Calibrating Harvestability

When Can Investors Reliably Capture the Risk Premium?

From the canonical horizon object to documented defaults and allocation use

Dr. Tamás Nagy

tnagyphd@gmail.com

Draft • March 2026

Paper Role in the Research Program

This paper is the **calibration and application companion** to the canonical fin_harvestability theory paper.

- topics/fin_harvestability/paper.md is the root paper: it defines fin_harvestability, derives it from the OU/HJB/Riccati backbone, and develops the proof-first lifecycle extension.
- This paper does not replace that root theory paper. It imports the object and its interpretation boundary, then specializes to calibration, documented defaults, and practical use.
- topics/fin_pricing_is_allocation/paper.md and its submission use the same fin_harvestability object at the bridge level, where investor-specific filtering is layered onto a common benchmark theorem.
- This paper should therefore be read as the reviewer-friendly quantitative note in the fin_harvestability family, not as the foundational derivation paper.

Executive Summary (Non-Technical)

Investors are routinely told that risky assets reward patience. But the central practical question is usually left vague: **how long must an investor hold before the expected premium becomes reliably capturable rather than mostly noise?** A one-year horizon and a thirty-year horizon are clearly different, yet portfolio theory rarely names the object that measures that difference.

This paper does not serve as the root derivation of that object. Instead, it takes **fin_harvestability** from the canonical theory paper and asks how the object should be interpreted, approximated, documented, and calibrated for practical quantitative use. Harvestability is the fraction of an asset's expected excess return that can be treated as realistically capturable at a given horizon. If fin_harvestability is near zero, volatility still dominates. If it is near one, the premium is close to being fully available to a patient investor.

The paper's main contribution is to make the concept **operational, documented, and calibratable**. We restate fin_harvestability as $h(T,\tau) = 1 - e^{-T/\tau}$, where T is the investment horizon and τ is the asset's characteristic time scale, then show how τ can be interpreted and computed outside the full proof spine. In a mean-reverting model, τ is the half-life of the mode. In an implementable asset-class approximation, τ becomes $1/\text{Sharpe}^2 = \sigma^2/\pi^2$, where σ is volatility and π is the risk premium.

That second route matters operationally. It lets us calibrate fin_harvestability from familiar quantities and document default values for Cash, Bonds, Equity, and Alternatives. The result is a practical bridge from return statistics to horizon-aware allocation.

The paper is deliberately narrow. It does **not** claim to solve the full lifecycle problem or to explain all portfolio heterogeneity. Its aim is more focused: to define fin_harvestability precisely, derive it from a simple theoretical backbone, show how τ can be computed, and make the concept usable in portfolio design and advisory work.

Abstract

Starting from the canonical fin_harvestability object $h(T, \tau) = 1 - e^{-T/\tau}$, this paper studies the calibration problem: when does expected return begin to dominate volatility strongly enough that the premium can be treated as genuinely available to the investor, and how should the time scale τ be documented in practical use? The root derivation lives in the companion theory paper topics/fin_harvestability/paper.md; the role of the present paper is narrower. We derive τ from two usable routes. Under Ornstein-Uhlenbeck dynamics, τ is the mean-reversion half-life. Under a geometric Brownian motion approximation, $\tau = \sigma^2/\pi^2 = 1/\text{Sharpe}^2$, where π is excess return and σ is volatility. The second route yields a simple calibration formula when full eigenmode estimation is unavailable. We document default τ values for Cash, Bonds, Equity, and Alternatives, discuss empirical calibration from return series, and show how fin_harvestability enters lifecycle allocation and advisory methodology. The paper is therefore a calibration and application companion to the proof-first fin_harvestability root, not a replacement for that root theory paper.

---

1. Introduction

1.1 The Problem

The risk premium is the reward for bearing risk, but it is not paid in a deterministic stream. At short horizons, realized outcomes are dominated by volatility. At longer horizons, the expected premium becomes more visible. Everyone agrees with this qualitatively. The problem is that the bridge from that intuition to an operational quantity is usually missing.

This paper asks two concrete questions:

1. What fraction of the expected premium is reliably harvestable at horizon T?
2. How should the associated time scale τ be defined and calibrated?

1.2 Our Contribution

This paper contributes four things:

1. **An explicit operational restatement** of fin_harvestability: $h(T, \tau) = 1 - e^{-T/\tau}$, imported from the canonical root paper and used here as a practical quantitative object.
2. **Two routes for** τ: an OU route, where τ is structural, and a GBM route, where $\tau = \sigma^2/\pi^2$.
3. **A calibration layer** for common asset classes with documented defaults and reporting discipline.

4. **A practical interpretation layer** for lifecycle allocation and advisory use.

1.3 Relation to the Companion Papers

The full first-principles derivation appears in topics/fin_harvestability/paper.md, the canonical fin_harvestability root paper, where $h(T, \tau)$ emerges from a CRRA investor's HJB problem under OU eigenmode dynamics. The contribution here is different: we take that canonical object and proof boundary as given, make the interpretation explicit, derive an implementable approximation for τ, and turn the object into something that can be calibrated and used. In that sense, this paper sits between theory and application.

The benchmark-side companion is topics/fin_pricing_is_allocation/paper.md, which uses fin_harvestability as the investor-specific filter layered on top of a common equilibrium benchmark. The present paper does not attempt to reproduce that bridge theorem. Its role is narrower: define the calibration route clearly enough that the fin_harvestability object can travel into reviewer-facing quantitative work.

---

2. Definition

2.1 Formal Definition

Definition 1 (Harvestability). The **fin_harvestability** of mode k at horizon T is

$$h_k(T, \tau_k) = 1 - e^{-T/\tau_k}$$

where $T \geq 0$ is the investment horizon (years) and $\tau_k > 0$ is the mode's characteristic time scale (years).

Interpretation. $h_k(T)$ is the fraction of mode k's risk premium that can be treated as realistically capturable at horizon T. It is not the realized return itself, and it is not a probability. It is a horizon-dependent scaling factor on the premium.

Properties (proved in LeanProofs/Harvestability/HarvestabilityFunction.lean):

- $h(0, \tau) = 0$
- $0 \leq h(T, \tau) < 1$ for finite T
- $h(T, \tau) \to 1$ as $T \to \infty$
- h is strictly increasing in T and strictly decreasing in τ

2.2 The Risk Premium

The **risk premium** (kockázati prémium) is

$$\pi_k = \mu_k - r$$

where μ_k is the expected return of mode k and r is the risk-free rate. Harvestability answers the question:

Of the premium π_k, how much should a horizon-T investor treat as genuinely available rather than still dominated by noise?

2.3 How to Read h

The object is easiest to understand by examples:

- $h = 0$: effectively none of the premium is harvestable at that horizon.
- $h = 0.25$: only a small part of the premium is reliably available; volatility still dominates.
- $h = 0.50$: roughly half of the premium is in the "reliably capturable" region.
- $h = 0.90$: most of the premium is available to a sufficiently patient investor.

If an asset has annual premium $\pi = 5\%$ and $h(T) = 0.40$, then the horizon-T investor should think of roughly 2% of that premium as reliably harvestable and the remainder as still too noisy to treat as fully usable in allocation.

---

3. Derivation of τ

3.1 Route A: Ornstein–Uhlenbeck (Primary)

When returns follow OU dynamics per eigenmode:

$$dA_k = -\frac{1}{\tau_k}(A_k - \bar{A}_k)\,dt + \sigma_k\,dW_k$$

the parameter τ_k is the **mean-reversion half-life**. The fin_harvestability function then appears in the HJB solution: the intertemporal hedging term is proportional to $1 - e^{-(T-t)/\tau_k}$. In this route, τ_k is a structural feature of the return process, not a tuning parameter.

When to use: Eigenmode decomposition of asset returns with estimable autocorrelation.

3.2 Route B: Geometric Brownian Motion (Asset-Class Proxy)

When returns are approximated as iid with annual excess return π_k and annual volatility σ_k, the cumulative return over T years has:

- Mean: $(\mu_k - r)T = \pi_k T$
- Std dev: $\sigma_k\sqrt{T}$

The **coefficient of variation** (noise-to-signal) is

$$\text{CV}(T) = \frac{\sigma_k\sqrt{T}}{\pi_k T} = \frac{1}{S_k\sqrt{T}}$$

where $S_k = \pi_k/\sigma_k$ is the Sharpe ratio.

Definition of : The time at which CV = 1 (std dev equals mean):

$$1 = \frac{1}{S_k\sqrt{\tau_k}} \quad \Rightarrow \quad \tau_k = \frac{1}{S_k^2} = \frac{\sigma_k^2}{\pi_k^2}$$

Interpretation: τ is the horizon at which the cumulative expected premium and the cumulative standard deviation are of the same order. For $T < \tau$, noise dominates. For $T > \tau$, signal begins to dominate.

This gives an operational definition of τ even when no eigenmode mean-reversion estimate is available. The same functional form $h(T) = 1 - e^{-T/\tau}$ then turns the reliability time scale into a fin_harvestability curve.

3.3 Route C: Cash (Special Case)

Cash has $\pi \approx 0$, $\sigma \approx 0$. Set $\tau_{\text{cash}} = \tau_{\text{min}}$ (e.g. 0.5 years) so that $h_{\text{cash}}(T) \approx 1$ for any $T \geq 1$.

---

4. Calibration

4.1 Formula

$$\tau_k = \frac{\sigma_k^2}{\pi_k^2}$$

4.2 Illustrative Values

Asset class	(p.a.)	(p.a.)	S = /	(years)
Cash	0	0.01	—	0.5 (min)
Bonds	0.01	0.035	0.29	12.3
Equity	0.05	0.16	0.31	10.2
Alternatives	0.03	0.10	0.30	11.1 + lock-up

Alternatives: Add lock-up years (e.g. 5) for illiquid vehicles.

4.3 Documented Defaults

Mode	(years)	Provenance
Cash	0.5	_min; cash is instantly harvestable
Bonds	12	1%, 3.5%
Equity	10	5%, 16%
Alternatives	16	3%, 10% + 5y lock-up

Population: Developed markets. **Period:** Representative long-horizon (e.g. 2000–2024 proxy). Recalibrate from live data when available.

4.4 Empirical Calibration

From historical returns, estimate

$$\hat{\pi}_k = \bar{R}_k - r, \qquad \hat{\sigma}_k = \sqrt{\text{Var}(R_k)},$$

then compute

$$\hat{\tau}_k = \frac{\hat{\sigma}_k^2}{\hat{\pi}_k^2}.$$

The calibration is only meaningful if the paper or implementation also documents:

- the asset universe,
- the sample period,
- the frequency,
- the risk-free benchmark,
- and any illiquidity adjustment such as lock-up years.

4.5 What the Defaults Mean

The default values summarize a simple intuition:

- **Cash**: effectively immediately harvestable.
- **Bonds**: small premium and moderate volatility imply a relatively long reliability horizon.
- **Equity**: higher premium compensates for higher volatility, so τ is not as large as raw volatility alone would suggest.
- **Alternatives**: moderate Sharpe plus illiquidity produces a long τ.

---

5. Application

5.1 Lifecycle Allocation

The optimal allocation to mode k is

$$w_k^* = w_k^{\text{benchmark}} \cdot h_k(T_{\text{eff}}, \tau_k) \cdot s_k$$

where $T_{\text{eff}} = T_{\text{own}} + \delta \cdot T_{\text{heir}}$ (effective horizon including bequest) and s_k is a safety multiplier.

5.2 Advisory Methodology

Harvestability feeds into the NDVR advisory methodology through the rule

$$w_k^* = w_k^{\text{benchmark}} \cdot h_k(T_{\text{eff}}, \tau_k) \cdot s_k,$$

where the benchmark is common, the effective horizon is client-specific, and s_k captures safety and implementation constraints. The point is not to infer the client's full utility function. The point is to explain which premia are realistically harvestable for that client.

5.3 Example: Bitcoin

For Bitcoin with $\pi \approx 30\%$ and $\sigma \approx 70\%$, the GBM route gives $S \approx 0.43$ and $\tau \approx 5.4$ years. Then:

- $h(1) \approx 0.17$
- $h(5) \approx 0.61$
- $h(10) \approx 0.85$

So under these assumptions, Bitcoin has low short-horizon fin_harvestability and much higher long-horizon fin_harvestability. The caveat is obvious: parameter uncertainty is large, so this example is illustrative rather than definitive.

---

6. Limitations and Non-Claims

- **OU vs GBM:** The OU route is conceptually primary. The GBM route is a practical approximation.
- **Estimated inputs:** π and σ are estimated, not known. Small changes in Sharpe can move τ materially.
- **Regime dependence:** τ is not necessarily stable across regimes.
- **Not a full lifecycle theory:** This paper does not rederive the full bequest, mortality, Bayesian, and safety extensions.
- **Not a universal explanation:** The paper does not claim that all portfolio heterogeneity or all horizon effects reduce to fin_harvestability.

---

7. Conclusion

Harvestability provides a clean answer to a simple but under-specified question: when does a risky premium become reliably available to an investor with horizon T? The answer is a horizon object $h(T, \tau)$, together with a time scale τ that can be interpreted structurally under mean reversion and computed operationally as σ^2/π^2 under a GBM approximation.

The value of the framework is not that it replaces all of lifecycle finance. Its value is that it introduces a portable object linking horizon, volatility, Sharpe ratio, and allocation. That makes it useful both as a theory companion and as a practical allocation input.

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

- Nagy, T. (2026). Harvestability. *Working paper*.
- 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.
- 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

Appendix: Implementation

The calibration module examples/ndvr_harvestability_calibration.py implements:

- tau_from_sharpe(sigma, pi, lock_up) — from GBM formula
- tau_from_returns(returns, rf, lock_up) — empirical calibration
- fin_harvestability (T, tau) — $h(T, \tau) = 1 - e^{-T/\tau}$
- DEFAULT_MODE_PARAMS — documented defaults

See research_notes/harvestability_definition_and_calibration.md for full documentation.