Harvestability

A Proof-First Foundation Note on Horizon-Dependent Correction

Samuelson's Horizon Independence as a Derived Knife-Edge Consequence

Dr. Tamás Nagy

`tnagyphd@gmail.com`

Working Paper • March 2026

Executive Summary (Non-Technical)

Finance has long known that investment horizon can matter in practice, especially under predictability and mean reversion, but those effects often appear in scattered model-specific form rather than as one reusable object for expressing **how much of a risky premium is actually available at a given horizon**. This paper makes that object explicit within one narrow formal setting. It introduces **fin_harvestability** as a reusable horizon quantity for this paper family.

The paper is **proof-first by design**. Its main job is to define the fin_harvestability object, derive it from the Ornstein-Uhlenbeck and HJB-Riccati spine, and make the proof boundary explicit for later finance papers in the family. The goal is not primarily calibration, productization, or benchmark-bridge positioning. Read in its narrowest external form, this manuscript is a foundation note on horizon-dependent correction, not a general re-founding of long-horizon allocation.

Its main allocation reading is benchmark-centered. The myopic Merton allocation remains the baseline, and fin_harvestability governs the horizon-dependent correction around that baseline. A simpler scaled-weight surface appears in some core propositions, but it should be read here as an auxiliary stylized layer rather than as the paper's definitive optimal-control formula.

Its scope is intentionally hierarchical. First comes the main object. Second comes the derivation that makes the object non-ad hoc. Third comes the lifecycle extension, where horizon, bequest, mortality, Bayesian caution, and safety constraints enter multiplicatively. That extension layer should be read as a disciplined structured extension of the core object, not as a claim to solve the full lifecycle problem in one unified economic model. Only after those layers are in place does the paper turn to the **Samuelson implication**, which is treated here as a consequence of the fin_harvestability theory rather than as the paper's primary identity.

This role matters for the surrounding paper family. The benchmark-side paper fin_pricing_is_allocation imports the investor-specific harvesting layer developed here when it explains how common pricing can coexist with heterogeneous holdings. The companion fin_harvestability_calibration paper imports the object and proof boundary established here, then specializes to calibration, reporting discipline, and narrower reviewer-facing use.

Abstract

This paper studies **fin_harvestability** as a horizon object for horizon-dependent allocation within a CRRA investor facing Ornstein-Uhlenbeck eigenmodes. The main object is the fin_harvestability function $h(T, \tau) = 1 - e^{-T/\tau}$, which measures how much of a risky mode's premium is reliably capturable at horizon T when the mode has characteristic time scale τ. The paper's primary role is proof-first: it isolates this object, derives it from an OU/HJB/Riccati spine, and makes the proof boundary explicit. Its central allocation claim is benchmark-centered: fin_harvestability organizes the horizon-dependent correction around the myopic Merton benchmark rather than replacing that benchmark with a pure multiplicative rescaling. The theory then records a downstream lifecycle-filter form, in which horizon, bequest, mortality, Bayesian caution, and safety constraints enter multiplicatively; this should be read as a structured extension layer rather than as an exhaustive lifecycle model. Within this hierarchy, the **Samuelson Error** $\varepsilon(T, \tau) = e^{-T/\tau}$ and the failure of horizon independence are treated as derived consequences rather than as the paper's foundational identity. The paper does not claim to settle the broader long-horizon allocation literature or to solve the full lifecycle problem in one calibrated economic model. In the surrounding paper family, the manuscript plays the role of a proof-oriented foundation note that downstream benchmark and calibration notes can import from without re-establishing the derivation. Its Lean contribution should be read as checking the structural backbone of the argument; some surrounding analytic estimates remain classical.

Paper Role in the Research Program

Within this repo's finance-paper family, this paper serves as the **proof-first foundation note** for the fin_harvestability line.

Its scope hierarchy is:

1. the main object h(T, \tau),
2. the OU/HJB/Riccati derivation spine,
3. the lifecycle-filter extension,
4. the Samuelson implication as a derived consequence.

Within the finance-paper family, fin_pricing_is_allocation imports this paper's investor-specific harvesting layer when it explains how common benchmark pricing can coexist with heterogeneous holdings. The companion fin_harvestability_calibration paper imports the object and proof boundary established here, then specializes to calibration, documented defaults, and narrower reviewer-facing use. This is a paper-family role, not a claim that the present manuscript exhausts the literature on long-horizon allocation. The provocative Samuelson framing should therefore be read as a consequence of this paper, not as a separate foundational dependency.

---

1. Introduction

1.1 Why a Reusable Horizon Object Is Useful

Financial economics has long contained two neighboring ideas without a single stable object linking them. On one side, the Samuelson-Merton benchmark under independent returns implies horizon independence. On the other side, the literature on predictability, mean reversion, and lifecycle investing implies that horizon can matter materially for optimal holdings. What is often missing is a reusable quantity within one disciplined model that says how much of a risky premium is actually available to an investor at horizon T.

This paper introduces that quantity. It treats **fin_harvestability** as a reusable horizon object within the present OU/CRRA theory family.

1.2 The Main Object and Main Claim

The main object of the paper is the fin_harvestability function

$$h(T, \tau) = 1 - e^{-T/\tau},$$

where T is the investor's horizon and τ is the mode's characteristic time scale.

The main claim is narrow and explicit: within the OU/CRRA setting, fin_harvestability is a useful horizon object that can be defined precisely, derived from a first-principles OU/HJB/Riccati spine, and used to organize the horizon-dependent correction around the myopic Merton benchmark. The broader lifecycle layer is secondary and downstream to that claim. The paper does not claim to be the first demonstration that horizon matters, nor does it claim to solve the full lifecycle allocation problem in one calibrated end-to-end framework.

Canonical theorem sentence for the public-facing manuscript: within the OU/CRRA eigenmode setting, $h(T, \tau)$ is the closed-form horizon-dependent correction term around the myopic Merton benchmark.

Relative to the nearest prior-art set, the paper's claim is also deliberately limited. Samuelson (1969) and Merton (1969, 1971) remain the benchmark references for IID and continuous-time lifetime portfolio choice, while Campbell and Viceira (1999, 2002), Barberis (2000), Wachter (2002), and Brennan and Xia (2010) remain the closest long-horizon predictable-return comparators. The present manuscript does not claim a new economic domain beyond that literature, nor does it claim that the exponential transient itself was previously unknown in mean-reverting control problems. Its strongest distinctiveness claim is narrower: within one OU/CRRA setting, it names a reusable horizon object, packages the derivation around the benchmark-centered correction that object governs, and makes the proof boundary explicit enough for downstream papers to reuse.

The comparison should be stated plainly. Relative to Campbell and Viceira, this paper is not trying to supersede the broader long-lived consumption and intertemporal-hedging program; it works in a narrower CRRA/OU eigenmode frame and emphasizes the reusable correction object rather than the full strategic-allocation architecture. Relative to Wachter, it is not claiming the strongest exact-solution consumption result in complete markets; it is claiming that, inside a simpler local setting, the horizon-dependent transient can be isolated, named, and carried with an unusually explicit verification boundary. Read this way, the manuscript's value is more in canonicalization, proof-first packaging, and downstream reuse than in asserting a new primitive finance mechanism.

1.3 Scope Hierarchy

The paper's scope is intentionally hierarchical.

1. **Main object**: define the fin_harvestability function and its basic properties.
2. **Derivation**: derive the object from the CRRA control problem with Ornstein-Uhlenbeck modes.
3. **Lifecycle extension**: extend the object into the multiplicative lifecycle filter with horizon, bequest, mortality, Bayesian caution, and safety constraints.
4. **Derived Samuelson implication**: show that horizon independence is a knife-edge special case and quantify its failure through the Samuelson Error.

This order matters. The paper is not primarily a provocation, a calibration note, or a benchmark-bridge manuscript. It is the proof-oriented foundation note for the fin_harvestability line.

1.4 Samuelson as a Derived Consequence

Paul Samuelson's 1969 result remains a central benchmark because it isolates the IID case cleanly. Under independent returns, the myopic allocation is horizon-independent. But once returns exhibit economically relevant time-scale structure, that conclusion stops being generic.

This paper therefore treats the Samuelson result as an important **derived boundary case**, not as the paper's primary identity. The knife-edge nature of horizon independence and the associated Samuelson Error emerge from the fin_harvestability theory developed here; they do not define the paper's scientific role.

1.5 Relation to Literature and Downstream Papers

The literature on horizon-dependent allocation already contains important ingredients: Campbell and Viceira (1999, 2002) analyze long-horizon allocation under return predictability, Barberis (2000) incorporates Bayesian learning, Wachter (2002) derives exact solutions under mean-reverting expected returns, and Brennan and Xia (2010) extend the setting further. The closest economics neighborhood for this paper is therefore exact predictable-return and intertemporal-hedging theory, not a new finance domain. In particular, the manuscript should be read as adjacent to Wachter- and Campbell-Viceira-style exact-solution territory rather than as a claim that the underlying economic mechanism was absent from prior finance theory. The present paper differs less by claiming that horizon matters for the first time and more by isolating a single proof-organizing object, deriving it in a proof-first way within one OU/CRRA model, and making the proof boundary explicit through a Lean-checked structural backbone. Some surrounding analytic estimates and economic interpretations remain classical in the manuscript.

This also clarifies what the manuscript is *not* trying to win. It is not the paper one would cite for the broadest consumption-based exact solution, nor for the richest recursive-utility strategic-allocation environment. It is the paper one would cite if one wanted the narrow correction object itself, the local derivation route that exposes it, and an explicit statement of which parts of that route are mechanically checked versus merely discussed in prose.

This foundation-note role also determines how downstream papers use it. The benchmark-side fin_pricing_is_allocation line imports the investor-specific harvesting layer from this paper when explaining how common pricing can coexist with heterogeneous holdings. The fin_harvestability_calibration line imports the object defined here and the proof boundary established here, then specializes to calibration, documented defaults, and narrower reviewer-facing application.

1.6 Paper Organization

Section 2 introduces the spectral framework for multi-asset returns. Section 3 defines the fin_harvestability object, the derived Samuelson Error, and the main theorem. Section 4 gives the first-principles OU/HJB/Riccati derivation. Section 5 records downstream lifecycle extensions with explicit caveats. Section 6 records selected illustrative readings rather than main claims. Section 7 documents the Lean verification architecture, and Section 8 concludes.

---

2. Spectral Decomposition of Multi-Asset Returns

2.1 Eigenvalue Decomposition

Consider n risky assets with covariance matrix Σ and expected excess returns μ. The eigenvalue decomposition $\Sigma = V\Lambda V^\top$ defines n orthogonal "modes," each an independent source of risk and return. In mode space:

- **Mode premium**: $\rho_k = v_k^\top \mu$ (the premium associated with mode k)
- **Mode variance**: $\sigma_k^2 = \lambda_k$ (the variance of mode k, equal to its eigenvalue)
- **Mode Sharpe ratio**: $\mathrm{SR}_k = \rho_k / \sigma_k$

The Pythagorean decomposition holds:

$$\mathrm{Var}[\text{portfolio}] = \sum_{k=1}^{n} w_k^2 \sigma_k^2$$

This orthogonality — verified in ModePremiumVariance.lean as variance_pythagorean — is the foundation of the entire framework. It implies that the multi-asset allocation problem decomposes into n independent one-dimensional problems, one per mode.

2.2 Ornstein–Uhlenbeck Dynamics per Mode

Each mode follows an Ornstein–Uhlenbeck (OU) process:

$$dX_k = -\frac{X_k}{\tau_k} dt + \sigma_k \, dW_k$$

where $\tau_k > 0$ is the characteristic time scale (half-life of mean reversion) and $\sigma_k > 0$ is the instantaneous volatility. The structure OUMode with fields tau, sigma, and positivity hypotheses is defined in CRRAOUInvestor.lean:

```
structure OUMode where
  tau :
  sigma :
  h_tau : 0 < tau
```

```
  h_sigma : 0 < sigma
```

The empirical eigenvalue spectrum of asset returns exhibits a characteristic pattern: a few dominant modes with large τ_k (slow, macro-driven) and many modes with small τ_k (fast, mean-reverting). This power-law structure is the empirical counterpart of our spectral framework.

2.3 The CRRA-OU Investor

We consider a CRRA investor with: - Risk aversion parameter $\gamma > 0$ - Investment horizon $T > 0$ - Access to N independent OU modes

The investor structure OUModeInvestor (extending CRRAInvestor) is defined in CRRAOUInvestor.lean, bundling the investor's risk aversion, horizon, mode dynamics, and risk premia into a single well-defined object. The theorem crra_ou_well_defined verifies that all parameters are positive:

```
theorem crra_ou_well_defined {N : } (inv : OUModeInvestor N) :
    0 < inv.gamma   0 < inv.horizon
    ( k, 0 < (inv.modes k).tau)   ( k, 0 < (inv.modes k).sigma)
```

---

3. The Harvestability Object and the Derived Samuelson Error

3.1 The Harvestability Function

Definition 1 (Harvestability). The **fin_harvestability function** is

$$h(T, \tau) = 1 - e^{-T/\tau}$$

where $T \geq 0$ is the investment horizon and $\tau > 0$ is the mode's characteristic time scale.

This is defined in Harvestability/HarvestabilityFunction.lean:

```
noncomputable def fin_harvestability (T tau :  ) :   :=
  1 - Real.exp (-T / tau)
```

Proposition 1 (Basic Properties; Lean-verified). For $T > 0$ and $\tau > 0$:

(a) $0 < h(T, \tau) < 1$ (harvestability_bounded)

(b) $h(0, \tau) = 0$ (harvestability_at_zero)

(c) $h(T, \tau) \to 1$ as $T \to \infty$ (merton_recovery)

(d) h is strictly increasing in T (harvestability_increasing_in_T)

(e) h is strictly decreasing in τ (harvestability_decreasing_in_tau)

(f) $h(T, \tau) \leq T/\tau$ (harvestability_le_linear)

Interpretation. $h(T, \tau)$ measures the fraction of a mode's stationary variance that an investor with horizon T can effectively capture. At $T = 0$, nothing is captured. As $T \to \infty$, everything is captured — recovering the Merton limit. The monotonicity in τ establishes a "pecking order": fast modes (small τ) are easier to harvest than slow modes (large τ).

The bound $h \leq T/\tau$ (Proposition 1f, proved in ConcavityBound.lean as harvestability_le_linear) has practical significance: linear rules like "100 minus your age" overestimate the true fin_harvestability, providing a conservative upper bound. The concavity of h in T means that the simple linear rule is always on the safe side.

3.2 The Samuelson Error

Definition 2 (Samuelson Error). The **Samuelson Error** for mode k at horizon T is

$$\varepsilon_k(T) = e^{-T/\tau_k}$$

This is defined in HarvestabilityDerivation/SamuelsonError.lean:

```
noncomputable def samuelsonError (T tau :   ) :    :=
  Real.exp (−T / tau)
```

Proposition 2 (Samuelson Error Properties; Lean-verified). For $T > 0$ and $\tau > 0$:

(a) $\varepsilon + h = 1$ (harvestability_plus_error)

(b) $0 < \varepsilon(T, \tau) < 1$ (samuelsonError_bounded)

(c) $\varepsilon(0, \tau) = 1$ (samuelsonError_at_zero)

(d) ε is strictly decreasing in T (samuelsonError_decreasing)

(e) Slow modes have larger error: $\tau_1 < \tau_2 \implies \varepsilon(T, \tau_1) < \varepsilon(T, \tau_2)$ (slow_modes_larger_error)

(f) Doubling: $\varepsilon(2T, \tau) = \varepsilon(T, \tau)^2$ (samuelsonError_doubling)

Interpretation. The Samuelson Error is the fraction of mode variance that a finite-horizon investor fails to capture. It equals 1 at $T = 0$ (no harvesting) and decays exponentially toward 0 as $T \to \infty$. The doubling property (f) is striking: doubling the horizon squares the error, making it exponentially effective to extend the investment period.

The key insight: Samuelson's claim that horizon doesn't matter is equivalent to claiming $\varepsilon = 0$ for all modes — which holds only when $\tau_k = 0$ (no mean reversion, i.e., IID returns). Under any positive mean reversion ($\tau_k > 0$), the Samuelson Error is strictly positive for finite horizons, and horizon-dependent allocation is optimal.

3.3 The Mode Ordering (Pecking Order of Risk Harvesting)

Proposition 3 (Mode Ordering; Lean-verified). For $T > 0$ and $0 < \tau_1 < \tau_2$:

$$h(T, \tau_1) > h(T, \tau_2)$$

Faster modes are more harvestable. The ordering is transitive (if $\tau_1 < \tau_2 < \tau_3$, then $h(T, \tau_1) > h(T, \tau_3)$), proved as mode_ordering_transitive in ModeOrdering.lean.

This establishes a **pecking order** of risk harvesting. Consider an investor building a portfolio from eigenmode exposures:

1. **First harvest**: fast mean-reverting modes (small τ) — "easy alpha"

2. **Then harvest**: medium modes — “carry trades”
3. **Last harvest**: slow structural modes (large τ) — accessible only to long-horizon investors

This ordering is not a heuristic; it is a mathematical theorem, formally verified.

3.4 Harvestability as a Benchmark-Centered Correction

The derivation-backed allocation identity preview is benchmark-centered:

$$w_k^*(t) = w_k^{\text{Merton}} + \eta(\gamma)\,\frac{\rho_k}{\sigma_k^2}\,h(T-t,\tau_k), \qquad \eta(\gamma) = \frac{\gamma-1}{\gamma^2}.$$

In this reading, fin_harvestability does not replace the myopic benchmark; it parameterizes the time-dependent correction around that benchmark. This full decomposition is formalized in FullAllocation.lean and derived in Section 4.6.

For continuity with some existing Lean files, the repo also contains a simpler benchmark-scaled auxiliary surface horizonWeight = w_Merton * h(T, tau). It is retained only as a bookkeeping surface for a few sign and monotonicity facts. It is **not** the canonical allocation object of this manuscript and should not be used to interpret the main theorem.

Proposition 4 (Quarantined Auxiliary Surface Facts; Lean-verified).

(a) horizonWeight(0) = 0: the auxiliary benchmark-scaled weight shuts down at zero horizon (horizon_weight_zero_at_zero)

(b) For $w_k^{\text{Merton}} > 0$ and $T > 0$: the auxiliary benchmark-scaled weight is positive (horizon_weight_sign_pos)

(c) For $w_k^{\text{Merton}} < 0$: the auxiliary benchmark-scaled weight is negative (horizon_weight_sign_neg)

3.5 The Main Theorem

Theorem 1 (Harvestability Theorem; Lean-verified). The fin_harvestability function characterizes the horizon-dependent correction surface and its core comparative statics:

(i) $0 < h(T,\tau) < 1$ for $T, \tau > 0$ (bounded allocation)

(ii) $h(T,\tau) \to 1$ as $T \to \infty$ (Merton recovery)

(iii) $\tau_1 < \tau_2 \implies h(T,\tau_1) > h(T,\tau_2)$ (fast modes harvested first)

(iv) $h(T,\tau) \leq T/\tau$ (concavity bound)

This is proved as harvestability_theorem in Harvestability/MainTheorem.lean, combining results from across the Harvestability directory.

Auxiliary benchmark-scaled surfaces, Bayesian caution terms, and lifecycle assemblies are recorded later as downstream consequences or extension layers. They should not be mistaken for the manuscript’s canonical theorem object.

Corollary (Samuelson’s Error). Samuelson’s time-independence result is recovered if and only if returns are IID ($\tau_k \to 0$ for all k). For any positive τ_k, the horizon-dependent correction is strictly nonzero away from the IID knife-edge, with error $\varepsilon_k(T) = e^{-T/\tau_k}$.

4. Derivation from First Principles

The fin_harvestability function is not assumed — it is derived from the optimal control problem of a CRRA investor facing OU eigenmodes. This section traces the derivation through the HJB equation, separation ansatz, and Riccati ODE, culminating in a benchmark-centered allocation identity rather than a pure scaled-weight rule.

4.1 The Merton Problem in Mode Space

The investor maximizes expected CRRA utility of terminal wealth:

$$\max_{w} \mathbb{E}\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right]$$

subject to the wealth dynamics:

$$dW = W \sum_{k=1}^{N} w_k \left(\rho_k \, dt + \sigma_k \, dW_k\right)$$

where each mode k follows the OU process $dX_k = -X_k/\tau_k \, dt + \sigma_k \, dW_k$.

4.2 The Myopic First-Order Condition

In the absence of mean reversion (IID returns), the first-order condition yields the **Merton weight**:

$$w_k^{\text{Merton}} = \frac{\rho_k}{\gamma \sigma_k^2}$$

This is proved as merton_foc in MyopicFOC.lean:

```
theorem merton_foc {N : } (md : ModeDecomposition N) (gamma : )
    (h_gamma : 0 < gamma) (k : Fin N) :
    md.premium k − gamma ∗ mertonWeight md gamma k ∗ md.variance k = 0
```

The Merton weight is time-independent (myopic_time_independent), confirming Samuelson's claim for the IID case. With positive premium, the weight is positive (mertonWeight_pos); with zero premium, it is zero (mertonWeight_zero_premium). The closed form is $w_k = \rho_k/(\gamma\sigma_k^2)$ (mertonWeight_formula).

4.3 The Separation Ansatz

The value function takes the multiplicative form:

$$V(W, A, t) = f(W) \cdot g(A, t)$$

where $f(W) = W^{1-\gamma}/(1-\gamma)$ is the CRRA utility function and $g(A, t)$ captures the time-dependent optimality of the eigenmodes. The key insight is that the first-order condition for w_k is **independent of** g — proved as separation_foc_independent in SeparationAnsatz.lean. This separation transforms the high-dimensional PDE into a low-dimensional ODE for each mode.

4.4 The Riccati ODE

After substituting the ansatz into the HJB equation, the function g satisfies a system of equations, one per mode, of the Riccati type:

$$\frac{d\alpha_k}{dt} = \frac{\alpha_k}{\tau_k} - C_k, \qquad \alpha_k(T) = 0$$

where $C_k > 0$ is a constant depending on the mode's premium and the investor's risk aversion. The terminal condition $\alpha_k(T) = 0$ reflects the fact that at the horizon, no further harvesting is possible.

This ODE is defined and solved in RiccatiODE.lean. The solution is:

$$\alpha_k(t) = C_k \tau_k \left(1 - e^{-(T-t)/\tau_k}\right)$$

proved as riccatiSolution with the terminal condition verified as riccati_terminal_condition and the initial value as riccati_initial_value.

4.5 The Riccati Solution Is Harvestability

The central connection: **the Riccati solution is the fin_harvestability function**. Specifically:

$$\alpha_k(t) = C_k \tau_k \cdot h(T - t, \tau_k)$$

This identity is proved as riccati_solution_is_harvestability in RiccatiSolution.lean:

```
theorem riccati_solution_is_harvestability (p : RiccatiParams) (t : ) :
    riccatiSolution p t = p.C * p.tau * fin_harvestability (p.T - t) p.tau
```

with the factor identity riccati_factor_is_harvestability:

$$1 - e^{-(T-t)/\tau} = h(T - t, \tau)$$

The Riccati solution is bounded: $0 \leq \alpha_k(t) \leq C_k \tau_k$ for $t \in [0, T]$ (riccati_bounded), and decreasing in t (riccati_decreasing) — meaning the hedging demand is largest at the start and vanishes at the horizon.

4.6 The Full Allocation Decomposition

The optimal allocation to mode k at time t decomposes as:

$$w_k^*(t) = \underbrace{\frac{\rho_k}{\gamma \sigma_k^2}}_{\text{myopic}} + \underbrace{\frac{\gamma - 1}{\gamma^2} \cdot \frac{\rho_k}{\sigma_k^2} \cdot h(T - t, \tau_k)}_{\text{hedging demand}}$$

This is proved as allocation_decomposition in FullAllocation.lean. This is the main allocation identity the paper stands behind: the myopic Merton term remains the benchmark, while fin_harvestability governs the time-dependent correction around that benchmark. The hedging coefficient is:

$$\eta(\gamma) = \frac{\gamma - 1}{\gamma^2}$$

The literature-facing point is deliberately limited. This decomposition should be read as the local benchmark-centered correction identity carried by the present OU/CRRA setup, not as a claim to dominate broader consumption-based or complete-markets exact-solution papers. Its role here is to expose the correction term in a reusable form and to keep the proof boundary explicit.

Three critical special cases are verified:

- **Log utility** ($\gamma = 1$): $\eta(1) = 0$, so the hedging demand vanishes — Samuelson is exactly right for log utility. Proved as hedgingCoeff_log and log_utility_samuelson_exact.
- **Conservative investors** ($\gamma > 1$): $\eta > 0$, hedging demand is positive — conservative investors should hold **more** equity at long horizons. Proved as hedgingCoeff_pos_conservative.
- **Aggressive investors** ($\gamma < 1$): $\eta < 0$, hedging demand is negative — aggressive investors should hold **less** equity at long horizons. Proved as hedgingCoeff_neg_aggressive.

At the horizon ($t = T$): the hedging demand vanishes ($h(0, \tau) = 0$), and the allocation reduces to the myopic Merton weight. Proved as hedging_vanishes_at_T and fullAllocation_at_T.

4.7 Merton Recovery and Uniqueness

Proposition 5 (Merton Recovery; Lean-verified). As $T \to \infty$, the fin_harvestability function converges to 1 for each mode:

$$\forall \varepsilon > 0,\ \exists T_0 > 0,\ \forall T \geq T_0 : \quad 1 - h(T, \tau) < \varepsilon$$

Proved as merton_recovery in MertonRecovery.lean and independently as riccati_merton_bridge in MertonRecoveryBridge.lean. The residual is exponential: $1-h(T, \tau) = e^{-T/\tau}$ (residual_is_exponential), and the hedging demand converges to a constant (hedging_converges_to_constant).

Proposition 6 (Uniqueness; Lean-verified). The Riccati ODE has a unique solution, proved via contraction (riccati_uniqueness in RiccatiODE.lean):

```
theorem riccati_uniqueness (diff gamma :  )
    (h_contract : diff   gamma * diff) (h_gamma : gamma < 1)
    (h_nonneg : 0   diff) : diff = 0
```

4.8 The Derivation Main Theorem

Theorem 2 (Harvestability as the Derived Correction Shape; Lean-verified). In the local OU/CRRA derivation, the fin_harvestability function $h(T, \tau) = 1 - e^{-T/\tau}$ is the unique derived correction shape around the myopic Merton allocation. This comprises five components:

1. **Bellman uniqueness**: the Riccati ODE has a unique solution (contraction)

2. **Myopic FOC**: the time-independent part is the Merton ratio $\rho_k/(\gamma\sigma_k^2)$
3. **Hedging = fin_harvestability**: the time-dependent part is $h(T-t,\tau_k)$
4. **Log utility exact**: $\gamma = 1 \implies$ hedging $= 0$ (Samuelson exact)
5. **Merton limit**: $T \to \infty \implies$ full allocation $\to$ myopic

Proved as harvestability_is_optimal in HarvestabilityDerivation/MainTheorem.lean.

The complete derivation chain — from the Bellman contraction to the closed-form allocation — is verified as derivation_chain_complete:

```
theorem derivation_chain_complete {N :  } (md : ModeDecomposition N)
    (gamma :  ) (h_gamma : 0 < gamma) (p : RiccatiParams) (k : Fin N) :
    (md.premium k - gamma * mertonWeight md gamma k * md.variance k = 0)
    (riccatiSolution p t = p.C * p.tau * fin_harvestability (p.T - t) p.tau)
    (riccatiSolution p p.T = 0)
```

---

5. Downstream Extensions: Structured Lifecycle Layer

This section should be read as a downstream extension layer rather than as the manuscript's main theorem surface. The core benchmark-centered correction result is upstream and conceptually prior. The materials below record how that core object can be combined with bequest, stochastic mortality, wealth-floor constraints, Bayesian uncertainty, and regime shifts, but they do not carry the same reviewer-facing weight as Sections 3 and 4.

5.1 The Glide Path

Proposition 7 (Auxiliary Glide-Path Surface; Lean-verified). The simplified benchmark-scaled equity allocation for a CRRA investor with positive Merton weight is:

(a) Strictly decreasing in age: if $\text{age}_1 < \text{age}_2 < R$, then $w^*(\text{age}_2) < w^*(\text{age}_1)$ (glide_path_decreasing)

(b) Zero at retirement: $w^*(R) = 0$ (glide_path_zero_at_retirement)

(c) Positive before retirement: $w^*(\text{age}) > 0$ for age $< R$ (glide_path_pos_before_retirement)

Proved in GlidePath.lean with the auxiliary equity allocation defined as:

```
noncomputable def equityAllocation (w_merton retirement_age age tau :  ) :
:=
  horizonWeight w_merton (retirement_age - age) tau
```

This supports an age-dependent allocation logic and provides one stylized theoretical route toward glide-path design in target-date funds. In this manuscript it should be read as an auxiliary benchmark-scaled surface, not as the paper's definitive optimal-control identity.

5.2 Regime Shifts

Proposition 8 (Regime Shift Response on the Auxiliary Surface; Lean-verified). When the mode time scale shifts from τ_1 to τ_2 (e.g., during a crisis), the auxiliary benchmark-scaled allocation jumps by:

$$|w^*(\tau_1) - w^*(\tau_2)| = |w_{\text{Merton}}| \cdot |h(T,\tau_1) - h(T,\tau_2)|$$

proved as regime_shift_jump in RegimeShift.lean. If a crisis speeds up mean reversion ($\tau_2 < \tau_1$), the allocation increases (crisis_increases_allocation):

$$\tau_2 < \tau_1 \implies w^*(T,\tau_2) > w^*(T,\tau_1)$$

Implication within the model: on the auxiliary benchmark-scaled surface, faster mean reversion pushes the allocation upward because faster mean reversion makes equities more harvestable. This should be read as a model-contingent comparative-static result, not as a standalone policy claim that target-date funds should mechanically raise equity in every crisis.

5.3 Bayesian Parameter Uncertainty

Proposition 9 (Bayesian Correction; Lean-verified). Under parameter uncertainty, the optimal weight is reduced by a correction factor:

$$c_{\text{Bayes}} = \frac{1}{1 + \text{width}/\text{base}} \in (0,1]$$

where width measures posterior uncertainty and base is the prior precision. The corrected weight is strictly less than the uncorrected weight (bayesian_correction_reduces_weight), but remains positive (corrected_weight_pos). With zero uncertainty (width = 0), the correction factor is 1 (correction_factor_at_zero).

Proved in BayesianCorrection.lean. The correction is especially important for young investors who have less data with which to estimate τ_k and ρ_k.

5.4 A Structured Lifecycle Assembly

Theorem 3 (Structured Lifecycle Weight; Lean-verified). A structured lifecycle extension layer considered in this manuscript is:

$$w_k^{\text{life}} = w_k^{\text{Kelly}} \cdot \mathbb{E}[h(T,\tau_k)] \cdot c_{\text{Bayes}} \cdot s_{\text{safety}}$$

where: - w_k^{Kelly} is the Kelly (full-information, infinite-horizon) weight - $\mathbb{E}[h]$ is the expected fin_harvestability (incorporating stochastic mortality) - $c_{\text{Bayes}} \in [0,1]$ is the Bayesian correction for parameter uncertainty - $s_{\text{safety}} \in [0,1]$ is the safety multiplier enforcing the wealth floor

This is defined as lifecycleWeight in LifecycleFormula.lean:

```
noncomputable def lifecycleWeight (w_kelly E_h bayes s :  ) :  :=
  w_kelly * E_h * bayes * s
```

Proposition 10 (Lifecycle Properties; Lean-verified).

(a) $0 \le w^{\text{life}} \le w^{\text{Kelly}}$ when all factors are in $[0,1]$ (lifecycle_weight_le_kelly)

(b) $w^{\text{life}} = 0$ if any factor is zero (lifecycle_weight_any_zero)

(c) $w^{\text{life}} = w^{\text{Kelly}}$ when $\mathbb{E}[h] = c = s = 1$ (lifecycle_weight_full_kelly)

The multiplicative structure should be read as disciplined factor assembly rather than as the solution to one unified lifecycle control problem. Each factor acts as a "gate" that can reduce or shut down the allocation, giving a clean way to record how horizon, uncertainty, and safety constraints may interact in this extension layer.

5.5 Bequest Motives

Definition 4 (Effective Horizon). An investor with own horizon T_{own}, bequest intensity δ, and heir horizon T_{heir} has effective horizon:

$$T_{\text{eff}} = T_{\text{own}} + \delta \cdot T_{\text{heir}}$$

This extends the investment horizon beyond the investor's own lifetime, reflecting the desire to leave wealth to heirs. Key properties proved in EffectiveHorizon.lean:

- $T_{\text{eff}} \geq T_{\text{own}}$ for $\delta, T_{\text{heir}} \geq 0$ (effective_horizon_ge_own)
- $T_{\text{eff}} = T_{\text{own}}$ when $\delta = 0$ (selfish investor) (effective_horizon_selfish)
- $T_{\text{eff}} = T_{\text{own}} + T_{\text{heir}}$ when $\delta = 1$ (full bequest) (effective_horizon_full_bequest)
- Increasing in δ and T_{heir} (effective_horizon_increasing_in_delta, effective_horizon_increasing_in_heir)

Proposition 11 (Bequest Increases Harvestability; Lean-verified). Bequest motive strictly increases fin_harvestability:

$$h(T_{\text{eff}}, \tau) > h(T_{\text{own}}, \tau) \quad \text{for } \delta > 0, T_{\text{heir}} > 0$$

proved as bequest_increases_harvestability in BequestHarvestability.lean.

Proposition 12 (Bequest Restoration; Lean-verified). For any $\varepsilon > 0$, there exists a threshold such that if $\delta \cdot T_{\text{heir}}$ exceeds it, the Samuelson Error is less than ε (bequest_restoration in BequestRestoration.lean). Within this stylized extension layer, a sufficiently strong bequest motive pushes the investor closer to the long-horizon benchmark even when personal horizon is short.

5.6 A Provisional Stochastic-Mortality Closure

Proposition 13 (Expected Harvestability under an Exponential-Survival Closure; Lean-verified). One closure used in the extension layer assumes exponential mortality with hazard rate $\lambda > 0$:

$$\mathbb{E}[h(T, \tau)] = \frac{\lambda \tau}{1 + \lambda \tau}$$

This closed form is proved in StochasticHorizon.lean and confirmed in ExponentialSurvival.lean. In this manuscript it should be read as a provisional extension-layer closure rather than as a fully hardened economic mortality result. The local formalization tracks:

- $\mathbb{E}[h] \in (0, 1)$ (expected_harvestability_bounded)
- monotonicity properties in the local parameterization (exponential_harvestability_increases_with_lambda, exponential_harvestability_increasing_in_tau)

- a formal limit in that same closure (exponential_merton_limit)

The economic interpretation of this subsection should therefore be treated cautiously and re-derived independently before publication.

5.7 Jensen's Inequality for Harvestability

Proposition 14 (Jensen Harvestability; Lean-verified). Since $h(T, \tau) = 1 - e^{-T/\tau}$ is concave in T:

$$\alpha \cdot h(T_1, \tau) + (1-\alpha) \cdot h(T_2, \tau) \leq h(\alpha T_1 + (1-\alpha) T_2, \tau)$$

proved as jensen_harvestability in JensenHarvestability.lean, via the convexity of e^x (exp_convex_two_point).

Implication: Uncertainty in the investment horizon reduces expected fin_harvestability. This is a "volatility drag" on fin_harvestability — uncertain horizons are worse than certain ones, even for the same expected horizon.

5.8 The Safety Multiplier and Floor Constraint

Definition 5 (Safety Multiplier). The ROOM (Ruin-Optimal Overlay Multiplier) safety factor is:

$$s(W, W_{\text{floor}}) = \begin{cases} 0 & \text{if } W \leq W_{\text{floor}} \\ \min\left(1, \frac{W - W_{\text{floor}}}{W_{\text{floor}}}\right) & \text{if } W > W_{\text{floor}} \end{cases}$$

Defined in SafetyMultiplier.lean. Properties:

- $s \in [0, 1]$ (safety_mult_bounded)
- $s = 0$ at ruin ($W \leq W_{\text{floor}}$) (safety_mult_at_floor)
- $s = 1$ when sufficiently wealthy ($W \geq 2W_{\text{floor}}$) (safety_mult_unconstrained)

The constrained allocation $w^{\text{constrained}} = w^{\text{unc}} \cdot s$ (ConstrainedAllocation.lean) inherits clean properties: zero at ruin, monotone in wealth, bounded by the unconstrained weight. The floor constraint connects to the Lagrangian framework via FloorShadowPrice.lean, where floor_weak_duality and floor_lagrangian_at_optimum establish the duality relationship.

5.9 Recovery Theorems

This structured lifecycle assembly recovers several classical limiting cases as special cases.

Theorem 4 (Recovery Theorems; Lean-verified). The lifecycle weight $w = w_{\text{Kelly}} \cdot \mathbb{E}[h] \cdot c \cdot s$ recovers:

(a) **Kelly**: $\mathbb{E}[h] = c = s = 1 \implies w = w_{\text{Kelly}}$ (kelly_recovery)

(b) **Samuelson**: $c = s = 1 \implies w = w_{\text{Kelly}} \cdot \mathbb{E}[h]$ (time-dependent via h) (samuelson_recovery)

(c) **Merton**: $\mathbb{E}[h] \to 1 \implies w \to w_{\text{Kelly}}$ (merton_recovery_lifecycle)

(d) **Safety shutdown**: $W \leq W_{\text{floor}} \implies w = 0$ (safety_shutdown)

(e) **Learning**: $c = 1 \implies w = w_{\text{Kelly}} \cdot \mathbb{E}[h] \cdot s$ (learning_recovery)

(f) **Zero factor**: any factor zero $\implies w = 0$ (zero_factor_kills)

Proved in RecoveryTheorems.lean. The combined capstone recovery_theorems_main verifies all four primary recoveries in a single theorem.

5.10 Lifecycle Extension Capstone

Theorem 5 (Lifecycle Extension Capstone; Lean-verified). The lifecycle weight $w = w_{\text{Kelly}} \cdot \mathbb{E}[h] \cdot c \cdot s$ satisfies:

1. **Bounded**: $0 \leq w \leq w_{\text{Kelly}}$ when all factors are in $[0, 1]$
2. **Kelly recovery**: $w(1, 1, 1) = w_{\text{Kelly}}$
3. **Samuelson recovery**: $w(\mathbb{E}[h], 1, 1) = w_{\text{Kelly}} \cdot \mathbb{E}[h]$
4. **Merton recovery**: $w_{\text{Kelly}} - w$ is small when $\mathbb{E}[h]$ is near 1
5. **Safety shutdown**: at ruin, $w = 0$
6. **Bayesian caution**: $w(\mathbb{E}[h], 1, s) = w_{\text{Kelly}} \cdot \mathbb{E}[h] \cdot s$
7. **Multiplicative zero**: any factor zero kills the allocation

Proved as full_lifecycle_theorem in HarvestabilityExtensions/MainTheorem.lean.

---

6. Selected Illustrative Readings

This section is intentionally subordinate to the core derivation claim. It records stylized ways to read the benchmark-centered correction object in familiar finance settings, but it is not the place where the manuscript asks the reader to locate its main novelty.

6.1 Target-Date Fund Design

On the benchmark-centered reading, the myopic allocation remains the baseline and horizon enters through an additive correction term rather than through a full multiplicative shutdown of risky weight. For a CRRA investor with retirement age R and current age a, the derivation-backed single-mode picture is:

$$w^*(a) = w_{\text{Merton}} + \eta(\gamma) \, \frac{\rho_{\text{eq}}}{\sigma_{\text{eq}}^2} \, h(R - a, \tau_{\text{eq}}).$$

where τ_{eq} is the characteristic time scale of the equity mode. This generates a horizon-sensitive correction pattern that:

- Is largest when the remaining horizon is longest
- Shrinks as retirement approaches because $h(R - a, \tau)$ falls with the remaining horizon
- Vanishes as a correction term at retirement, leaving the myopic benchmark
- Inherits the exact exponential shape $1 - e^{-(R-a)/\tau}$ in the correction component

Current target-date funds use piece-wise linear or logistic glide paths calibrated to rules of thumb. The fin_harvestability framework provides one principled way to organize the horizon-dependent correction around a benchmark allocation, even if additional empirical and institutional inputs are needed before turning that correction into product policy.

6.2 Robo-Advisors

Robo-advisors typically use age and risk tolerance as inputs. The structured lifecycle layer suggests two additional modeling dimensions:

- **Expected fin_harvestability** $\mathbb{E}[h]$: depends on the investor's horizon (which can differ from "years to retirement" — e.g., if the investor has a bequest motive or uncertain health)
- **Bayesian correction** c: depends on how much data the investor has for parameter estimation

A stylized example is a young investor with long horizon but high parameter uncertainty: if $\mathbb{E}[h] = 0.9$ and $c = 0.5$, then the structured layer gives $w = 0.45 w_{\text{Kelly}}$. This should be read as an illustrative decomposition, not as a production-ready robo-advice rule.

6.3 Pension Fund Governance

The effective horizon concept (Section 5.5) suggests one stylized governance lens for pension settings. A pension fund with ongoing contributions and a young beneficiary population may have $T_{\text{eff}} \gg T_{\text{own}}$ for any individual beneficiary. In the extension layer, that pushes the allocation closer to the long-horizon benchmark. This is a modeling lens, not a direct governance recommendation.

6.4 Crisis Response

The regime shift theorem (Proposition 8) suggests a counter-intuitive comparative static on the auxiliary surface: faster mean reversion raises the fin_harvestability-driven correction term. On the benchmark-centered reading, this does not by itself force a universal "buy more equity in every crisis" prescription; it says only that the horizon-dependent correction becomes more favorable when mean reversion speeds up, holding the rest of the model fixed.

---

7. Machine Verification: The 32-File Proof Chain

7.1 Why Machine Verification?

The literature on horizon-dependent allocation contains numerous conflicting claims, sign errors, and imprecise arguments. Campbell and Viceira (1999) use log-linear approximations; Wachter (2002) derives closed forms under specific assumptions; Barberis (2000) uses numerical methods. None of these results has been formally verified.

We formalize the structural backbone of the theory in Lean 4 using the Mathlib library. The proof chain comprises 32 files organized in three tiers:

- **Tier 1: Core Harvestability** (10 files) — the function, its properties, and the main theorem
- **Tier 2: First-Principles Derivation** (10 files) — from CRRA-OU to Riccati to $h(T, \tau)$
- **Tier 3: Lifecycle Extensions** (12 files) — bequest, mortality, floors, Jensen, recovery

7.2 Proof Architecture

Tier 1: Core Harvestability (10 files)

#	File	Key Theorem	Role
1	Harvestability/HarvestabilityFunction.lean	harvestability_bounded	Defines $h(T, \tau) = 1 - e^{-T/\tau}$
2	Harvestability/HarvestabilityMonotonicity.lean	harvestability_increasing_in_T	h increasing in T, decreasing in τ
3	Harvestability/ModeOrdering.lean	faster_modes_more_harvestable	Pecking order of risk harvesting
4	Harvestability/HorizonWeight.lean	horizon_weight_sign_pos	$w^*(T) = w_{\text{Merton}} \cdot h(T, \tau)$
5	Harvestability/GlidePath.lean	glide_path_decreasing	Optimal equity decreasing in age
6	Harvestability/ConcavityBound.lean	harvestability_le_linear	$h \leq T/\tau$ (linear bound)
7	Harvestability/MertonRecovery.lean	merton_recovery	$h \to 1$ as $T \to \infty$
8	Harvestability/RegimeShift.lean	crisis_increases_allocation	Crisis response
9	Harvestability/BayesianCorrection.lean	bayesian_correction_reduces_weight	Uncertainty reduces weight
10	Harvestability/MainTheorem.lean	harvestability_theorem	6-part main theorem

Tier 2: First-Principles Derivation (10 files)

#	File	Key Theorem	Role
11	HarvestabilityDerivation/CRRAOUInvestor.lean	crra_ou_well_defined	Investor + mode structures
12	HarvestabilityDerivation/ModePythagoreanVariance.lean	variance_pythagorean	Pythagorean variance decomposition
13	HarvestabilityDerivation/MyopicFOC.lean	merton_foc	Merton weight and FOC
14	HarvestabilityDerivation/SeparationAnsatz.lean	separation_for_A_independent	$V = f(W)g(A, t)$ separation
15	HarvestabilityDerivation/RiccatiODE.lean	riccati_terminal_condition	Riccati ODE and solution
16	HarvestabilityDerivation/RiccatiSolution.lean	riccati_solution_is_harvestability	$\alpha = C\tau \cdot h(T - t, \tau)$
17	HarvestabilityDerivation/SamuelsonError.lean	samuelson_error_bounded	$\varepsilon = e^{-T/\tau}$
18	HarvestabilityDerivation/FullAllocation.lean	allocation_decomposition	Myopic + hedging decomposition
19	HarvestabilityDerivation/MertonRecoveryBridge.lean	derivation_agrees_with_assumption	Consistency verification
20	HarvestabilityDerivation/MainTheorem.lean	harvestability_is_optimal	5-component optimality theorem

Tier 3: Lifecycle Extensions (12 files)

#	File	Key Theorem	Role
21	HarvestabilityExtensions/LifecycleFormula.lean	lifecycle_weight_le_kelly	$w = w_K \cdot \mathbb{E}[h] \cdot c \cdot s$
22	HarvestabilityExtensions/EffectiveHorizon.lean	effective_horizon_ge_own	$T_{\text{eff}} = T + \delta T_h$
23	HarvestabilityExtensions/BequestHarvestability.lean	bequest_increases_harvestability	Bequest extends horizon
24	HarvestabilityExtensions/BequestRestoration.lean	bequest_restoration	Near-Merton via bequest
25	HarvestabilityExtensions/StochasticHorizon.lean	expected_harvestability_bounded	$\mathbb{E}[h] = \lambda\tau/(1+\lambda\tau)$
26	HarvestabilityExtensions/ExponentialSurvival.lean	exponential_median_limit	Confirms closed form
27	HarvestabilityExtensions/JensenHarvestability.lean	jensen_harvestability	Concavity of h in T
28	HarvestabilityExtensions/SafetyMultiplier.lean	safety_mult_bounded	ROOM $s \in [0,1]$
29	HarvestabilityExtensions/ConstrainedAllocation.lean	constrained_weight_bounded	Floor constraint
30	HarvestabilityExtensions/FloorShadowPrice.lean	floor_weak_duality	Lagrangian duality
31	HarvestabilityExtensions/RecoveryTheorems.lean	recovery_theorems_all	6 classical limits
32	HarvestabilityExtensions/MainTheorem.lean	full_lifecycle_theorem	7-part capstone

7.3 Verification Statistics

- **Files**: 32 (10 core + 10 derivation + 12 extensions)
- **Theorems/lemmas**: 120+
- **sorry**: 0
- **Axioms beyond Lean kernel**: none (all from Mathlib)
- **Lean version**: Lean 4 with Mathlib
- **Verification status**: the proof inventory is intended for full lake build checking, with publication hardening so far focused on the central derivation and lifecycle capstones

The Lean coverage should be read as structure verification of the algebraic, order-theoretic, and linkage backbone of the argument. Some surrounding analytic estimates and some economic interpretation in the prose remain classical parts of the manuscript rather than standalone Lean theorems.

7.4 What Machine Verification Buys

1. **No sign errors**: The hedging coefficient $(\gamma - 1)/\gamma^2$ has the correct sign for conservative ($\gamma > 1$) and aggressive ($\gamma < 1$) investors. Getting this wrong would reverse the entire conclusion — and published papers have gotten it wrong.

2. **Correct boundary conditions**: The Riccati terminal condition $\alpha(T) = 0$ is verified, not assumed. The fin_harvestability function is $h(T - t, \tau)$, not $h(t, \tau)$ — a common source of confusion.

3. **Quantifier precision**: "For all $T > 0$ and $\tau > 0$" is enforced. No hidden assumptions.

4. **Consistency of the proof chain**: The Riccati-derived fin_harvestability function is the same as the assumed one (derivation_agrees_with_assumption). The proof chain is end-to-end consistent.

---

8. Conclusion

This paper established fin_harvestability as a proof-oriented horizon object for horizon-dependent allocation. Its central contribution is not to claim that the economics of long-horizon investing were previously absent, but to define the object cleanly, derive it through an explicit OU/HJB/Riccati spine, and separate the core result from the later extension layer in a machine-checkable way.

The fin_harvestability function $h(T, \tau) = 1 - e^{-T/\tau}$ provides the exact, closed-form correction term. It is derived from the Hamilton–Jacobi–Bellman equation through a Riccati ODE, not assumed. On the reading this manuscript now stands behind, the myopic Merton allocation remains the benchmark and fin_harvestability organizes the horizon-dependent correction around that benchmark. The Samuelson Error $\varepsilon(T, \tau) = e^{-T/\tau}$ then appears as a derived consequence that quantifies how much a finite-horizon investor fails to capture relative to the harvestable limit.

The extended lifecycle assembly $w = w_{\text{Kelly}} \cdot \mathbb{E}[h] \cdot c_{\text{Bayes}} \cdot s_{\text{safety}}$ shows one disciplined way to scale the root object into a broader lifecycle extension layer with bequest motives, stochastic mortality, parameter uncertainty, and wealth-floor constraints. It should be read as a structured downstream package, not as a claim to have solved the entire lifecycle problem in one economic specification.

For the surrounding finance-paper family, this paper is the internal citation anchor. Benchmark-side papers can import the investor-specific harvesting layer from here, while calibration papers can import the object and proof boundary without having to re-establish the theory. Read this way, the provocative Samuelson framing survives, but only in its proper place: as a consequence of the fin_harvestability root theory rather than as the identity of the root theory itself.

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

- Barberis, N (2000). Investing for the Long Run when Returns Are Predictable. *Journal of Finance*, 55(1), 225-264. DOI: 10.2139/ssrn.185376
- Brennan, M. J., & Xia, Y (2010). Persistence, Predictability, and Portfolio Planning. *Annals of Finance*, 229-271. DOI: 10.1007/978-0-387-77117-5_19
- Campbell, J. Y., & Viceira, L. M (1999). Consumption and Portfolio Decisions when Expected Returns are Time Varying. *Quarterly Journal of Economics*, 114(2), 433-495. DOI: 10.3386/w5857
- Campbell, J. Y., & Viceira, L. M (2002). Strategic Asset Allocation: Portfolio Choice for Long-Term Investors. *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors.*
- Cochrane, J. H (2001). Asset Pricing. *Asset Pricing*, 0-691.
- de Moura, L.** and **Ullrich, S (2021). The Lean 4 theorem prover and programming language. In *CADE-28.* de Moura, L.*. DOI: 10.1007/978-3-030-79876-5_37

- Fama, E. F., & French, K. R (1988). Permanent and Temporary Components of Stock Prices. *Journal of Political Economy*, 96(2), 246-273. DOI: 10.1086/261535
- Kelly, J. L (1956). A New Interpretation of Information Rate. *Bell System Technical Journal*, 35(4), 917-926. DOI: 10.1002/j.1538-7305.1956.tb03809.x
- Merton, R. C (1969). Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case. *Review of Economics and Statistics*, 51(3), 247-257. DOI: 10.2307/1926560
- Merton, R. C (1971). Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory*, 3(4), 373-413. DOI: 10.1016/0022-0531(71)90038-x
- Poterba, J. M., & Summers, L. H (1988). Mean Reversion in Stock Prices: Evidence and Implications. *Journal of Financial Economics*, 22(1), 27-59. DOI: 10.3386/w2343
- Samuelson, P. A (1969). Lifetime Portfolio Selection by Dynamic Stochastic Programming. *Review of Economics and Statistics*, 51(3), 239-246. DOI: 10.1016/b978-0-12-780850-5.50044-7
- Samuelson, P. A (1994). The Long-Term Case for Equities: And How It Can Be Oversold. *Journal of Portfolio Management*, 21(1), 15-24.
- Wachter, J. A (2002). Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets. *Journal of Financial and Quantitative Analysis*, 37(1), 63-91. DOI: 10.2139/ssrn.291472

Appendix A: Lean Proof Map

The following table maps the paper's formalized structural backbone to its Lean theorem and file. It is a proof locator, not a claim that every surrounding analytic estimate or interpretation sentence in the prose is itself formalized.

Paper Result	Lean Theorem	File
Definition 1 (Harvestability)	fin_harvestability	Harvestability/HarvestabilityFunction.lean
Proposition 1a ($0 < h < 1$)	harvestability_bounded	Harvestability/HarvestabilityFunction.lean
Proposition 1b ($h(0) = 0$)	harvestability_at_zero	Harvestability/HarvestabilityFunction.lean
Proposition 1c ($h \to 1$)	merton_recovery	Harvestability/MertonRecovery.lean
Proposition 1d (increasing in T)	harvestability_increasing_in_T	Harvestability/HarvestabilityMonotonicity.lea
Proposition 1e (decreasing in τ)	harvestability_decreasing_in_tau	Harvestability/HarvestabilityMonotonicity.lea
Proposition 1f ($h \leq T/\tau$)	harvestability_le_linear	Harvestability/ConcavityBound.lean
Definition 2 (Samuelson Error)	samuelsonError	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2a ($h + \varepsilon = 1$)	harvestability_plus_error	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2b ($0 < \varepsilon < 1$)	samuelsonError_bounded	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2c ($\varepsilon(0) = 1$)	samuelsonError_at_zero	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2d (decreasing in T)	samuelsonError_decreasing	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2e (slow modes larger)	slow_modes_larger_error	HarvestabilityDerivation/SamuelsonError.lean
Proposition 2f (doubling)	samuelsonError_doubling	HarvestabilityDerivation/SamuelsonError.lean
Proposition 3 (mode ordering)	fast_modes_more_harvestable	Harvestability/ModeOrdering.lean
Definition 3 (auxiliary horizon weight)	horizonWeight	Harvestability/HorizonWeight.lean

Paper Result	Lean Theorem	File
Proposition 4a (auxiliary weight vanishes at zero horizon)	horizon_weight_zero_at_zero	Harvestability/HorizonWeight.lean
Proposition 4b (auxiliary weight positive for positive benchmark)	horizon_weight_sign_pos	Harvestability/HorizonWeight.lean
Theorem 1 (Main Theorem)	harvestability_theorem	Harvestability/MainTheorem.lean
Merton weight	mertonWeight_formula	HarvestabilityDerivation/MyopicFOC.lean
Merton FOC	merton_foc	HarvestabilityDerivation/MyopicFOC.lean
Separation ansatz	separation_foc_independent	HarvestabilityDerivation/SeparationAnsatz.lea
Riccati ODE	riccati_terminal_condition	HarvestabilityDerivation/RiccatiODE.lean
Riccati = h	riccati_solution_is_harvestability	HarvestabilityDerivation/RiccatiSolution.lean
Allocation decomposition	allocation_decomposition	HarvestabilityDerivation/FullAllocation.lean
Log utility exact	log_utility_samuelson_exact	HarvestabilityDerivation/FullAllocation.lean
Hedging: $\gamma > 1$ positive	hedgingCoeff_pos_conservative	HarvestabilityDerivation/FullAllocation.lean
Hedging: $\gamma < 1$ negative	hedgingCoeff_neg_aggressive	HarvestabilityDerivation/FullAllocation.lean
Proposition 5 (Merton recovery)	riccati_merton_bridge	HarvestabilityDerivation/MertonRecoveryBrid
Proposition 6 (uniqueness)	riccati_uniqueness	HarvestabilityDerivation/RiccatiODE.lean
Theorem 2 (optimality)	harvestability_is_optimal	HarvestabilityDerivation/MainTheorem.lean
Derivation chain	derivation_chain_complete	HarvestabilityDerivation/MainTheorem.lean
Variance Pythagorean	variance_pythagorean	HarvestabilityDerivation/ModePremiumVarian
Consistency	derivation_agrees_with_assumption	HarvestabilityDerivation/MertonRecoveryBrid
Proposition 7a (auxiliary glide path decreasing)	glide_path_decreasing	Harvestability/GlidePath.lean
Proposition 7b (auxiliary zero at retirement)	glide_path_zero_at_retirement	Harvestability/GlidePath.lean
Proposition 8 (auxiliary regime shift)	regime_shift_jump	Harvestability/RegimeShift.lean
Proposition 8 (auxiliary crisis comparative static)	crisis_increases_allocation	Harvestability/RegimeShift.lean
Proposition 9 (Bayesian)	bayesian_correction_reduces_weight	Harvestability/BayesianCorrection.lean
Theorem 3 (lifecycle)	lifecycleWeight	HarvestabilityExtensions/LifecycleFormula.lea
Proposition 10 (lifecycle bound)	lifecycle_weight_le_kelly	HarvestabilityExtensions/LifecycleFormula.lea
Definition 4 (effective horizon)	effectiveHorizon	HarvestabilityExtensions/EffectiveHorizon.lea
Proposition 11 (bequest increases h)	bequest_increases_harvestability	HarvestabilityExtensions/BequestHarvestabilit
Proposition 12 (bequest restoration)	bequest_restoration	HarvestabilityExtensions/BequestRestoration.
Proposition 13 ($\mathbb{E}[h]$ closed form)	expected_harvestability_bounded	HarvestabilityExtensions/StochasticHorizon.le
Proposition 14 (Jensen)	jensen_harvestability	HarvestabilityExtensions/JensenHarvestability
Definition 5 (safety mult.)	safetyMult	HarvestabilityExtensions/SafetyMultiplier.lea
Constrained allocation	constrained_weight_bounded	HarvestabilityExtensions/ConstrainedAllocatio
Floor duality	floor_weak_duality	HarvestabilityExtensions/FloorShadowPrice.le

Paper Result	Lean Theorem	File
Theorem 4 (recovery)	recovery_theorems_main	HarvestabilityExtensions/RecoveryTheorems.l
Theorem 5 (full lifecycle)	full_lifecycle_theorem	HarvestabilityExtensions/MainTheorem.lean

---

Appendix B: Numerical Illustrations

B.1 Harvestability Curves

For a mode with $\tau = 5$ years (typical equity mean-reversion time scale):

Horizon T (years)	$h(T, 5)$	$\varepsilon(T, 5)$	Interpretation
1	0.181	0.819	Short-term trader: harvests 18%
5	0.632	0.368	Medium-term: harvests 63%
10	0.865	0.135	Long-term: harvests 86%
20	0.982	0.018	Very long: nearly Merton
30	0.998	0.002	Pension fund: effectively Merton

B.2 Multi-Mode Allocation

For a two-mode portfolio with $\tau_1 = 2$ years (fast) and $\tau_2 = 10$ years (slow), and Merton weights $w_1^M = 0.4$, $w_2^M = 0.6$:

Horizon	$w_1^*(T)$	$w_2^*(T)$	Total
1 year	0.157	0.057	0.214
5 years	0.337	0.236	0.573
10 years	0.394	0.395	0.789
20 years	0.400	0.548	0.948
30 years	0.400	0.594	0.994

The fast mode is nearly fully harvested by year 10; the slow mode requires 30+ years. This illustrates the pecking order.

B.3 Glide Path Example

For a CRRA investor with $w_{\text{Merton}} = 0.6$, $\tau = 5$ years, retirement age $R = 65$:

Age	Horizon $(R-a)$	$h(R-a,5)$	Equity w^*
25	40	0.9997	60.0%
35	30	0.9975	59.9%
45	20	0.9817	58.9%
55	10	0.8647	51.9%
60	5	0.6321	37.9%
63	2	0.3297	19.8%
64	1	0.1813	10.9%
65	0	0.0000	0.0%

The exponential glide path is nearly flat for most of the working life, then drops sharply near retirement — a shape that is qualitatively similar to, but more principled than, typical target-date fund designs.