic, Pareto tail $\alpha > 1$ thm_pareto_gt1, positive rank-volatility thm_atlas_vol), but those hypotheses are *posited* features of the rank model, not derived. Of the two premia, only the rebalancing premium has a tier-**D** anchor: it equals the half-variance gap $(\bar{\sigma}^2 - \sigma_{\text{mkt}}^2)/2$ of Theorem 1, which *is* derived. The small-cap premium does not — and Section 7 shows it can be negative in realized terms.

7. Empirical validation

The formal layer proves a set of inequalities and identities. Two distinct empirical questions follow. First — are the proven inequalities *vacuous*? An inequality can be true with zero margin, or true only for inputs no real market produces. Section 7.1 checks every load-bearing theorem against twenty-four years of real equity data and finds the margins economically large. Second — does a

non-negative growth *rate* guarantee a realized *result*? Section 7.2 shows, by direct simulation, that it does not. The first section grounds the theorems; the second guards their interpretation.

7.1 The inequalities hold on real data — and are not vacuous

We tested the load-bearing theorems of Section 4 against a real equity cross-section: the **Ken French 49-Industry Portfolios** (value-weighted monthly returns), with industry market capitalizations reconstructed from the reported number of firms times average firm size, over **January 2000 – December 2023** (288 months). The covariance matrix is the sample covariance of monthly returns, annualized; the market portfolio uses the time-averaged capitalization weights. The study is fully reproducible with no API key from `empirical_validation.py`. The exact upstream ZIP is frozen locally and bound to SHA-256 `a0b23457eac619c8a3cce362de563b6f57acc3514779ceccdb99886edfa0a804`; the extracted CSV is independently bound to SHA-256 `1aa2bd245feb954362bc91a68ceea21c2020e08b0c194a966de20606c1aaa51`. A machine-readable receipt records the source URL, byte counts, analysis window, archive member, and both digests, so future replays fail closed instead of silently accepting revised source data.

Each theorem of Section 4 maps to a checkable number. All nine checks pass:

Quantity (real data)	Value	Governing theorem	Holds
Herfindahl H , market-cap	0.052 (≈ 19 effective holdings)	$1/n \leq H \leq 1$ (§4.5)	
Excess growth γ^* , market-cap	1.06%/yr	$\gamma^* \geq 0$ (§4.6)	
Excess growth γ^* , equal-weight	1.51%/yr	$\gamma^* \geq 0$ (§4.6)	
trace-form vs pairwise-form γ^*	agree to machine precision	identity (§4.6)	
Diversification return (EW)	1.51%/yr = γ_{EW}^*	$\frac{1}{2}(\bar{\sigma}^2 - \sigma_p^2) = \gamma^*$ (§4.5)	
Excess-growth ceiling, market	$1.06\% \leq 2.34\%$	$\gamma^* \leq \frac{1}{2} \sum_i w_i \sigma_{ii}$ (§4.8)	
Diversification ratio, market / EW	1.32 / 1.38 (> 1)	$DR > 1 \Rightarrow \gamma^* > 0$ (§4.8)	

Two features matter. First, the inequalities are **not tight against zero**: the equal-weight excess growth rate is a full 1.5% per year of structural drift, and even the capitalization-weighted market portfolio carries 1.06%/yr — the rebalancing premium is economically real on real data, not a rounding artifact. Second, the two *identities* — the trace-form/pairwise-form equality for γ^* , and the diversification-return identity $\frac{1}{2}(\bar{\sigma}^2 - \sigma_p^2) = \gamma_{EW}^*$ — hold to machine precision, exactly as the kernel proves them symbolically. The general- n sign theorem of Section 4.6 is here instantiated at $n = 49$ and confirmed. Both portfolios in this data set have $DR > 1$ and positive γ^* , consistent with the proved one-way implication $DR > 1 \Rightarrow \gamma^* > 0$; the converse is not claimed.

A real number for the honesty boundary. The *realized* log-growth of the equal-weight portfolio over the cap-weighted market across the full window is 2.03%/yr. Of this, the pure variance-driven SPT component — the excess-growth gap $\gamma_{EW}^* - \gamma_{\text{market}}^*$ — accounts for only about 0.45%/yr. The remaining $\sim 1.6\%$ /yr is the difference in the underlying industries' drifts (equal-weighting tilts

toward smaller, higher-mean industries over this period), which SPT’s rate guarantee says nothing about. This is the same boundary Section 7.2 draws with simulation, now visible in real data: the excess growth rate is a genuine, measurable component of realized performance, but it is not the whole of it.

7.2 A rate is not a result

Under the stated covariance and model premises, the formal layer establishes that the excess growth *rate* is non-negative. This section establishes, by direct simulation, what the formal layer pointedly does **not** say: a non-negative rate does not guarantee realized outperformance over a finite horizon. This is the paper’s honest boundary, and we put a number on it. All figures are from `numerical_validation.py` (seed 42); they are illustrative outputs of a chosen market, not empirical estimates.

Setup. A synthetic market is simulated and three portfolios are tracked against the capitalization-weighted index: equal-weighted, entropy-weighted, and diversity-weighted ($p = 0.5$). Each is run over 500 sample paths of five years.

A separate rate diagnostic behaves as proven. Before the path experiment, the script runs an independent deterministic ten-stock, diagonal-covariance calibration. For equal weights $w_i = 1/10$, it obtains $\gamma_{EW}^* = 9.59\%$ per year from $\frac{1}{2}(\overline{\sigma^2} - \sigma_{EW}^2)$. This number is **not** estimated from the twenty-stock, 500-path experiment below; it is a separate algebra check of the diagonal formula.

The realized result does not follow. Over the 500 paths:

Portfolio	Paths beating the market	Mean relative value V_π/V_μ
Diversity-weighted ($p = 0.5$)	53.6%	1.0058
Entropy-weighted	57.6%	1.0036

The corrected self-financing log-wealth simulation produces a modest majority of winning paths and mean relative values only slightly above one. Roughly 42–46% of paths still lose to the market. Thus a non-negative structural component does not imply pathwise or high-confidence finite-horizon outperformance. The boundary term $\log G(\mu(T)) - \log G(\mu(0))$ can offset the integrated growth on individual paths. The master-formula decomposition makes this transparent: realized log relative value is *leakage* plus *integrated growth*, and only the latter carries the sign guarantee.

The small-cap premium can invert. In the Atlas simulation the realized small-cap premium came out **negative**, $SCP \approx -0.225$ per year, even though Proposition 3 makes it positive *in the model*. This is the cleanest illustration of the tier boundary: a conditional theorem (tier **A**) whose premises the finite simulation does not honor.

The takeaway is not that SPT is wrong — it is exactly right about the rate. The takeaway is that the rate is a *favorable drift*, and a practitioner who reads “non-negative excess growth” as “beats the market” is reading a tier-**D/S** statement as a guarantee it never made.

The distributional picture behind the table is the point: realized relative value V_π/V_μ remains a cloud straddling the break-even line $V = 1$. The positive excess-growth integral pushes relative value upward, while the path-dependent boundary term can offset it. Over a five-year horizon the positive term does not decide every outcome. The full distribution is reproducible from `numerical_validation.py` (seed 42).

Where the structural benefit actually shows up. It is worth comparing the functionally generated portfolios against the standard alternatives a practitioner would consider, including minimum-variance. In a second simulation (20 stocks, 10 years, 200 paths) the picture is consistent and instructive:

Strategy	Mean 10-year log-return	Volatility (std)	Mean/std	Annualized mean difference vs market
Market (cap-weighted)	−0.0759	0.402	−0.189	—
Equal-weighted	−0.0415	0.253	−0.164	+0.34%/y
Diversity- weighted ($p = 0.5$)	−0.0434	0.297	−0.146	+0.33%/y
Minimum- variance (diagonal model)	−0.0313	0.201	−0.156	+0.45%/y

All three alternatives improve the mean log-return relative to the market in this particular simulation, while the functionally generated portfolios also cut dispersion materially (0.25–0.30 versus 0.40). The minimum-variance row now uses the exact diagonal-model solution $w_i \propto 1/\sigma_i^2$, not inverse-volatility weights. Every row accumulates self-financing log wealth with the quadratic-variation correction $-\frac{1}{2} \sum_i (w_i \sigma_i)^2 dt$. These are illustrative finite-sample outcomes, not a theorem of universal outperformance; the practitioner’s decision remains a risk-budget decision, not a “free alpha” decision.

8. The one remaining assumption

If the reader takes one structural fact from this paper, let it be this: after the work of Sections 4 and 5, the SPT framework rests on a **single** mathematical assumption — the master formula of Section 3.1 — and that assumption is an *identity*, not an inequality. Everything else is either derived (tier **D**) or sealed (tier **S**).

8.1 Why it cannot be discharged today

The master formula is Itô’s lemma applied to $\log G(\mu(t))$. To prove it inside a formal system one needs the stochastic integral, the quadratic-variation calculus, and Itô’s formula for continuous semimartingales. We checked the state of the art directly:

- **Mathlib (v4.28)** — the library against which Sections 5.1 and 5.2 are sealed — contains martingales, conditional expectation, and a general theory of stochastic processes, but **no stochastic integral, no Itô formula, and no Brownian motion** as named results. A search for an Ito-named theorem returns nothing.
- An **in-house** stochastic-calculus encoding exists in our kernel, but its “Itô lemma” is itself posited as a structural condition, not derived from measure theory. Discharging the master formula against it would be circular — moving the assumption, not removing it.

So we do not pretend. The master formula stays tier **A**, named once, used where needed, never dressed as a theorem.

8.2 The discharge path

There is a concrete route, and it is worth recording — and the relevant external frontier is moving unusually fast. The work is the formal—mathfin project (Coelho, arXiv:2606.15089 and 2606.01356), the first machine-checked Itô calculus in any proof assistant, which builds directly on the present author’s earlier discrete-Itô formalization. At the arXiv snapshot the Itô formula was *scalar*, Brownian-driven, and restricted to bounded-derivative C^3 functions on $[0, T]$, with the unrestricted- C^2 case (stopping-time localization) and integration against an Itô process both named as out of scope. As of this writing (mid-2026) the public development has already closed several of these: a localized, unrestricted Itô formula (the stopping-time “Summit C”) and an Itô formula against a constant-coefficient Itô process are now in the repository.

Against that moving frontier, what the SPT master formula still requires is two things, both on the project’s active roadmap rather than its completed core:

1. **A multidimensional Itô formula.** $\mu(t)$ is the n -vector of market weights; the development currently provides the scalar formula and a two-dimensional Itô lemma, with the move “from scalar to vector state space” listed as a future direction.
2. **Variable-coefficient integrands.** The weights $\mu_i(t)$ have state-dependent dynamics, whereas the available Itô-process formula is the constant-coefficient case; the general (variable-coefficient) integrand is the next gated step.

The localized formula addresses the boundary singularity of $\log G$, but the certifier exposed a more basic distinction. Itô’s formula for $\log G(\mu_t)$ proves a decomposition of the generator. The Fernholz master formula additionally identifies a self-financing portfolio and proves that its relative log-wealth equals the generator boundary term plus the correct drift process. The scalar Lean developments in Sections 8.4–8.5 establish the first statement only. They are valuable stochastic-calculus certificates, but they do not close the portfolio layer. The general master formula therefore remains tier **A** in this paper.

8.3 The two-asset control calculation: portfolio excess growth is $\tau/8$

The cleanest way to prevent the generator drift and the portfolio rate from being conflated is to calculate the latter directly. Let $\Sigma = (\sigma_{ij})_{i,j=1}^2$ be the instantaneous return-covariance matrix and let the equal-weight portfolio have $\pi = (1/2, 1/2)$. Its standard excess-growth rate is half the component-weighted average variance minus the portfolio variance:

$$\gamma_{\text{EW}}^* = \frac{1}{2} \left[\frac{\sigma_{11} + \sigma_{22}}{2} - \frac{\sigma_{11} + 2\sigma_{12} + \sigma_{22}}{4} \right].$$

Writing $\tau := \text{Var}(R_1 - R_2) = \sigma_{11} + \sigma_{22} - 2\sigma_{12}$ gives the exact identity

$$\gamma_{\text{EW}}^* = \frac{\tau}{8} \geq 0.$$

The canonical proof source derives the equality from the two displayed covariance definitions by exact real arithmetic. If $\tau > 0$, then $\gamma_{\text{EW}}^* > 0$; if $\tau = 0$, it vanishes. This value is independent of

the current market weights. Consequently, neither $\frac{1}{4}(\mu_1^2 + \mu_2^2)\tau$ nor $\frac{1}{2}\mu_1\mu_2\sigma^2$ can be identified with this portfolio excess-growth rate without an additional, explicitly proved change-of-variables and self-financing bridge.

8.4 The two-asset generator-side Itô decomposition

The bounded milestone of Section 8.2 is done. Reparametrise the two weights by their log-odds. The weights and equal-weight generator then have the closed forms

$$y = \log(\mu_1/\mu_2), \quad \mu_1 = \frac{e^y}{1 + e^y}, \quad \mu_2 = 1 - \mu_1,$$

and

$$h(y) := \log G(\mu(y)) = \frac{1}{2}y - \log(1 + e^y).$$

The logistic derivative identity makes all three derivatives polynomial in the first weight:

$$h'(y) = \frac{1}{2} - \mu_1, \quad h''(y) = -\mu_1\mu_2, \quad h'''(y) = -\mu_1\mu_2(1 - 2\mu_1).$$

Thus the function is three-times continuously differentiable with all three derivatives bounded by one, so it lies inside Coelho's at-most-exponential-growth hypothesis. For a constant-coefficient log-odds diffusion, `MathFin.ito_formula_itoProcess` yields, almost surely,

$$h(y_T) - h(y_0) = (\text{Itô integral}) + \int_0^T \left(h'(y_s) b + \frac{1}{2} h''(y_s) \sigma^2 \right) ds.$$

The negative of the Itô correction is strictly positive for non-zero volatility:

$$-\frac{1}{2} h''(y) \sigma^2 = \frac{1}{2} \mu_1 \mu_2 \sigma^2 > 0.$$

This decomposition is verified in `lean_ito_bridge/ItoBridge/SPTMasterN2.lean`, built against `formal-mathfin` in a separate Lean v4.31 project. Its native declaration names predate the present claim audit; the formal import registry now exposes it under the honest local name `log_generator_ito_decomposition_n2`. The axiom audit reports `{propext, Classical.choice, Quot.sound}` and no `sorryAx`.

Scope, stated as plainly as the rest of the paper. This theorem contains $h(y_T) - h(y_0)$, a stochastic integral, and the drift of $h(y_t)$. It contains no portfolio holdings, no self-financing wealth process, and no equality involving $\log(V_\pi/V_\mu)$. Therefore it does not prove a master formula or portfolio outperformance, even in this constant-coefficient two-asset model. Turning it into such a result requires a separate theorem connecting the generated trading strategy to relative wealth and reconciling the covariance clock with the log-odds volatility. That bridge remains open.

8.5 The three-asset rank-1 generator decomposition and the remaining bridge

The same distinction persists with three assets. A rank-1 constant-coefficient market is still driven by one scalar Brownian motion, so the generator can again be checked with a scalar Itô theorem.

The observation is that a **single-factor (rank-1)** constant-coefficient market with any number of assets is again a scalar problem. If every log-price is driven by the same Brownian motion, then every market weight and the log-generator are deterministic functions of that one scalar driver. The checked three-asset theorem uses the narrower zero-drift, zero-initial-log-price model $Y_0 = (0, 0, 0)$ (equivalently the Brownian driver starts at $B_0 = 0$); it is not a general constant-drift statement. Coelho's scalar Itô formula therefore applies without multidimensional machinery.

For three assets, choose antisymmetric loadings $\sigma = (1, 0, -1)$ and the exact paths $Y_1(t) = B_t$, $Y_2(t) = 0$, $Y_3(t) = -B_t$. Define the normalizing sum and equal-weight log-generator by

$$S(x) = e^x + 1 + e^{-x}, \quad \Psi(x) = -\log S(x).$$

The identity $e^x e^{-x} = 1$ gives

$$\Psi'(x) = -\mathbb{E}_\mu[\sigma], \quad \Psi''(x) = -\text{Var}_\mu[\sigma] < 0.$$

Coelho's unrestricted formula (`MathFin.ito_formula_unrestricted_local`) then yields the local-martingale decomposition

$$\log G(\mu_T) - \log G(\mu_0) = (\text{local martingale}) - \frac{1}{2} \int_0^T \text{Var}_{\mu_s}[\sigma] ds.$$

The quantity $-\frac{1}{2}\Psi''$ is strictly positive, but the displayed theorem is still a decomposition of $\log G(\mu_t)$, not a statement about a self-financing portfolio. It is verified in `lean_ito_bridge/ItoBridge/SPTMasterN3.lean`; the formal import registry exposes the declaration under `log_generator_ito_decomposition_n3_rank1`. The axiom audit again reports `{propext, Classical.choice, Quot.sound}` and no `sorryAx`.

The exercise identifies two separate walls. A genuinely multi-factor generator decomposition still needs stochastic integration against random adapted integrands. Even in the scalar rank-1 setting, however, a portfolio theorem also needs the self-financing strategy-to-wealth bridge. The old wording collapsed these two obligations into one and therefore overstated what the Lean types proved. The corrected trust map keeps both obligations explicit.

9. Discussion

9.1 What a practitioner should take away

SPT is one of the rare quantitative theories whose central qualitative claims survive mechanical verification essentially intact. The algebra of diversification is exact (tier **D**). The two analytic load-bearing inequalities are now checked against an audited library (tier **S**). What does *not* survive — and never claimed to — is the leap from a non-negative growth *rate* to realized outperformance. A practitioner can deploy functionally generated portfolios knowing the drift is genuinely in their

favor and genuinely structural, while sizing positions and horizons against the fact that leakage and noise routinely overwhelm that drift over five-year windows.

9.2 What the formalization adds

The contribution is not a theorem Fernholz did not have; it is a *trust-map* Fernholz’s readers did not have. Three things are now different. The excess-growth sign is a checked PSD-trace theorem rather than a “clearly” step. The diversity bound is a checked Jensen consequence. And the entire edifice has been reduced to one explicit, named assumption with a known discharge path — which is a far stronger statement than “the proofs look right.” A field that can say exactly where its one assumption lives is in better epistemic health than one that cannot.

9.3 Limitations

- The master formula remains assumed (Section 8). The two scalar Lean developments certify generator-side Itô decompositions, not self-financing portfolio identities. The remaining formal work has two parts: the strategy-to-relative-wealth bridge even in scalar models, and random-adapted-integrand stochastic calculus for genuinely multi-factor models.
- Most tier-**D** algebra is proved for the $n = 3$ representative with a per-summand certificate that carries to general n by inspection; the one exception is the *correlated excess-growth sign*, which Section 4.6 proves outright for all n (two $\forall N$ finite-sum theorems). The general-covariance sign is therefore covered twice — by the all- n derived theorems and by the tier-**S** PSD-trace seal — established by different means; we do not claim a single unified kernel derivation across the two routes.
- The Atlas premia are conditional on rank-model hypotheses; the small-cap premium has no tier-**D** anchor and inverts in simulation.
- The empirical section (7.1) confirms the proven inequalities hold with economically large margins on 24 years of real industry data, but it is a *validation* of the theorems on one historical cross-section — sample covariance, time-averaged weights — not an out-of-sample statistical study with standard errors or alternative estimators. The finite-horizon simulation (7.2) is illustrative — a chosen synthetic market, one seed — not an empirical estimate.

9.4 Outlook

The next formal target is the self-financing bridge in the two-asset constant-coefficient model: define the generated holdings, prove the wealth equation, and show that its finite-variation term matches the standard covariance rate $\tau/8$. Only after that scalar bridge closes does it make sense to extend the result to random adapted integrands and genuinely multi-factor markets. The same three-tier discipline — derive what is algebra, seal what is analysis, name what is modeling — remains a template for auditing other “obviously true” quantitative theories.

Acknowledgments

The formal development uses the Lean 4 theorem prover and the Mathlib library. The author thanks the Mathlib community, whose theorems on convexity (`negMulLog`, `Jensen`) and positive-semidefinite matrices (the continuous functional calculus square root, congruence closure, and trace positivity) are what make the tier-**S** seals possible.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used AI-based tools for manuscript drafting, literature search, symbolic computation verification, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Booth, David G. and Fama, Eugene F. (1992). Diversification Returns and Asset Contributions. *Financial Analysts Journal*, 48, 26–32. DOI: 10.2469/faj.v48.n3.26
- Choueifat, Yves and Coignard, Yves (2008). Toward Maximum Diversification. *The Journal of Portfolio Management*, 35, 40–51. DOI: 10.3905/JPM.2008.35.1.40
- Choueifat, Yves, Froidure, Tristan and Reynier, Julien (2013). Properties of the most diversified portfolio. *The Journal of Investment Strategies*, 2, 49–70. DOI: 10.21314/jois.2013.033
- Coelho, Raphael (2026). A Machine-Checked Itô Calculus for Brownian Motion. *arXiv preprint arXiv:2606.15089*. <https://arxiv.org/abs/2606.15089>
- Coelho, Raphael (2026). A Formally Verified Library of Mathematical Finance in Lean 4. *arXiv preprint arXiv:2606.01356*. <https://arxiv.org/abs/2606.01356>
- Elton, Edwin J. and Gruber, Martin J. (1977). Risk Reduction and Portfolio Size: An Analytical Solution. *The Journal of Business*, 50, 415–437. DOI: 10.1086/295964
- Fernholz, E. Robert (2002). *Stochastic Portfolio Theory*. Springer. DOI: 10.1007/978-1-4757-3699-1
- Fernholz, Robert, Garvy, Robert and Hannon, John (1998). Diversity-Weighted Indexing. *The Journal of Portfolio Management*, 24, 74–82. DOI: 10.3905/jpm.24.2.74
- French, Kenneth R. (2024). 49 Industry Portfolios. Kenneth R. French Data Library.
- Karatzas, Ioannis and Fernholz, Robert (2009). Stochastic Portfolio Theory: an Overview. *Handbook of Numerical Analysis*, 15, 89–167. DOI: 10.1016/S1570-8659(08)00003-3
- Karatzas, Ioannis and Ruf, Johannes (2017). Trading strategies generated by Lyapunov functions. *Finance and Stochastics*, 21, 753–787. DOI: 10.1007/s00780-017-0332-8
- Maeso, Jean-Michel and Martellini, Lionel (2020). Maximizing an equity portfolio excess growth rate: a new form of smart beta strategy? *Quantitative Finance*, 20, 1185–1197. DOI: 10.1080/14697688.2020.1729398
- The mathlib Community (2020). The Lean mathematical library. *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP)*, 367–381. DOI: 10.1145/3372885.3373824

Appendix A: Index of formal statements

All 100 paper-facing kernel statements are rendered, anchor-by-anchor, from the kernel source in the companion file `formal_statements.md`. Internal algebraic helper lemmas remain in the proof environment but are intentionally excluded from this reader-facing index. Each theorem cited in the body carries a `formal_ref` anchor into that file. The tier labels in this paper are not part of the kernel render; they are the author’s epistemic classification, defined in the “Epistemic status” table.

Appendix B: The Lean 4 / Mathlib bridges

Three hand-authored Lean files discharge tier-**S** content; all compile under Lean 4.28.0 / Mathlib v4.28.0, contain no sorry, and depend only on {propext, Classical.choice, Quot.sound}.

- `lean_lab/LeanLab/SPTEntropyBound.lean` proves the diversity bound via Jensen’s inequality. Provenance: `SPTEntropyBound.stamp.json`.
- `lean_lab/LeanLab/SPTEccessGrowth.lean` proves the PSD-trace inequality and the conditional log-concave trace-model sign. Provenance: `SPTEccessGrowth.stamp.json`.
- `lean_lab/LeanLab/SPTMasterFormulaN2.lean` proves the $n = 2$ drift coefficient $x^2\varphi''(x) = -\frac{1}{2}$ for $\varphi = \frac{1}{2}\log$, discharging the calculus premise of the §8.3 reduction. Provenance: `SPTMasterFormulaN2.stamp.json`.

Appendix C: Reproduction

- Kernel proofs: `python3 proofs/stochastic_portfolio_theory/spt_foundations_proof.py` (and `spt_arbitrage_proof.py`, `stochastic_portfolio_theory_proof.py`). Run the three scripts from the repository root for a fresh replay.
- Empirical validation: `python3 proofs/stochastic_portfolio_theory/empirical_validation.py` replays against the hash-verified frozen input. `--freeze-input` is a one-time bootstrap action and refuses to overwrite a different frozen archive.
- Lean seals: from `lean_lab/`, run `lake env lean` on each of the three files listed in Appendix B.
- Lean generator-side Itô decompositions (historical native declaration names retained): `lean_ito_bridge/ItoBridge/SPTMasterN2.lean` and `lean_ito_bridge/ItoBridge/SPTMasterN3.lean`.
- Numerical study: `python3 proofs/stochastic_portfolio_theory/numerical_validation.py` (seed 42).
- Real-data validation: `python3 proofs/stochastic_portfolio_theory/empirical_validation.py`.