

Common Pricing, Different Portfolios

Why Horizon and Bequest Change Allocation Without Changing Equilibrium Pricing

A mode-space theory of equilibrium benchmarks and investor-specific fin_harvestability

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft • March 2026

Paper Role in the Research Program

This paper is the **canonical benchmark-side root paper** in the current finance architecture.

- It is the source of truth for the common-benchmark theorem, the pricing-allocation equivalence, and the market-as-benchmark interpretation in mode space.
- `topics/fin_harvestability/paper.md` is the companion root paper on the investor-specific side: it defines `fin_harvestability`, derives the horizon filter, and develops the lifecycle extension.
- The core bridge claim of the research program is obtained by reading these two papers together:
 - this paper supplies the common benchmark layer,
 - the `fin_harvestability` paper supplies the investor-specific filtering layer.
- `topics/fin_pricing_is_allocation/submission.md` is the reviewer-facing bridge artifact derived from this benchmark paper plus imported `fin_harvestability` consequences.

Executive Summary (Non-Technical)

Modern finance usually tells two different stories. **Equilibrium pricing** says the market portfolio is the efficient benchmark. **Lifecycle investing** says investors with different ages, horizons, mortality profiles, or bequest motives should hold different portfolios. Those claims look inconsistent: if the market is the right benchmark, why should the young, the old, the selfish, and the dynastic hold different things?

This paper argues that the conflict is largely artificial. In a mode decomposition of returns, equilibrium pricing fixes a **common fully harvested benchmark**. Once returns are resolved into orthogonal modes, the pricing restriction and the mean-variance first-order condition become inverse readings of the same identity. In the scalar case, the optimizer reproduces the market benchmark mode by mode.

The economic heterogeneity appears one layer above that benchmark. When this static equilibrium object is paired with companion lifecycle results from the same research program, investor-specific holdings are obtained by filtering the common benchmark through `fin_harvestability`, Bayesian caution, and safety constraints. Different investors therefore do not need different pricing theories

in order to hold different portfolios. They can share one equilibrium pricing law and disagree only in **what fraction of each priced mode they can realistically harvest**.

That is the paper’s strongest claim. It does not say that finance is solved, or that every orthogonal basis is economically meaningful, or that the scalar pricing law has already been derived from all deeper primitives. What the current formal spine shows is narrower but still important: several standard economic routes generate a common mode-pricing family, the scalar benchmark appears as a collapse case of that family, and the market portfolio can be interpreted as the fully harvested equilibrium object against which lifecycle tilts should be read.

If this view is correct, several familiar distinctions become cleaner. Passive investing becomes the common benchmark. Active management becomes a claim about mode-level alpha. Lifecycle tilts become investor-specific harvesting filters rather than anti-equilibrium behavior. Samuelson-style horizon independence becomes a limit case of sufficiently long effective horizons rather than the universal baseline.

Abstract

We propose a mode-space formulation in which equilibrium pricing and portfolio choice are organized around a common benchmark rather than treated as separate layers. Let returns be resolved into orthogonal modes c_k with per-mode premiums π_k , variances $v_k = \text{Var}(c_k)$, and market weights w_k^{mkt} . The strongest static restriction used in the main theorem is

$$\pi_k = \lambda w_k^{\text{mkt}} v_k$$

and the mean-variance first-order condition

$$w_k^* = \frac{\pi_k}{\lambda v_k}$$

are inverse readings of the same identity, implying the fully harvested benchmark allocation

$$w_k^\infty = w_k^{\text{mkt}}.$$

The static theorem therefore says that, in equilibrium, the market portfolio is the common benchmark exposure in mode space.

Paired with companion lifecycle results from the same research program, this benchmark admits the intertemporal reading

$$w_k^{\text{life}} = w_k^\infty \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k = w_k^{\text{mkt}} \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k,$$

where $\mathbb{E}[h_k]$ is expected fin_harvestability at the investor’s effective horizon, b_k is a Bayesian caution term, and s_k is a safety multiplier. The central economic implication is that investors can share one equilibrium pricing law and still rationally hold different portfolios, because what differs across investors is not necessarily priced return but fin_harvestability.

The formal spine is machine-verified in Lean 4. It establishes mode-wise variance and premium decomposition, diagonal Markowitz optimality, the pricing-allocation equivalence, and derived CCAPM, SDF, and Arrow-Debreu routes into a broader family of mode-specific prices of risk λ_k , with explicit bridge theorems for when that family collapses to the scalar case. The paper’s boundary is equally clear: it does not yet identify the canonical market basis, prove universal scalar

collapse, or claim empirical validation. Its contribution is structural and economic at once: it identifies the common equilibrium benchmark and shows how lifecycle, horizon, and bequest effects can be interpreted as investor-specific harvesting filters layered on top of that benchmark.

1. Introduction

1.1 The Real Tension

Financial economics inherited two powerful but uncomfortable ideas.

Asset-pricing theory asks: what return should a risky asset earn in equilibrium? CAPM, APT, intertemporal pricing models, and stochastic discount factor methods all address that question. Their output is a pricing relation.

Portfolio theory asks: how should an investor allocate capital? Markowitz, Black-Litterman, risk parity, and related methods take expected return and covariance as inputs and output optimal weights.

Lifecycle finance then adds a third pressure. Horizon, mortality, labor income, wealth floors, and bequest motive are all supposed to matter for actual portfolio choice. Young and old investors should not hold the same thing. Selfish and dynastic investors should not hold the same thing. Near-ruin and comfortably funded investors should not hold the same thing.

Put bluntly, the field seems to say two conflicting things at once. Equilibrium language says the market portfolio is the efficient benchmark. Lifecycle language says the right portfolio depends on the investor. If both claims are taken literally and in the same coordinates, they appear to collide.

This paper argues that the collision is mostly a modeling artifact. The two claims operate on different margins. Equilibrium pricing fixes a common benchmark. Investor-specific horizon, bequest, mortality, learning, and safety constraints determine how much of that benchmark is actually harvestable.

That distinction is the paper's central economic payoff.

1.2 The Common Benchmark

The static theorem starts from one mode-space equation:

$$\pi_k = \lambda w_k^{\text{mkt}} v_k.$$

Combined with the mean-variance first-order condition

$$w_k^* = \frac{\pi_k}{\lambda v_k}.$$

it yields

$$w_k^\infty = w_k^{\text{mkt}}.$$

This is the benchmark statement of the paper. In the scalar-collapse regime, the market portfolio is the fully harvested equilibrium allocation in mode space. The market is not merely an observed

aggregate. It is the benchmark exposure that the optimizer recovers once the common pricing law is imposed.

The deeper Lean spine now shows that this scalar equation is not the only possible object. A weaker route-level statement is

$$\pi_k = \lambda_k w_k^{\text{mkt}} v_k,$$

with a mode-specific price of risk λ_k . The scalar equation is the special case in which the route-specific family collapses to one common λ .

The point is therefore not that every model already lives in the scalar regime. The point is that the static benchmark can now be located inside a more disciplined hierarchy: universal mode pricing first, scalar collapse second.

1.3 From Benchmark to Heterogeneity

The economic question is what to do with the benchmark once it is identified.

Companion lifecycle results from the same research program imply the multiplicative allocation formula

$$w_k^{\text{life}} = w_k^\infty \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k = w_k^{\text{mkt}} \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k,$$

where $\mathbb{E}[h_k]$ is expected fin_harvestability at the investor's effective horizon, b_k is a Bayesian caution term, and s_k is a safety multiplier.

Substituting the benchmark identity gives

$$w_k^{\text{life}} = w_k^{\text{mkt}} \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k.$$

This is the larger paper-level message. Investors can share one equilibrium pricing law and still rationally hold different portfolios, because they differ in fin_harvestability, not necessarily in beliefs about the priced object itself.

The interpretation is especially sharp for bequest. A selfish finite-horizon investor scales down slow modes because they are weakly harvestable over the investor's effective horizon. A strong-bequest or dynastic investor extends effective horizon, pushing fin_harvestability toward one and recovering the Merton or market benchmark more closely. Samuelson-style horizon independence is therefore not the universal baseline. It is a limiting case of sufficiently long effective horizon.

1.4 Main Claim and Supporting Claims

The paper has one main claim and one immediate economic implication:

Expected return is not an external input to allocation. It is the dual variable implied by equilibrium structure in mode space.

And the key implication is:

Investors can share one equilibrium pricing law and still rationally hold different portfolios, because what differs across investors is not necessarily priced return but fin_harvestability.

Everything else supports that statement.

1. orthogonal mode decomposition diagonalizes risk accounting and makes allocation local in each mode;
2. the aggregate premium decomposes additively across modes;
3. mean-variance optimization becomes mode-wise and produces a simple first-order condition;
4. equilibrium pricing is exactly the condition under which the optimizer reproduces the market portfolio;
5. CAPM, aggregate clearing, alpha, and diversification then become consequences or interpretations of the same general framework;
6. lifecycle and bequest effects can be read as investor-specific filters on a common equilibrium benchmark rather than as rejections of equilibrium pricing.

1.5 Contributions

The paper makes five contributions.

First, it identifies a single mode-wise equilibrium equation

$$\pi_k = \lambda w_k^{\text{mkt}} v_k$$

as the strong-form structural object behind pricing, allocation, and market clearing in the model. This identifies the market portfolio as the fully harvested benchmark in the scalar-collapse regime.

Second, it shows that mean-variance allocation is the dual reading of that same equation:

$$w_k^* = \frac{\pi_k}{\lambda v_k}.$$

The usual “expected returns as optimizer inputs” view is therefore incomplete. In equilibrium, the inputs are not free primitives. They are shadow prices implied by the same structure.

Third, it embeds that scalar statement in a stronger formal hierarchy. Several economic routes now generate a route-level family $\pi_k = \lambda_k w_k^{\text{mkt}} v_k$, after which additional collinearity or homogeneity conditions collapse the family to one scalar λ .

Fourth, it connects the static benchmark to companion lifecycle results. The paper’s main economic interpretation is that lifecycle, horizon, mortality, learning, and bequest effects act as investor-specific harvesting filters on a common equilibrium benchmark rather than as evidence against equilibrium pricing itself.

Fifth, it gives the result a machine-verified formal spine. The claim is not left at the level of narrative synthesis. What is checked in Lean 4 is the algebraic chain under the paper’s stated structural assumptions, together with companion mode-fin_harvestability results elsewhere in the same research program.

1.6 Relation to the Literature

This paper does not compete with the classical literature by offering yet another factor model or another optimizer. It differs at the level of architecture.

Markowitz gives the portfolio first-order condition but takes expected return as a primitive input. CAPM gives an equilibrium pricing relation but does not collapse portfolio choice and pricing into one mode-wise identity. Mean-variance equilibrium and SDF-based formulations connect pricing to risk, but usually not in a representation where the optimizer and the equilibrium condition become

literal inverses. Factor models such as APT and Fama-French enlarge the priced object set, but they still leave open the relation between priced expected return and chosen optimal exposure.

The novelty here is to place those objects in one orthogonal mode space and show that the equations invert into each other there. The strengthened version of that claim is now more disciplined than before: the universal object is a family of route-level mode prices of risk, and the scalar theorem is the economically stronger collapse case rather than the only conceivable formulation.

That is why the theorem is stronger than a reinterpretation of CAPM and different from a repackaged Markowitz argument. The claim is not that one can *derive* familiar formulas in another basis. The claim is that the deep split between pricing and allocation disappears once the right structural object is written down.

1.7 Why This Is Not Just a Basis Change

A natural objection is that any quadratic optimization problem looks cleaner after diagonalization, so perhaps the paper is only a linear-algebra restatement.

That objection would be decisive if the paper only simplified the optimizer. But the result here is stronger: the equilibrium pricing relation and the allocation first-order condition collapse onto the same equation in the same coordinates. The gain is not merely computational. It is conceptual.

Still, the objection exposes the true frontier. The theorem is structural once a mode decomposition is fixed. The hard economic question is which decomposition is canonical, stable, and empirically meaningful. The algebra is finished earlier than the economics.

1.8 What This Paper Does Not Claim

The result is deliberately narrower than its rhetorical implications might suggest.

This paper does not prove that any orthogonal basis is economically canonical. It does not prove that Fourier, PCA, wavelets, or any other basis is *the* market basis. It does not prove that empirical expected returns can be estimated without error. And it does not yet deliver a large-sample validation paper establishing that the strongest empirical version of the theory dominates competing approaches across markets.

What it does claim is simpler and more foundational: once a meaningful orthogonal mode decomposition exists, the split between pricing and allocation disappears.

1.9 Why the Claim Matters

If correct, the paper changes the interpretation of several standard objects.

Markowitz is no longer “an optimizer that needs expected returns from somewhere else.” It is the allocation-side face of a pricing equation.

CAPM is no longer the master theory. It is the flat-spectrum special case in which the price of risk is constant across modes.

An EMH-style reading is not the statement that every local object is perfectly priced at every moment. In this paper it is the aggregate clearing condition implied by the same equation.

Active management is then a claim about mode-level alpha. Passive management is the equilibrium solution. The old separation between “valuation” and “portfolio construction” becomes conceptually artificial.

The same reframing also changes how lifecycle investing should be interpreted. If the mode-pricing law is common, then horizon-dependent portfolios do not automatically signal disagreement about what risk is priced. They may instead reflect disagreement only about what fraction of each priced mode is harvestable for a given investor. That turns an apparent conflict between equilibrium finance and lifecycle finance into a layered theory: common pricing underneath, investor-specific harvesting above it.

1.10 Paper Structure

The paper proceeds in seven steps.

1. Section 2 defines the mode-space objects and the decomposition assumptions.
2. Section 3 states the equilibrium pricing restriction and derives the allocation equation.
3. Section 4 states and interprets the pricing-allocation unity theorem.
4. Section 5 shows how CAPM, aggregate clearing, alpha, and diversification follow from it.
5. Section 6 explains the economic meaning of the theorem and its boundaries.
6. Section 7 summarizes the formal-verification scope.
7. Section 8 concludes with implications and open problems.

2. Mode Space

2.1 Model Objects and Notation

We work with a one-period mean-variance setting.

Let $r \in \mathbb{R}^n$ denote the vector of risky excess returns over the risk-free asset on a fixed probability space. Let $x \in \mathbb{R}^n$ denote a portfolio in asset space. The portfolio excess return is

$$R(x) = x^\top r.$$

Assume there is an orthogonal decomposition operator that maps asset-space returns into mode-space coordinates. We write the resulting mode random variables as

$$c_1, \dots, c_N,$$

and we assume they span the relevant priced subspace of portfolio returns. A portfolio can therefore be represented either by asset weights x or by mode exposures $w = (w_1, \dots, w_N)$, where

$$R(w) = \sum_{k=1}^N w_k c_k.$$

This paper does not derive the decomposition operator. It assumes one has been fixed and asks what follows structurally once pricing and allocation are both written in that representation.

2.2 Orthogonal Return Decomposition

Let aggregate excess return $R - r_f$ admit an orthogonal decomposition into modes:

$$R - r_f = \sum_{k=1}^N c_k.$$

We assume the modes are pairwise orthogonal in the second-moment sense, so that

$$\text{Cov}(c_j, c_k) = 0 \quad (j \neq k),$$

and therefore

$$\text{Var}(R) = \sum_{k=1}^N \text{Var}(c_k).$$

Write

$$v_k := \text{Var}(c_k) > 0$$

for the variance of mode k , and

$$\pi_k$$

for the premium carried by that mode.

The first structural fact is simply additive accounting:

$$\Pi := \mathbb{E}[R] - r_f = \sum_{k=1}^N \pi_k.$$

This is the mode decomposition of the aggregate risk premium.

2.3 Why Mode Space Is the Right Location

In raw asset space, covariance couples everything to everything else. Optimization is global and hard to interpret economically. In mode space, the covariance structure is diagonal. Each mode becomes one local risk-bearing direction with its own premium and variance.

That is not yet an economic theorem. It is a representation theorem. But it is precisely the representation in which pricing and allocation can be compared term by term.

2.4 Mode Sharpe Ratios and a Boundary

Define the mode Sharpe ratio by

$$\text{SR}_k := \frac{\pi_k}{\sqrt{v_k}}.$$

The formal kernel tracks the quantity

$$\sum_{k=1}^N \text{SR}_k^2.$$

What this draft does **not** claim is that the aggregate asset-space Sharpe ratio always equals that sum without additional assumptions. The point of the mode Sharpe vector here is diagnostic: it measures how priced return is distributed across orthogonal directions and makes diversification a statement about the number and quality of active modes, not merely the number of assets.

3. From Pricing to Allocation

3.1 The Allocation Equation

Consider a mean-variance investor who chooses mode exposures w_k to maximize

$$\sum_{k=1}^N w_k \pi_k - \frac{\lambda}{2} \sum_{k=1}^N w_k^2 v_k,$$

where $\lambda > 0$ is risk aversion.

Because mode space diagonalizes the covariance structure, the problem separates mode by mode. The first-order condition is

$$\pi_k - \lambda w_k v_k = 0.$$

Hence the optimizer chooses

$$w_k^* = \frac{\pi_k}{\lambda v_k}.$$

This is Markowitz in diagonal form. The usual portfolio problem has not disappeared; it has been rewritten in the coordinates in which its economic logic is transparent.

3.2 The Pricing Equation and Its Derived Routes

The paper's central equilibrium relation in mode space is

$$\pi_k = \lambda w_k^{\text{mkt}} v_k.$$

This equation is the mode-space analogue of the statement that priced return is proportional to marginal risk times the price of risk.

There are now three formal routes in the Lean spine that generate this restriction rather than merely postulate it.

1. A **CCAPM route**, where per-mode premiums are generated by consumption covariance and market exposure agrees with consumption loading.
2. An **SDF route**, where per-mode premiums are generated by an SDF loading and that loading agrees with market exposure.
3. An **Arrow-Debreu route**, where a complete market spans every mode and Arrow-Debreu pricing holds on each spanned mode.

What remains unproved is not whether these routes exist, but whether they exhaust the economically relevant cases and whether one canonical decomposition is selected by market data. So the current contribution remains structural, but it is no longer fair to describe the pricing restriction as merely an unsupported assumption.

3.3 The Inversion

Substituting the pricing equation into the allocation equation gives

$$w_k^* = \frac{\lambda w_k^{\text{mkt}} v_k}{\lambda v_k} = w_k^{\text{mkt}}.$$

So in equilibrium, the optimizer reproduces the market allocation exactly, mode by mode.

Conversely, if the optimizer reproduces market allocation, then rearranging the first-order condition yields the pricing equation.

This is the central duality:

$$\pi_k = \lambda w_k^{\text{mkt}} v_k \iff w_k^* = w_k^{\text{mkt}}.$$

That equivalence is the heart of the paper.

4. The Pricing-Allocation Unity Theorem

4.1 Formal Statement and Assumptions

The core theorem is a one-period mean-variance statement. The assumptions are:

1. There is a fixed orthogonal mode decomposition of the priced return space.
2. A portfolio can be represented by mode exposures w_k , with per-mode premium π_k and variance $v_k > 0$.
3. The investor solves the diagonal mean-variance problem

$$\max_w \sum_{k=1}^N w_k \pi_k - \frac{\lambda}{2} \sum_{k=1}^N w_k^2 v_k$$

with scalar risk-aversion parameter $\lambda > 0$.

4. The market portfolio has mode exposures w_k^{mkt} .
5. The mode-space pricing relation takes the scalar form

$$\pi_k = \lambda w_k^{\text{mkt}} v_k.$$

Assumption 5 can be imposed directly as a structural equilibrium restriction. In the current Lean spine it can also be obtained along three formal routes with additional hypotheses: a CCAPM route, an SDF-loading route, and a complete-market Arrow-Debreu route. What is not yet proved is that every economically relevant model must collapse to this exact scalar- λ form.

4.2 Main Theorem

Under Assumptions 1-5, the following are equivalent within the model:

1. **Pricing:** $\pi_k = \lambda w_k^{\text{mkt}} v_k$ for all k .
2. **Allocation:** $w_k^* = w_k^{\text{mkt}}$ for all k .
3. **Mean-variance market clearing:** the market portfolio is the mean-variance optimum in mode space.

In the normalized special case $w_k^{\text{mkt}} = 1$, this reduces to

$$\pi_k = \lambda v_k \iff w_k^* = 1,$$

which is the form that first appears in the formal kernel.

4.3 Proof Sketch

The proof has one algebraic core and one economic wrapper.

1. The investor first-order condition gives

$$w_k^* = \frac{\pi_k}{\lambda v_k}.$$

2. The pricing relation gives

$$\pi_k = \lambda w_k^{\text{mkt}} v_k.$$

3. Substituting the pricing equation into the optimizer gives

$$w_k^* = \frac{\lambda w_k^{\text{mkt}} v_k}{\lambda v_k} = w_k^{\text{mkt}}.$$

4. Conversely, if the optimizer reproduces market allocation, rearranging the same first-order condition yields the pricing equation.

That is the complete algebraic argument. The economic weight of the theorem lies in how Assumption 5 is justified. In the current formal kernel, that justification is supplied conditionally by three derived routes:

1. **CCAPM**: if mode premiums are generated by consumption covariance and market exposure agrees with consumption loading, then the general pricing restriction follows.
2. **SDF**: if mode premiums are generated by an SDF loading and the loading agrees with market exposure, then the same restriction follows.
3. **Arrow-Debreu**: if every mode is spanned in a complete market and Arrow-Debreu pricing holds on each spanned mode, then the same restriction follows globally.

So the theorem is algebraically simple but no longer hangs on a single unsupported premise. What remains open is the broader economic question of when these routes are exhaustive and when the scalar- λ specialization is the right one.

4.4 Interpretation

The theorem says that pricing and allocation are primal-dual views of one equilibrium object.

Pricing tells us what premium each mode must earn if the market holds it.

Allocation tells us what exposure an optimizer would choose if that premium is offered.

Efficiency is the consistency condition under which these two views coincide.

The theorem is therefore not “pricing affects allocation” in the loose causal sense. It is stronger: the equations are algebraic inverses.

4.5 The Seven-Step Argument

The argument can be read as a seven-step chain.

1. Orthogonal decomposition diagonalizes risk and turns covariance into mode-wise variance accounting.
2. Aggregate excess return decomposes additively into mode premiums.

3. Mean-variance optimization becomes local in each mode and yields $w_k^* = \pi_k / (\lambda v_k)$.
4. Market clearing prices each mode by $\pi_k = \lambda w_k^{\text{mkt}} v_k$.
5. Substitution yields $w_k^* = w_k^{\text{mkt}}$.
6. Therefore CAPM, aggregate clearing, and Markowitz are not competing theories but different projections of the same equation.
7. Deviations from equality are mode-level alpha, not a collapse of the whole framework.

This is the paper’s full argument in its most compressed form.

4.6 Why This Is More Than a Coordinate Trick

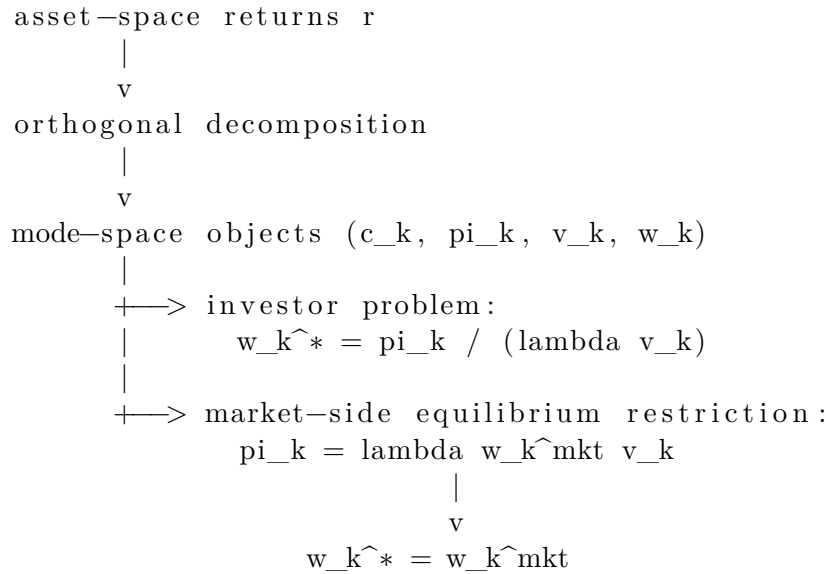
A legitimate objection is that every quadratic problem looks simpler after diagonalization. Why is this not just linear algebra dressed up as economics?

The answer is that the theorem does not merely say “the optimizer is easier to compute after diagonalization.” It says that the *pricing* relation and the *allocation* relation collapse onto the same equation in the same coordinates. The unification is not computational only. It is conceptual.

Still, the objection points to the next frontier: the economics lives in the choice of decomposition. The algebraic theorem is complete once the decomposition is fixed. The empirical and structural question is which decomposition the market itself actually uses or approximates.

4.7 Figure 1: Conceptual Map

The logic of the paper can be summarized schematically:



The point of the figure is not computational. It is architectural. Pricing and allocation are written as separate routes through the same mode-space object system, and the unity theorem says those two routes coincide when the market satisfies the equilibrium restriction.

4.8 A Two-Mode Toy Example

The algebra is easiest to see in a minimal non-normalized example.

Take two orthogonal modes with

$$v_1 = 1, \quad v_2 = 4,$$

market exposures

$$w_1^{\text{mkt}} = 0.4, \quad w_2^{\text{mkt}} = 0.6,$$

and risk-aversion parameter

$$\lambda = 2.$$

Under the paper's equilibrium restriction, the implied mode premiums are

$$\pi_1 = \lambda w_1^{\text{mkt}} v_1 = 2 \cdot 0.4 \cdot 1 = 0.8,$$

$$\pi_2 = \lambda w_2^{\text{mkt}} v_2 = 2 \cdot 0.6 \cdot 4 = 4.8.$$

The allocation equation then gives

$$w_1^* = \frac{\pi_1}{\lambda v_1} = \frac{0.8}{2 \cdot 1} = 0.4, \quad w_2^* = \frac{\pi_2}{\lambda v_2} = \frac{4.8}{2 \cdot 4} = 0.6.$$

So the optimizer reproduces the market exactly without any normalization such as $w_k^{\text{mkt}} = 1$. The point is not that the market weights are trivial. The point is that once premiums satisfy the pricing restriction, the optimizer and the market speak the same coordinates.

The same example also clarifies alpha. If an investor believes mode 2 should earn $\tilde{\pi}_2 = 5.2$ instead of 4.8, then the perceived mode alpha is

$$\alpha_2 = \tilde{\pi}_2 - \lambda w_2^{\text{mkt}} v_2 = 5.2 - 4.8 = 0.4,$$

and the investor's optimal exposure becomes

$$\tilde{w}_2^* = \frac{5.2}{2 \cdot 4} = 0.65,$$

which is an active tilt away from the market allocation. In this sense, active management is a claim that one or more mode premiums differ from the equilibrium values implied by market holdings.

5. Classical Finance as Interpretations and Corollaries

5.1 Markowitz

Markowitz is the allocation-side reading of the unity theorem. In mode space, the covariance matrix is diagonal and the first-order condition is exact:

$$w_k^* = \frac{\pi_k}{\lambda v_k}.$$

So the classical mean-variance problem survives unchanged, but the role of expected return changes. It is no longer an arbitrary exogenous input. In equilibrium it must be the premium implied by the same priced-risk equation.

5.2 CAPM

CAPM appears when the relevant price-of-risk object is constant across the priced directions, yielding a flat equilibrium condition in the appropriate spectral sense. In that regime, beta pricing is the asset-space shadow of the mode-space equation.

CAPM is therefore a special case, not the root theorem. It is what the general unity reduces to when the market's risk-pricing structure has the flat-spectrum form required by the model.

5.3 Aggregate Clearing and the EMH Interpretation

The aggregate equilibrium statement is

$$\Pi = \lambda \sum_{k=1}^N v_k$$

in the normalized case, or its market-weighted generalization in the full formulation.

This can be read as an EMH-style interpretation only with care. What is proved in the present draft is an aggregate clearing identity inside the mode-space model. A stronger informational-efficiency interpretation would require extra economic assumptions that are not derived here. Local deviations can exist mode by mode without invalidating the aggregate equilibrium architecture.

5.4 Alpha and Anomalies

Define mode alpha by

$$\alpha_k := \pi_k - \lambda w_k^{\text{mkt}} v_k.$$

Then $\alpha_k = 0$ is simply equilibrium at mode k . Positive alpha means the mode is underpriced relative to the general equation; negative alpha means it is overpriced.

This yields a cleaner interpretation of active management:

- passive management accepts the equilibrium allocation;
- active management is a claim that some $\alpha_k \neq 0$;
- successful active management is a claim that those deviations will eventually decay.

5.5 Diversification

Because the optimizer becomes mode-wise in orthogonal coordinates, diversification is fundamentally about mode count and mode quality. Adding assets matters only insofar as it adds independent or weakly redundant priced directions.

This is a sharper statement than the textbook diversification slogan. The true scarce object is not “many assets.” It is many economically meaningful modes.

6. Economic Meaning and Boundaries

6.1 The Real Contribution

The most important economic consequence of the theorem is not a new optimizer. It is a reinterpretation of expected return.

The standard workflow in quantitative portfolio construction is:

1. estimate expected returns,
2. estimate covariance,
3. optimize portfolio weights.

The theorem here says that step 1 and step 3 cannot be conceptually independent. Once the market and the investor share the same mode space, expected return is the shadow price consistent with the same equilibrium structure that determines optimal holdings.

In that sense, expected return is not an input. It is an equilibrium object.

6.2 Pricing Is Common, Harvesting Is Investor-Specific

The most promising application layer is intertemporal rather than static.

In the scalar-collapse regime emphasized in this paper, the long-run benchmark allocation is

$$w_k^\infty = \frac{\pi_k}{\lambda v_k} = w_k^{\text{mkt}}.$$

So the equilibrium market allocation is not just an observed portfolio. It is the fully harvested benchmark exposure implied by the common mode-pricing law.

Now place that benchmark next to the lifecycle result developed elsewhere in this research program:

$$w_k^{\text{life}} = w_k^\infty \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k,$$

where $\mathbb{E}[h_k]$ is expected `fin_harvestability` at the investor's effective horizon, $b_k \in [0, 1]$ is a Bayesian caution term, and $s_k \in [0, 1]$ is a safety or floor multiplier.

Substituting the pricing-allocation identity gives

$$w_k^{\text{life}} = w_k^{\text{mkt}} \cdot \mathbb{E}[h_k] \cdot b_k \cdot s_k.$$

This is economically stronger than the static unity theorem by itself. It says investors need not disagree about equilibrium pricing in order to disagree about optimal holdings. They can share the same priced mode structure and still choose different allocations because they differ in horizon, mortality, floor constraints, and bequest motive.

That observation resolves a large practical confusion. Lifecycle investing and equilibrium pricing are often presented as competing views: one says the young should hold more risky assets because horizon matters, while the other says the market portfolio is the efficient benchmark. In the present framing they are not competitors. The market portfolio is the common fully harvested benchmark, and lifecycle tilts describe how much of each equilibrium-priced mode a particular investor can actually harvest.

This immediately yields a nontrivial bequest implication. A selfish finite-horizon investor will generally scale down slow modes because their `fin_harvestability` is low at the investor's effective horizon. By contrast, a dynastic or strong-bequest investor increases effective horizon, pushing `fin_harvestability` toward one and recovering the market or Merton benchmark more closely. Samuelson-type horizon independence is then not the universal truth. It is the limiting case of investors whose effective horizon is long enough that most relevant modes are almost fully harvestable.

In that sense, the right split is no longer pricing versus allocation. The deeper split is common pricing versus investor-specific harvesting.

6.3 What Remains Economically Hard

Three hard problems remain.

First, what is the economically canonical mode decomposition? A theorem that works for any orthogonal basis is mathematically strong, but economics needs more than basis invariance.

Second, how stable is the decomposition through time? If the relevant market modes drift, then the equation remains true conditionally but the state variable itself becomes dynamic.

Third, how should mode-level alpha be estimated under noise and belief disagreement? In practice, π_k is never observed directly. The theorem reorganizes the conceptual problem, but it does not eliminate statistical difficulty.

6.4 Non-Claims

This paper does not claim:

1. that every anomaly is a pure mode-level mispricing signal rather than an estimation artifact;
2. that the theorem alone identifies the best tradable basis;
3. that the current draft already resolves the empirical asset-pricing debate;
4. that the equilibrium relation is always static rather than intertemporal or state-dependent.

Those are important next questions. They are not needed for the core theorem.

6.5 The Natural Next Layer

The next theoretical layer is to connect the mode basis to deeper primitives:

1. the stochastic discount factor,
2. belief heterogeneity and disagreement,
3. dynamic programming and value-function contraction,
4. or the spectral structure of a market generator.

If that layer is completed, then the theorem would move from a powerful structural unification to a deeper theory of why the market basis has the form it does.

6.6 Roadmap Beyond the First Derived Routes

The most natural upgrade path for this paper is no longer to obtain the very first derivation, but to unify and strengthen the derivations that are now already present in the formal spine.

One route is stochastic-discount-factor based. Start with an SDF m and impose

$$\mathbb{E}[m(r_i - r_f)] = 0$$

for every asset i . The current Lean development now contains both a reduced-form mode-loading version of this route and a formal bridge from route-specific mode prices of risk to the scalar-collapse case. The next step is to bridge that reduced-form loading object to a more explicit

asset-level pricing-kernel derivation. In the richer version, the mode premiums should inherit a relation of the form

$$\pi_k = \lambda_k w_k^{\text{mkt}} v_k$$

or, in the simplest case, the scalar- λ specialization used here.

A second route is complete-market Arrow-Debreu pricing. The current Lean spine already shows that if every mode is spanned and Arrow-Debreu pricing holds on each spanned mode, then both a universal mode-pricing reading and the scalar general restriction follow through the complete-market bridge. The next step is to connect that complete-market theorem more tightly to observable asset structures and spanning results.

A third route is representative-agent or mean-variance equilibrium. There one derives market holdings from the aggregation of investor first-order conditions and then shows that the market portfolio must satisfy the same pricing-allocation identity mode by mode.

A fourth route is dynamic. The Bellman bridge developed elsewhere in this research program suggests that the static pricing-allocation identity may be the one-period shadow of a deeper contraction or value-function recursion.

A fifth route is intertemporal and lifecycle-based. The static theorem identifies the fully harvested equilibrium benchmark. The natural next step is then to show, in one integrated economic theorem, how finite horizon, stochastic mortality, Bayesian uncertainty, and bequest motives transform that benchmark into investor-specific holdings through multiplicative `fin_harvestability` filters rather than through disagreement about the underlying priced modes.

This paper therefore no longer sits at the stage of a single naked assumption. It now has several formal derived routes, a route-level λ_k formulation, and explicit collapse bridges back to the scalar theorem, but it still does not unify all those routes into one economically dominant theorem.

7. Formal Verification Scope

The formal contribution of this paper is not a vague claim of consistency. The core algebraic spine is machine-verified in Lean 4 under the paper’s stated structural assumptions.

7.1 Verification Status

Paper Claim	Status	Lean Support	Interpretation
Variance and premium decompose across orthogonal modes	Lean-proved	<code>parseval_variance</code> , <code>total_premium_decomposition</code>	Exact algebraic identities inside the mode decomposition model
Mean-variance optimization is diagonal in mode space	Lean-proved	<code>diagonal_markowitz_foc</code>	Exact first-order condition under the one-period mean-variance setup

Paper Claim	Status	Lean Support	Interpretation
Pricing restriction implies market allocation, and conversely	Lean-proved	<code>general_pricing_implies_allocation</code> <code>general_allocation_implies_pricing</code>	Theorem , core structural equivalence of the paper
CCAPM, SDF, and Arrow-Debreu routes generate the pricing restriction	Lean-proved under route-specific assumptions	<code>ccapm_implies_general_equilibrium</code> <code>sdf_implies_general_equilibrium</code> <code>complete_market_implies_general_equilibrium</code>	Theorem , conditional derivations, not yet generalized equilibrium theorem
Derived pricing routes imply market allocation	Lean-proved under route-specific assumptions	<code>ccapm_implies_market_allocation</code> <code>sdf_implies_market_allocation</code> <code>arrow_debreu_complete_allocation</code>	Theorem , bridge the economic routes back into the allocation theorem
CAPM, aggregate clearing, anomaly cancellation, and diversification are corollaries	Lean-proved	<code>capm_from_equilibrium</code> <code>spectral_emh_total</code> <code>anomaly_cancellation</code> <code>more_modes_strictly_better</code>	These depend on the model assumptions already stated
EMH as a full informational-efficiency doctrine	Interpretation only	None	Lean proves an aggregate clearing identity, not the whole empirical EMH program
Canonical basis selection and empirical validation	Future work	None	Not mechanized and not claimed as solved here

7.2 Minimal Reproducibility Path

The present paper has a theorem artifact, not an empirical replication package.

1. The core theoretical artifact is the Lean code under `LeanProofs/PricingAllocation/`.
2. The central files for the current paper are `PricingIsAllocation.lean`, `ConsumptionCAPM.lean`, `StochasticDiscountFactor.lean`, `ArrowDebreu.lean`, and `MainTheorem.lean`.
3. The current draft does **not** include a separate data or simulation artifact; accordingly, the paper does **not** claim empirical validation.

7.3 Scope Boundary

What the formal verification does **not** cover is equally important:

1. the empirical identification of the correct market basis;
2. large-scale market-data validation;
3. the claim that any one economic interpretation of the modes is already final.

The machine-checked result is therefore a structural theorem about the internal consistency of the mode-space formulation, not an empirical certification of the full finance program. It now includes several derived equilibrium routes, but it is not yet a universal derivation from all deeper primitives or a proof that one canonical market basis is singled out empirically.

8. Conclusion

This paper advances a simple but still deliberately bounded claim: within a stated mode-space equilibrium model, pricing and allocation are the same equation viewed from different sides.

The traditional split between asset pricing and portfolio construction comes from working in the wrong coordinates. In mode space, the split disappears. Premium, risk contribution, and optimal weight are tied together by one identity. CAPM, Markowitz, and aggregate-clearing interpretations are then not rival kingdoms of finance, but corollaries or readings of a common equilibrium structure.

The contribution is conceptual before it is computational. It tells us what a coherent joint mode-space theory must look like. It says that any architecture treating expected return as a free-standing upstream input to allocation is, at best, incomplete.

It also suggests a sharper application than the draft had before. Investors do not need different pricing theories in order to hold different portfolios. They may share one equilibrium mode-pricing law and differ only in what fraction of each priced mode they can harvest at their effective horizon. That gives the theorem a direct bridge to lifecycle investing, bequest motives, and horizon-dependent portfolio choice.

The main open question is not the internal algebraic consistency of the model. The open question is economic: what is the right market decomposition, how does it evolve, which routes generate only a mode-specific λ_k family and which genuinely collapse to a scalar λ , and how much of observed alpha is genuine mode-level disequilibrium rather than estimation noise or belief disagreement?

That is where the next papers belong. The present result isolates the theorem they will have to respect.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, 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

- Arrow, K. J., and Debreu, G (1954). Existence of an equilibrium for a competitive economy. *Econometrica*. DOI: 10.1017/ccol052123736x.005
- Black, F., and Litterman, R (1992). Global portfolio optimization. *Financial Analysts Journal*. DOI: 10.2469/faj.v48.n5.28
- Breeden, D. T (1979). An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*. DOI: 10.1142/9789812701022_0003
- Cochrane, J. H (2001). Asset Pricing. *Asset Pricing*, 0-691.

- Fama, E. F (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*. DOI: 10.2307/2325486
- Hansen, L. P., and Jagannathan, R (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy*. DOI: 10.3386/t0089
- Markowitz, H (1952). Portfolio selection. *Journal of Finance*, 7(1), 77-91.
- Merton, R. C (1973). An intertemporal capital asset pricing model. *Econometrica*. DOI: 10.2307/1913811
- Ross, S. A (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*. DOI: 10.1016/0022-0531(76)90046-6
- Sharpe, W. F (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*. DOI: 10.2307/2977928