

Certified Alpha Capacity:

STATISTICAL EVIDENCE, ECONOMIC LIFETIME, AND ARBITRAGE UNDER DECAY

NICOLÒ BONACORSI

Department of Applied Physics and Applied Mathematics

Columbia University

New York, NY 10027, USA

`nb3328@columbia.edu`

ORCID: [0009-0005-0479-3102](#)

30 September 2026

Abstract

In this paper we study whether a trading signal can accumulate enough statistical evidence for reliable deployment before its economic value decays. We solve the problem by measuring statistical evidence and remaining opportunity on a common Kullback–Leibler information scale, obtaining an exact feasibility threshold in a canonical Gaussian model and general information lower bounds beyond it. We show that finite information budgets generate survival frontiers, search penalties, and a market equilibrium in which arbitrage activity shortens the lifetime of certifiable opportunities; a retrospective funding-rate study illustrates the role of persistence.

AMS 2020 Subject Classification: Primary 62L10, 62L15. Secondary 60G40, 62M10.

Keywords: Certified Alpha Capacity, alpha decay, sequential testing, information theory, quantitative finance, optimal stopping, multiple testing.

1. Introduction

A trading signal can be statistically detectable and economically attractive at a given date and still fail as a research project: by the time enough evidence has accumulated to justify deployment, the opportunity may have decayed. The relevant constraint is therefore temporal as well as statistical. A validation rule demands evidence, while the signal can generate that evidence only for as long as it remains economically useful.

We introduce *Certified Alpha Capacity* (CAC) to compare these two quantities. The first is the statistical information that a signal can generate over its remaining economic lifetime. The second is the information required by a declared certification rule, including its false-deployment probability, power target, and multiplicity correction. The economic boundary is the point at which deploying the signal ceases to have positive incremental value. Measuring Kullback–Leibler (KL) information relative to that boundary puts statistical learning and economic decay on a common clock.

The question sits at the intersection of several established literatures. Sequential analysis studies how evidence accumulates and when a statistical decision should be made [2–6]. Stochastic-deadline models make the time available for learning part of the decision problem [14–17], while recent work studies sequential activation of alpha signals under decay and implementation costs

[41]. Empirical finance treats search over large strategy libraries as a multiple-testing and selection problem [9–11]. A separate literature documents post-discovery alpha decay and the role of arbitrage capital, implementation frictions, and market competition [19, 24–26, 32]. CAC links these strands by treating economic lifetime as a finite information budget and then allowing market activity to change that budget.

The analysis proceeds from this comparison. In the canonical Gaussian experiment, the lifetime-information budget yields an exact boundary for reliable certification; a KL data-processing inequality gives a necessary information bound beyond that benchmark. Mapping the information requirement into signal-strength and persistence coordinates produces the *Alpha Survival Frontier*, which also makes the cost of screening many candidates explicit. The same information clock converts statistical learning into economic value under quadratic trading costs. When arbitrage activity shortens signal lifetime, the certification boundary becomes an equilibrium restriction on entry and crowding. Heterogeneous implementation costs and capacity then determine how that equilibrium relation appears in cross-sectional data.

Two empirical exercises illustrate the distinction between lifetime information and implementation technology. A cross-sectional anomaly benchmark examines the equilibrium prediction when the implementation state is unobserved. A funding-rate study asks whether a pre-outcome estimate of signal lifetime contains information about subsequent persistence beyond current signal magnitude. These exercises are used to study separate implications of the theory.

Section 2 introduces the information clock and its economic interpretation. Section 3 derives the certification boundary. Section 4 links information to economic value, and Section 5 develops the survival frontier and search effects. Section 6 makes lifetime endogenous through arbitrage and develops the cross-sectional identification result. Section 7 reports the funding-rate illustration. Constrained stopping, unknown scale, predictable regressors, and random economic lifetimes are treated in the appendices.

2. Model

2.1. Probabilistic setup and information time

We fix a measurable space $(\Omega, \mathcal{F})$ carrying an observed return process $Y = (Y_t)_{t \geq 0}$ and a real-valued signal process $X = (X_t)_{t \geq 0}$. For $t \geq 0$ let

$$\mathcal{F}_t^0 := \sigma\{(X_s, Y_s) : 0 \leq s \leq t\}.$$

We consider two probability measures $\mathbb{P}_0$ and $\mathbb{P}_1$ on $(\Omega, \mathcal{F})$ and assume that their restrictions to $\mathcal{F}_t^0$ are equivalent for every finite t . We write

$$\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$$

for the usual right-continuous augmentation of the raw observation filtration; local equivalence ensures that the null sets used in this augmentation are the same under the two laws. All stopping times below are taken with respect to the filtration that is explicitly indicated. Expectations under $\mathbb{P}_0$ and $\mathbb{P}_1$ are denoted by $\mathbb{E}_0$ and $\mathbb{E}_1$, respectively.

The process X is assumed to be $\mathbb{F}$ -predictable. Under the alternative law $\mathbb{P}_1$ the observed channel satisfies

$$dY_t = \theta_t X_t dt + \sigma dB_t^{(1)}, \quad \sigma > 0, \quad (2.1)$$

where $B^{(1)}$ is a $(\mathbb{P}_1, \mathbb{F})$ -Brownian motion and $\theta = (\theta_t)_{t \geq 0}$ is predictable. Under the zero-boundary null law $\mathbb{P}_0$,

$$dY_t = \sigma dB_t^{(0)}, \quad (2.2)$$

where $B^{(0)}$ is a $(\mathbb{P}_0, \mathbb{F})$ -Brownian motion. We impose the local square-integrability and exponential-martingale conditions needed for Girsanov's theorem; on a bounded horizon, Novikov's condition is a sufficient assumption.

The accumulated signal-to-noise clock is

$$A_t := \frac{1}{\sigma^2} \int_0^t \theta_s^2 X_s^2 ds, \quad t \geq 0. \quad (2.3)$$

It is continuous and nondecreasing. Let T denote the end of the opportunity's economic lifetime; we refer to T as the *economic death time* of the opportunity. In the main deterministic-horizon development T may be deterministic; when random, we assume it is an $\mathbb{F}$ -stopping time. Put

$$H := A_T.$$

When H is deterministic, we denote its constant value by A . Thus A_t always denotes the information-clock process, whereas unsubscripted A denotes a deterministic terminal information horizon. For every $s \geq 0$ define the right-continuous generalized inverse

$$\tau_s := \inf\{t \geq 0 : A_t > s\} \wedge T. \quad (2.4)$$

and set

$$\mathcal{G}_s := \mathcal{F}_{\tau_s}, \quad \mathbb{G} = (\mathcal{G}_s)_{s \geq 0}.$$

Thus $\mathbb{G}$ is the observation filtration expressed in information time; it is constant after the realized terminal clock H , and calendar intervals on which A is flat are collapsed by the time change.

Let

$$\Lambda_t := \log \left. \frac{d\mathbb{P}_1}{d\mathbb{P}_0} \right|_{\mathcal{F}_t}$$

be the log-likelihood ratio. Under the two laws it has the representations

$$\Lambda_t = \int_0^t \frac{\theta_u X_u}{\sigma} dB_u^{(0)} - \frac{1}{2} A_t, \quad \mathbb{P}_0\text{-a.s.}, \quad (2.5)$$

$$\Lambda_t = \int_0^t \frac{\theta_u X_u}{\sigma} dB_u^{(1)} + \frac{1}{2} A_t, \quad \mathbb{P}_1\text{-a.s.} \quad (2.6)$$

For $0 \leq s \leq H$, set $Z_s := \Lambda_{\tau_s}$. The time-change theorem yields

$$Z_s = \tilde{B}_s^{(1)} + \frac{s}{2} \quad (\mathbb{P}_1), \quad Z_s = \tilde{B}_s^{(0)} - \frac{s}{2} \quad (\mathbb{P}_0), \quad (2.7)$$

where $\tilde{B}^{(i)}$ is a $(\mathbb{P}_i, \mathbb{G})$ -Brownian motion. Equivalently,

$$\left. \frac{d\mathbb{P}_1}{d\mathbb{P}_0} \right|_{\mathcal{G}_s} = e^{Z_s}. \quad (2.8)$$

Let $\mathcal{G}_\infty := \sigma(\bigcup_{s \geq 0} \mathcal{G}_s)$. For any bounded $\mathbb{G}$ -stopping time S , the stopped sigma-field is

$$\mathcal{G}_S := \{E \in \mathcal{G}_\infty : E \cap \{S \leq s\} \in \mathcal{G}_s \text{ for every } s \geq 0\}.$$

We write $\mathbb{P}_i^S$ for the restriction of $\mathbb{P}_i$ to $(\Omega, \mathcal{G}_S)$, $i \in \{0, 1\}$. When $A < \infty$ is deterministic, we use

$$\mathcal{T}_A := \{S : S \text{ is a } \mathbb{G}\text{-stopping time and } 0 \leq S \leq A\}$$

for the admissible information-time stopping rules.

We use

$$D_{\text{KL}}(\mathbb{P} \parallel \mathbb{Q}) := \mathbb{E}_{\mathbb{P}} \left[\log \frac{d\mathbb{P}}{d\mathbb{Q}} \right]$$

for Kullback–Leibler divergence whenever $\mathbb{P} \ll \mathbb{Q}$, and

$$\text{kl}(p, q) := p \log \frac{p}{q} + (1 - p) \log \frac{1 - p}{1 - q}, \quad p, q \in (0, 1),$$

for binary relative entropy. All logarithms are natural. The standard normal distribution function is denoted by Φ , with $\bar{\Phi} := 1 - \Phi$ and $z_u := \Phi^{-1}(u)$ for $u \in (0, 1)$. For $x \in \mathbb{R}$ write $x_+ := \max\{x, 0\}$ and $\mathbf{1}_E$ for the indicator of an event E . In $\mathbb{R}^d$, $x^\top$ denotes transpose, I_d the identity matrix, and $\mathbb{S}_{++}^d$ the cone of symmetric positive-definite $d \times d$ matrices; $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ denote the extreme eigenvalues of a symmetric matrix M . We write $\mathbb{N} := \{1, 2, \dots\}$, $a \wedge b := \min\{a, b\}$, $a \vee b := \max\{a, b\}$, $N(m, v)$ for a Gaussian law with mean m and variance v , and $\text{sgn}(x)$ for the sign of x (with $\text{sgn}(0) = 0$). Conditional expectations, variances, covariances, and correlations are always taken under the probability law explicitly stated in the surrounding paragraph. We use $o(\cdot)$, $O(\cdot)$, and $\sim$ in their standard asymptotic senses, and $\lfloor x \rfloor$ and $\lceil x \rceil$ for the floor and ceiling of x , respectively. For a symmetric matrix M , the notation $M \succ 0$ means that M is positive definite; when $M \succ 0$, $M^{1/2}$ denotes its unique symmetric positive-definite square root and $M^{-1/2}$ its inverse. Symbols introduced within a subsection are local to that subsection unless explicitly carried forward.

2.2. Economic value on the information clock

At a fixed calendar time, suppress the time subscript and write X and θ for the current signal value and drift loading. For the quadratic instantaneous trading objective

$$q\theta X - \frac{\gamma}{2}q^2, \quad \gamma > 0,$$

where $q \in \mathbb{R}$ is the position and γ is the quadratic trading-cost coefficient, the full-information optimal position is $q^* = \theta X / \gamma$. Its value rate is

$$\frac{\theta^2 X^2}{2\gamma} dt = \frac{\sigma^2}{2\gamma} dA_t.$$

Thus statistical information and perfect-information economic value are carried by the same clock. When $H = A$ is deterministic, define the lifetime path information and integrated perfect-information value by

$$\mathcal{I}_{\text{life}} := D_{\text{KL}}(\mathbb{P}_1^A \parallel \mathbb{P}_0^A), \quad V^{PI} := \int_0^T \frac{\theta_t^2 X_t^2}{2\gamma} dt,$$

where the superscript PI stands for perfect information.

Theorem 2.1 (Statistical information and full-information value). *In the deterministic information-horizon benchmark $H = A$,*

$$\mathcal{I}_{\text{life}} = \frac{A}{2}, \quad V^{PI} = \frac{\sigma^2}{2\gamma} A = \frac{\sigma^2}{\gamma} \mathcal{I}_{\text{life}}.$$

More generally, for any bounded $\mathbb{G}$ -stopping time S satisfying $S \leq H$ almost surely,

$$D_{\text{KL}}(\mathbb{P}_1^S \parallel \mathbb{P}_0^S) = \mathbb{E}_1[Z_S] = \frac{1}{2} \mathbb{E}_1[S].$$

Proof. For every bounded stopping time S , the stopped likelihood-ratio identity gives

$$\frac{d\mathbb{P}_1}{d\mathbb{P}_0} \Big|_{\mathcal{G}_S} = e^{Z_S}.$$

Therefore

$$D_{\text{KL}}(\mathbb{P}_1^S \| \mathbb{P}_0^S) = \mathbb{E}_1[Z_S].$$

Under $\mathbb{P}_1$, $Z_s = \tilde{B}_s^{(1)} + s/2$. Boundedness of S makes optional sampling applicable to $\tilde{B}^{(1)}$, so $\mathbb{E}_1[\tilde{B}_S^{(1)}] = 0$ and hence $\mathbb{E}_1[Z_S] = \mathbb{E}_1[S]/2$. For deterministic $S = A$, this gives $\mathcal{I}_{\text{life}} = A/2$. The economic identity follows from the full-information quadratic value calculation immediately above. $\square$

Under this reduction, signal amplitude, feature units, calendar half-life and lookback enter the certification problem through the total information generated before the economic endpoint.

3. Reliable certification before decay

Fix error parameters $\alpha, \beta \in (0, 1)$, where α is the maximum false-deployment probability under $\mathbb{P}_0$ and $1 - \beta$ is the required deployment probability under $\mathbb{P}_1$. Unless stated otherwise we consider the nondegenerate regime $1 - \beta \geq \alpha$.

With only a type-I constraint, a procedure can deploy immediately with probability α without learning anything about the signal. A power target rules out this lottery and makes certification depend on evidence accumulated under the alternative.

For deterministic information lifetime A , a certification/deployment rule is a pair (S, δ) , where $S \in \mathcal{T}_A$ and $\delta \in \{0, 1\}$ is $\mathcal{G}_S$ -measurable. The event $\{\delta = 1\}$ means that the signal is certified and deployed. On $\{\delta = 0\}$ we may, without loss, set $S = A$. Define

$$c_{\alpha, \beta}(A) := \sup \mathbb{E}_1[(A - S)\delta]$$

over all such rules satisfying

$$\mathbb{P}_0(\delta = 1) \leq \alpha, \quad \mathbb{P}_1(\delta = 1) \geq 1 - \beta.$$

The corresponding expected post-certification value, measured in economic units, is

$$\text{CAC}_{\alpha, \beta}(A) = \frac{\sigma^2}{2\gamma} c_{\alpha, \beta}(A).$$

3.1. Exact Gaussian certification threshold

Theorem 3.1 (Exact Gaussian certification threshold). *In the canonical Brownian likelihood experiment, the largest possible probability of deployment by information time A among all level- α rules is*

$$\pi_\alpha(A) = \Phi\left(\sqrt{A} - z_{1-\alpha}\right).$$

Consequently a false-deployment target α and power target $1 - \beta$ are statistically feasible by death if and only if

$$A \geq A_{\text{crit}}(\alpha, \beta) = (z_{1-\alpha} + z_{1-\beta})^2.$$

Equivalently, in KL units,

$$\mathcal{I}_{\text{life}} \geq \mathcal{I}_{\text{crit}}(\alpha, \beta) := \frac{1}{2}(z_{1-\alpha} + z_{1-\beta})^2.$$

Positive post-certification value requires $A > A_{\text{crit}}$: at the critical point the full lifetime is needed merely to attain the requested statistical discrimination.

Proof. Every decision made by a rule with $S \leq A$ is $\mathcal{G}_A$ -measurable, so the Neyman–Pearson lemma applies to the full experiment observed up to A . The full path likelihood ratio at time A is e^{Z_A} , and Z_A is sufficient for testing the two simple path laws. Under $\mathbb{P}_0$ and $\mathbb{P}_1$,

$$Z_A \sim N(-A/2, A), \quad Z_A \sim N(A/2, A),$$

respectively. Neyman–Pearson therefore gives the level- α threshold

$$c_A := -\frac{A}{2} + \sqrt{A} z_{1-\alpha}, \quad Z_A \geq c_A,$$

whose $\mathbb{P}_1$ -probability equals $\Phi(\sqrt{A} - z_{1-\alpha})$. Solving $\pi_\alpha(A) \geq 1 - \beta$ gives the displayed threshold. If $A > A_{\text{crit}}$, choose a deterministic test time $t \in (A_{\text{crit}}, A)$; the level- α Neyman–Pearson test at t has power greater than $1 - \beta$ and leaves positive time $A - t$ whenever it deploys. At equality, feasibility forces the rule to attain the extremal Neyman–Pearson receiver-operating-characteristic (ROC) point. Because the terminal likelihood ratio has a continuous distribution, the level- α most-powerful test is unique up to null sets; therefore $\delta = \mathbf{1}_{E_A}$ almost surely, where $E_A := \{Z_A \geq c_A\}$. If $\mathbb{P}_1(S < A) > 0$, then on $\{S < A\}$ the conditional increment $Z_A - Z_S$ is Gaussian with strictly positive variance $A - S$, so $0 < \mathbb{P}_1(E_A \mid \mathcal{G}_S) < 1$. Since δ is $\mathcal{G}_S$ -measurable and $\delta = \mathbf{1}_{E_A}$, this would imply $\mathbb{P}_1(E_A \mid \mathcal{G}_S) = \delta \in \{0, 1\}$, a contradiction. Thus $S = A$ almost surely at the critical point and the post-certification value is zero. $\square$

For example, with $\alpha = 0.05$ and power 0.90,

$$A_{\text{crit}} = 8.5638473507\dots, \quad \mathcal{I}_{\text{crit}} = 4.2819236753\dots \text{ nats}.$$

If $\mathcal{I}$ denotes the KL information available to an arbitrary binary decision experiment, the generic data-processing bound is $\mathcal{I} \geq \text{kl}(0.9, 0.05) = 2.3762054\dots$ nats, which is below the canonical Gaussian requirement.

A constructive fixed-time lower bound for $A > A_{\text{crit}}$ is

$$c_{\alpha,\beta}(A) \geq \sup_{A_{\text{crit}} \leq t < A} (A - t) \Phi(\sqrt{t} - z_{1-\alpha}).$$

It tests once at time t . The optimal stopping rule can only improve on this benchmark.

3.2. A general information lower bound

Proposition 3.2 (Information lower bound for reliable certification). *Suppose $A > 0$ and $1 - \beta \geq \alpha$. Every feasible rule satisfies*

$$\frac{A - c_{\alpha,\beta}(A)}{2} \geq \text{kl}(1 - \beta, \alpha).$$

Hence

$$\frac{\text{CAC}_{\alpha,\beta}(A)}{V^{PI}} \leq 1 - \frac{\text{kl}(1 - \beta, \alpha)}{\mathcal{I}_{\text{life}}}.$$

Feasibility requires the right-hand side to be nonnegative; otherwise no rule satisfies the declared power and type-I targets.

Proof. Take any feasible rule and write $R = \mathbb{E}_1[(A - S)\delta]$, $p = \mathbb{P}_1(\delta = 1)$, and $q = \mathbb{P}_0(\delta = 1)$. Setting $S = A$ on nondeployment gives

$$R = A - \mathbb{E}_1 S.$$

The stopped likelihood ratio is e^{Zs} , so data processing from the stopped path to the binary decision δ yields

$$\frac{1}{2}\mathbb{E}_1 S = D_{\text{KL}}(\mathbb{P}_1^S \|\mathbb{P}_0^S) \geq \text{kl}(p, q) \geq \text{kl}(1 - \beta, \alpha),$$

where the last inequality uses $p \geq 1 - \beta \geq \alpha \geq q$ and the monotonicity of binary relative entropy on this region. Hence every feasible rule satisfies $R \leq A - 2 \text{kl}(1 - \beta, \alpha)$. Taking the supremum over feasible rules proves the result. $\square$

The exact Gaussian threshold gives the feasibility boundary in the canonical experiment. The KL inequality extends the information-budget argument beyond that model and also bounds the economic value consumed by certification.

4. Economic boundary and the value of information

The economic conversion below is related to value-of-information and decision-oriented experimental-design methods, which evaluate information through the decision it changes [28–31]. Let the economically relevant null boundary be a predictable process $b = (b_t)_{t \geq 0}$ with $b_t < \theta_t$, and assume the same local integrability and Girsanov conditions as in Section 2.1 for the drift gap $\theta - b$. Write $\mathbb{P}_\theta$ and $\mathbb{P}_b$ for the laws on $(\Omega, \mathcal{F}, \mathbb{F})$ of the same observed channel with drift loadings θ and b , respectively, and common noise scale σ , and write $\mathbb{E}_\theta$ for expectation under $\mathbb{P}_\theta$. For $t \geq 0$, let $\mathbb{P}_\theta^t$ and $\mathbb{P}_b^t$ denote their restrictions to $\mathcal{F}_t$. Their relative information clock is

$$A_t^{\theta, b} := \frac{1}{\sigma^2} \int_0^t (\theta_s - b_s)^2 X_s^2 ds.$$

On every bounded horizon on which the corresponding Girsanov condition holds,

$$D_{\text{KL}}(\mathbb{P}_\theta^t \|\mathbb{P}_b^t) = \frac{1}{2} \mathbb{E}_\theta[A_t^{\theta, b}],$$

and when the clock is deterministic this reduces to $A_t^{\theta, b}/2$. The likelihood experiment is therefore canonical after the same information-time change.

If the trader remains completely flat before certification, the perfect-information value forgone is

$$dC_t^{\text{flat}} = \frac{\theta_t^2 X_t^2}{2\gamma} dt = \frac{\sigma^2}{2\gamma} \left(\frac{\theta_t}{\theta_t - b_t} \right)^2 dA_t^{\theta, b}.$$

Since one KL nat equals $dA/2$, the corresponding economic price per nat is

$$\Pi_t^{\text{flat}} = \frac{\sigma^2}{\gamma} \left(\frac{\theta_t}{\theta_t - b_t} \right)^2,$$

which diverges as $b_t \uparrow \theta_t$.

The divergence comes from requiring zero exposure during learning. The statistical experiment remains regular as the drift approaches the economic boundary.

For the next one-dimensional statement, fix a calendar time and suppress its subscript: $X \in \mathbb{R}$ is the current signal value, $\theta \in \mathbb{R}$ is the true loading, $b \in \mathbb{R}$ is the economic-boundary loading, and $q \in \mathbb{R}$ is the position. Write $\Pi_t^{\text{flat}} := \Pi_t^{\text{flat}}$ at this fixed time. The infinitesimal relative information clock is $dA^{\theta, b} := (\theta - b)^2 X^2 dt / \sigma^2$.

Theorem 4.1 (Value relative to the economic boundary). *Suppose the economic comparison is made relative to the position that is optimal at the economic boundary*

$$q_b = \frac{bX}{\gamma}, \quad q_\theta = \frac{\theta X}{\gamma}.$$

The regret of using q_b under drift θ is

$$\frac{(\theta - b)^2 X^2}{2\gamma} dt = \frac{\sigma^2}{2\gamma} dA^{\theta, b}.$$

Hence the incremental economic price per KL nat is exactly

$$\Pi^{\text{relative}} = \frac{\sigma^2}{\gamma},$$

independent of $\theta - b$. The divergence of Π^{flat} therefore measures the cost of insisting on zero exposure while the statistical separation from the boundary model shrinks.

Proof. The instantaneous objective under the true drift is $u(q; \theta) = q\theta X - \gamma q^2/2$. Completing the square gives

$$u(q_\theta; \theta) - u(q_b; \theta) = \frac{(\theta - b)^2 X^2}{2\gamma}.$$

The KL rate between the two drift models is $(\theta - b)^2 X^2/(2\sigma^2)$. Divide. $\square$

The economic value consumed per nat depends on the admissible pre-certification policy. Under zero exposure before certification, the relevant cost is Π^{flat} . Relative to the boundary-optimal policy, the incremental value per nat is the finite constant σ^2/γ . The analysis below uses this boundary-relative conversion.

4.1. Several assets: value per unit of information

In multiple dimensions the scalar factor σ^2/γ is replaced by a direction-dependent economic value per nat. Let $Y_t \in \mathbb{R}^d$ satisfy

$$dY_t = \mu dt + \Sigma^{1/2} dB_t, \quad \mu \in \mathbb{R}^d, \quad \Sigma \in \mathbb{S}_{++}^d,$$

where B is a standard d -dimensional Brownian motion. For a position $q \in \mathbb{R}^d$, let the instantaneous quadratic trading value be

$$u(q; \mu) = q^\top \mu - \frac{1}{2} q^\top \Gamma q, \quad \Gamma \in \mathbb{S}_{++}^d.$$

For the true drift μ and economic-boundary drift $b \in \mathbb{R}^d$, let $\Delta := \mu - b$, $q_\mu := \Gamma^{-1}\mu$, and $q_b := \Gamma^{-1}b$. Write $\mathbb{P}_\mu$ and $\mathbb{P}_b$ for the corresponding path laws on any fixed finite horizon; they share the covariance matrix Σ .

Theorem 4.2 (Multivariate value–information bound). *Using the boundary-optimal position q_b when the drift is μ has instantaneous regret relative to the full-information optimum*

$$dR = \frac{1}{2} \Delta^\top \Gamma^{-1} \Delta dt,$$

whereas the path-law KL divergence between the two Gaussian drift models accrues at rate

$$d\mathcal{I} = \frac{1}{2} \Delta^\top \Sigma^{-1} \Delta dt.$$

Set

$$\Psi := \Sigma^{1/2} \Gamma^{-1} \Sigma^{1/2}.$$

Hence, for $\Delta \neq 0$,

$$\frac{dR}{d\mathcal{I}} = \frac{\Delta^\top \Gamma^{-1} \Delta}{\Delta^\top \Sigma^{-1} \Delta} \in [\lambda_{\min}(\Psi), \lambda_{\max}(\Psi)].$$

Writing I_d for the $d \times d$ identity matrix, the price per information nat is direction-independent if and only if $\Psi = cI_d$ for some $c > 0$, equivalently $\Gamma = c^{-1}\Sigma$, in which case $dR/d\mathcal{I} = c$. The one-dimensional identity is recovered with $c = \sigma^2/\gamma$.

Proof. Completing the square gives $u(q_\mu; \mu) - u(q_b; \mu) = \frac{1}{2}\Delta^\top \Gamma^{-1} \Delta$. Girsanov's formula for Gaussian drift changes gives the stated KL rate. With $x = \Sigma^{-1/2}\Delta$, the ratio is the Rayleigh quotient $x^\top \Psi x / (x^\top x)$, which is bounded by the extreme eigenvalues of Ψ . Equality of the quotient in every direction is equivalent to $\Psi = cI_d$. $\square$

These eigenvalue bounds are the multivariate version of the CAC conversion factor: statistical directions are valuable according to both their distinguishability under Σ^{-1} and their economic curvature under Γ^{-1} . The spectrum therefore identifies which market directions are cheap or expensive to certify economically.

5. Alpha Survival Frontiers

Alpha decay has been studied both as a post-discovery empirical phenomenon and as an input to trading and strategy-durability models [32–35, 40]. The object here is the amount of statistical information that remains before the signal reaches its economic boundary. For the local terminal asymptotics in this section, take the drift and covariance paths to be deterministic; the same formulas hold pathwise on any realization for which the stated expansions are valid. Let the observed d -dimensional return channel, under the true market model and an economic-boundary model, differ by drift gap

$$\Delta_t = \mu_t - b_t \in \mathbb{R}^d,$$

with instantaneous noise covariance $\Sigma_t \in \mathbb{S}_{++}^d$. Throughout this section, a superscript (b) on an information quantity means that KL divergence is measured relative to the economic-boundary model $b = (b_t)$. Let T be the end of the opportunity's economic lifetime. Assume that as $t \uparrow T$, for some $p > 0$ and nonzero $c \in \mathbb{R}^d$,

$$\Delta_t = c(T - t)^p + o((T - t)^p), \quad \Sigma_t \rightarrow \Sigma_T \succ 0.$$

Define remaining KL information from time t until death by

$$\mathcal{I}_{\text{rem}}^{(b)}(t) := \frac{1}{2} \int_t^T \Delta_s^\top \Sigma_s^{-1} \Delta_s ds.$$

Theorem 5.1 (Remaining information near the economic endpoint). *Under the assumptions above,*

$$\mathcal{I}_{\text{rem}}^{(b)}(t) = \frac{c^\top \Sigma_T^{-1} c}{2(2p+1)} (T - t)^{2p+1} + o((T - t)^{2p+1}).$$

In particular, for a transverse crossing $p = 1$,

$$\mathcal{I}_{\text{rem}}^{(b)}(t) \sim \frac{c^\top \Sigma_T^{-1} c}{6} (T - t)^3.$$

Proof. Put $u = T - s$. By the assumed expansion and continuity of matrix inversion on the positive-definite cone,

$$\Delta_{T-u}^\top \Sigma_{T-u}^{-1} \Delta_{T-u} = u^{2p} c^\top \Sigma_T^{-1} c + o(u^{2p}).$$

Integrating from $u = 0$ to $u = T - t$ gives

$$\mathcal{I}_{\text{rem}}^{(b)}(t) = \frac{1}{2} c^\top \Sigma_T^{-1} c \frac{(T-t)^{2p+1}}{2p+1} + o((T-t)^{2p+1}).$$

□

5.1. Economic value has the same terminal scaling

Let the instantaneous trading objective be

$$q^\top \mu_t - \frac{1}{2} q^\top \Gamma_t q, \quad \Gamma_t \succ 0,$$

and compare the position $q_t^\mu = \Gamma_t^{-1} \mu_t$, optimal under μ_t , with the boundary-optimal position $q_t^b = \Gamma_t^{-1} b_t$. The resulting instantaneous value gap is

$$\frac{1}{2} \Delta_t^\top \Gamma_t^{-1} \Delta_t.$$

Define the remaining boundary-relative value by

$$V_{\text{rem}}^{\mu:b}(t) := \frac{1}{2} \int_t^T \Delta_s^\top \Gamma_s^{-1} \Delta_s ds.$$

If $\Gamma_t \rightarrow \Gamma_T \succ 0$, then

$$V_{\text{rem}}^{\mu:b}(t) = \frac{c^\top \Gamma_T^{-1} c}{2(2p+1)} (T-t)^{2p+1} + o((T-t)^{2p+1}).$$

Consequently,

$$\frac{V_{\text{rem}}^{\mu:b}(t)}{\mathcal{I}_{\text{rem}}^{(b)}(t)} \rightarrow \frac{c^\top \Gamma_T^{-1} c}{c^\top \Sigma_T^{-1} c}.$$

The limit is the local terminal counterpart of the economic-information spectrum above.

5.2. Latest feasible start time near the economic endpoint

If on a terminal interval the gap is exactly

$$\Delta_{T-u} = cu^p$$

and Σ is constant, then, for remaining calendar time $\ell \geq 0$, the preceding formula is exact:

$$\mathcal{I}_{\text{rem}}^{(b)}(T-\ell) = K_I \ell^{2p+1}, \quad K_I := \frac{c^\top \Sigma^{-1} c}{2(2p+1)}.$$

For reliable certification target $(\alpha, 1-\beta)$, define

$$\mathcal{I}_* := \mathcal{I}_{\text{crit}}(\alpha, \beta) = \frac{1}{2} [z_{1-\alpha} + z_{1-\beta}]^2.$$

For a researcher who first begins observing the candidate at time $T-\ell$, the terminal interval in which the entire future path contains insufficient information has exact width

$$\ell_* = \left(\frac{2(2p+1)\mathcal{I}_*}{c^\top \Sigma^{-1} c} \right)^{1/(2p+1)}.$$

For a transverse crossing,

$$\ell_* = \left(\frac{6\mathcal{I}_*}{c^\top \Sigma^{-1} c} \right)^{1/3}.$$

Thus there is generally a nonempty terminal region in which the alpha remains economically positive while a newly started certification experiment cannot reach the declared reliability target before the economic endpoint. This boundary is the latest time at which a new certification experiment can still succeed. A procedure that began earlier carries its accumulated likelihood state and remaining error budget, so its continuation problem is state-dependent.

For a general smooth crossing, the same expression gives the local asymptotic scale of this latest-start boundary; the exact global boundary depends on the full path geometry.

5.3. Exponential decay with a positive economic hurdle

Take the scalar model

$$\mu_t = \mu_0 e^{-\lambda t}, \quad 0 < b < \mu_0,$$

and define the economic endpoint by $\mu_T = b$. Then

$$T = \frac{1}{\lambda} \log \frac{\mu_0}{b}.$$

Compare the true drift with the boundary drift b . All KL quantities in this subsection are relative to that boundary model. The total lifetime KL is

$$\mathcal{I}_{\text{life}}^{(b)} = \frac{1}{2\sigma^2} \int_0^T (\mu_0 e^{-\lambda t} - b)^2 dt.$$

Writing $r_b := b/\mu_0 \in (0, 1)$, direct integration gives the exact identity

$$\mathcal{I}_{\text{life}}^{(b)} = \frac{\mu_0^2}{2\sigma^2 \lambda} \left[\frac{1}{2} - 2r_b + \frac{3}{2}r_b^2 + r_b^2 \log \frac{1}{r_b} \right].$$

5.3.1. Cubic scaling near the economic hurdle

Let $r_b = 1 - \varepsilon$, $\varepsilon \downarrow 0$. Since

$$\frac{1}{2} - 2r_b + \frac{3}{2}r_b^2 + r_b^2 \log(1/r_b) = \frac{\varepsilon^3}{3} + \frac{\varepsilon^4}{12} + O(\varepsilon^5),$$

we obtain

$$\mathcal{I}_{\text{life}}^{(b)} \sim \frac{\mu_0^2}{6\sigma^2 \lambda} \left(1 - \frac{b}{\mu_0} \right)^3.$$

The same cubic law therefore appears in the exponential model near the economic hurdle.

A more local expression follows from $\mu_{T-u} = be^{\lambda u}$:

$$\mathcal{I}_{\text{rem}}^{(b)}(T - \ell) = \frac{b^2}{2\sigma^2} \left[\frac{e^{2\lambda\ell} - 1}{2\lambda} - \frac{2(e^{\lambda\ell} - 1)}{\lambda} + \ell \right] \sim \frac{b^2 \lambda^2}{6\sigma^2} \ell^3.$$

Thus the statistical evidence still available before the economic endpoint can disappear much faster than the remaining calendar time itself.

5.4. Zero-hurdle exponential model: the Alpha Survival Frontier

If the economic boundary is zero and $\mu_t = \mu_0 e^{-\lambda t}$ for all $t \geq 0$, the calendar horizon is unbounded while total lifetime information remains finite:

$$\mathcal{I}_{\text{life}} = \frac{1}{2\sigma^2} \int_0^\infty \mu_0^2 e^{-2\lambda t} dt = \frac{\mu_0^2}{4\sigma^2\lambda}.$$

Let

$$S := \frac{\mu_0}{\sigma}, \quad h_{1/2} := \frac{\log 2}{\lambda}$$

be the instantaneous Sharpe (with time units fixed consistently) and alpha half-life. Then

$$\mathcal{I}_{\text{life}} = \frac{S^2 h_{1/2}}{4 \log 2}.$$

Reliable certification is feasible only if

$$\frac{S^2 h_{1/2}}{4 \log 2} \geq \frac{1}{2} [z_{1-\alpha} + z_{1-\beta}]^2,$$

or equivalently

$$h_{1/2} \geq h_{\min}(S; \alpha, \beta) := \frac{2 \log 2}{S^2} [z_{1-\alpha} + z_{1-\beta}]^2.$$

We refer to this boundary as the canonical *Alpha Survival Frontier*.

For $\alpha = 0.05$, power $1 - \beta = 0.90$:

instantaneous Sharpe S	minimum half-life
1.0	11.872
1.5	5.276
2.0	2.968
3.0	1.319

These thresholds use the calendar-time unit implicit in the scaling of S and are implied by the stated exponential-decay model.

5.5. Multiple testing and search breadth

Large-scale strategy search creates the familiar multiplicity and selection problem in empirical finance [9–11]. Suppose $M \in \mathbb{N}$ candidate signals are screened under familywise false-deployment budget $\alpha_F \in (0, 1)$. Under Bonferroni, assign

$$\alpha_M = \frac{\alpha_F}{M}.$$

The canonical information threshold for one selected signal is then

$$\mathcal{I}_{\text{crit}}^{(M)} = \frac{1}{2} [z_{1-\alpha_F/M} + z_{1-\beta}]^2.$$

Under a Bonferroni allocation, this is the exact per-test Gaussian threshold. Bonferroni remains valid under arbitrary dependence; additional dependence information can support less conservative allocations.

The zero-hurdle exponential survival frontier becomes

$$h_{\min}^{(M)}(S) = \frac{2 \log 2}{S^2} [z_{1-\alpha_F/M} + z_{1-\beta}]^2.$$

At $\alpha_F = 0.05$, power 0.90:

M	$h_{\min}, S = 1$	$S = 1.5$	$S = 2$	$S = 3$
1	11.872	5.276	2.968	1.319
10	20.627	9.168	5.157	2.292
100	28.979	12.880	7.245	3.220
1000	37.085	16.482	9.271	4.121

5.5.1. Maximum search breadth under Bonferroni correction

Fix lifetime information $\mathcal{I} \geq 0$. For power $1 - \beta$, the smallest level at which the canonical experiment can achieve the requested power is

$$\alpha_{\min}(\mathcal{I}, \beta) = \bar{\Phi}(\sqrt{2\mathcal{I}} - z_{1-\beta}).$$

Therefore Bonferroni can support at most

$$M_{\max}^{\text{Bonf}}(\mathcal{I}) = \left\lfloor \frac{\alpha_F}{\bar{\Phi}(\sqrt{2\mathcal{I}} - z_{1-\beta})} \right\rfloor,$$

with zero feasible candidates if the right-hand side is below one.

Under independent candidate tests, Sidak replaces α_F/M by $1 - (1 - \alpha_F)^{1/M}$, yielding the exact inversion

$$M_{\max}^{\text{Sidak}}(\mathcal{I}) = \left\lfloor \frac{\log(1 - \alpha_F)}{\log(1 - \alpha_{\min}(\mathcal{I}, \beta))} \right\rfloor.$$

The Sidak formula requires independence. Bonferroni supplies the dependence-robust guarantee used elsewhere in the paper.

5.5.2. Unequal allocation of the familywise error budget

Different lifetime information budgets generally call for unequal Bonferroni allocations. Candidate j , with lifetime information $\mathcal{I}_j \geq 0$ and power target $1 - \beta_j$, with $\beta_j \in (0, 1)$, needs marginal type-I level at least

$$\alpha_j^* = \bar{\Phi}(\sqrt{2\mathcal{I}_j} - z_{1-\beta_j}).$$

This expression exactly inverts the one-candidate Gaussian power envelope.

Theorem 5.2 (Heterogeneous Bonferroni feasibility). *Consider separable candidate tests whose marginal type-I budgets α_j are combined by the dependence-robust union bound $\sum_j \alpha_j \leq \alpha_F$. Then all candidates can meet their declared power targets if and only if*

$$\sum_{j=1}^M \alpha_j^* \leq \alpha_F.$$

Moreover, among a fixed list of candidates with unit value per certified candidate, the maximum feasible cardinality is obtained by sorting α_j^ increasingly and taking the longest prefix whose cumulative sum is at most α_F .*

Proof. Necessity within the separable Bonferroni architecture follows because candidate j cannot attain power $1 - \beta_j$ at any level below α_j^* ; hence every feasible allocation obeys $\alpha_j \geq \alpha_j^*$, and summing gives the inequality. Sufficiency follows by setting $\alpha_j = \alpha_j^*$ and applying the union bound. For maximum cardinality, any feasible set of size k has total cost at least the sum of the k smallest α_j^* ; therefore the cheapest prefix is feasible whenever any k -set is feasible. $\square$




Figure 1: Canonical Alpha Survival Frontier for familywise false-deployment level 0.05 and power 0.90. Each curve gives the minimum half-life required at a given instantaneous Sharpe for a Bonferroni search over M candidates. The highlighted point is $S = 2$, $M = 100$, for which the minimum half-life is 7.245 years.

This converts candidate research into an explicit error-budget allocation problem. If candidates carry heterogeneous economic weights, choosing the subset with maximum total weight subject to $\sum_j \alpha_j^* \leq \alpha_F$ becomes a 0–1 knapsack problem. Finite alpha lifetime therefore turns familywise error control into a scarce allocation budget.

5.5.3. Large-search asymptotics

Let $\ell_M := \log(M/\alpha_F)$ and hold β fixed. Using the standard Gaussian upper-quantile expansion,

$$z_{1-\alpha_F/M} = \sqrt{2\ell_M} - \frac{\log \ell_M + \log(4\pi)}{2\sqrt{2\ell_M}} + o(\ell_M^{-1/2}).$$

Hence

$$\mathcal{I}_{\text{crit}}^{(M)} = \ell_M + z_{1-\beta} \sqrt{2\ell_M} - \frac{1}{2} \log \ell_M - \frac{1}{2} \log(4\pi) + \frac{1}{2} z_{1-\beta}^2 + o(1).$$

In particular, the leading information cost of search breadth is logarithmic in M :

$$\mathcal{I}_{\text{crit}}^{(M)} \sim \log M.$$

The Gaussian tail therefore makes the leading search cost logarithmic in M . CAC compares this growing requirement with the finite lifetime-information budget generated by the opportunity.

5.6. A dimensionless feasibility ratio

Define, for a declared correction rule,

$$\text{CR} := \frac{\mathcal{I}_{\text{life}}}{\mathcal{I}_{\text{crit}}(\alpha, \beta)}.$$

With multiplicity, replace α by the declared per-test budget.

Interpretation in the canonical experiment:

- $\text{CR} < 1$: reliable deployment at the requested error/power pair is statistically infeasible before the opportunity exhausts its information budget;
- $\text{CR} = 1$: terminal feasibility boundary; the entire information lifetime can be consumed by discrimination;
- $\text{CR} > 1$: reliable certification is statistically feasible under the declared model; realized profitability is a separate economic question.

The ratio is specific to the statistical experiment, economic boundary, error targets, decay model, and multiplicity rule used in its numerator and denominator.

6. Market equilibrium with certification costs

Information acquisition and arbitrage are endogenous in several market-equilibrium models [18, 20–22]. More recent work studies competition among factor investors and the empirical role of post-publication attenuation, arbitrage capital, and implementation frictions [23–27]. The preceding sections take opportunity lifetime as given. We now allow arbitrage activity to shorten it. An explicit crowding map links aggregate research intensity to the remaining information budget, so the CAC threshold becomes an equilibrium constraint.

Retain the reliable-information threshold

$$\mathcal{I}_* := \mathcal{I}_{\text{crit}}(\alpha, \beta) = \frac{1}{2} [z_{1-\alpha} + z_{1-\beta}]^2.$$

Let $v > 0$ denote economic value per KL nat. In the canonical scalar model, Theorem 4.1 gives $v = \sigma^2/\gamma$. Define the optimal reliable CAC economic surplus at lifetime information $\mathcal{I} \geq 0$ by

$$C^*(\mathcal{I}) := \frac{v}{2} c_{\alpha, \beta}(2\mathcal{I}).$$

The equilibrium arguments below use only the resulting value function and its stated regularity; Appendix B gives the corresponding randomized constrained-stopping representation. For the equilibrium result we require the natural regularity that C^* is continuous, with

$$C^*(\mathcal{I}) = 0 \quad (\mathcal{I} \leq \mathcal{I}_*), \quad C^*(\mathcal{I}) > 0 \quad (\mathcal{I} > \mathcal{I}_*).$$

In the canonical model, the CAC value is strictly increasing above the frontier. Indeed, if one KL nat of perfect-information value is worth v , then embedding any ε -optimal rule at $\mathcal{I}_1$ in the longer horizon $\mathcal{I}_2 > \mathcal{I}_1 \geq \mathcal{I}_*$ adds $v(\mathcal{I}_2 - \mathcal{I}_1)$ on every deployment, whose probability under the alternative law is at least $1 - \beta$. Letting $\varepsilon \downarrow 0$ gives

$$C^*(\mathcal{I}_2) - C^*(\mathcal{I}_1) \geq v(1 - \beta)(\mathcal{I}_2 - \mathcal{I}_1).$$

Thus the constructive linear surplus

$$C_0(\mathcal{I}) = v(1 - \beta)(\mathcal{I} - \mathcal{I}_*)_+$$

is a lower envelope of the optimal CAC value.

Let $\mathcal{I}_0 > 0$ be the no-crowding lifetime-information endowment and let $x \geq 0$ denote aggregate arbitrage/research intensity. The remaining information budget is

$$\mathcal{I} = G(x),$$

where $G : [0, \infty) \rightarrow [0, \mathcal{I}_0]$ is continuous and strictly decreasing, with $G(0) = \mathcal{I}_0$. The canonical exponential-crowding benchmark uses a crowding sensitivity $\chi > 0$ and is

$$G(x) = \frac{\mathcal{I}_0}{1 + \chi x},$$

which arises when $\lambda_0 > 0$ denotes the baseline exponential decay rate and crowding accelerates it from λ_0 to $\lambda_0(1 + \chi x)$.

Suppose each unit of active arbitrage capital pays marginal research or implementation cost $\psi > 0$. In a competitive market with many individually small arbitrageurs, aggregate post-certification value $C^*(G(x))$ is allocated proportionally across active capacity. Positive-entry equilibrium therefore satisfies

$$\frac{C^*(G(x))}{x} = \psi,$$

or equivalently

$$C^*(G(x)) = \psi x.$$

Theorem 6.1 (Competitive equilibrium under certification costs). *Assume C^* and G have the properties above. If $\mathcal{I}_0 > \mathcal{I}_*$ and there is a finite x_* with $G(x_*) = \mathcal{I}_*$, then for every $\psi > 0$ there is a unique positive equilibrium $x_\psi \in (0, x_*)$ satisfying*

$$C^*(G(x_\psi)) = \psi x_\psi.$$

Under the entry dynamic

$$\dot{x} = \eta \{C^*(G(x)) - \psi x\}, \quad \eta > 0,$$

this equilibrium is globally asymptotically stable. Its residual information $\mathcal{I}_\psi = G(x_\psi)$ satisfies

$$\mathcal{I}_* < \mathcal{I}_\psi < \mathcal{I}_0, \quad \mathcal{I}_\psi \downarrow \mathcal{I}_* \quad \text{as } \psi \downarrow 0.$$

If $\mathcal{I}_0 \leq \mathcal{I}_$, the no-entry equilibrium $x = 0$ is globally stable.*

Proof. For $\mathcal{I}_0 > \mathcal{I}_*$ define

$$F_\psi(x) := C^*(G(x)) - \psi x, \quad x \in [0, x_*].$$

Because C^* is continuous and strictly increasing above $\mathcal{I}_*$ while G is continuous and strictly decreasing, F_ψ is continuous and strictly decreasing. Moreover,

$$F_\psi(0) = C^*(\mathcal{I}_0) > 0, \quad F_\psi(x_*) = -\psi x_* < 0.$$

The intermediate value theorem therefore gives a unique root $x_\psi \in (0, x_*)$. The sign of F_ψ is positive below x_ψ and negative above it; for $x \geq x_*$, $G(x) \leq \mathcal{I}_*$ and hence $F_\psi(x) = -\psi x < 0$. Because F_ψ is continuous and decreasing, the scalar vector field is one-sided Lipschitz: $(x - y)(F_\psi(x) - F_\psi(y)) \leq 0$. Solutions are therefore unique. The vector field points toward x_ψ from every initial condition in $[0, \infty)$ and is strictly negative for sufficiently large x , which proves global asymptotic stability on the economically relevant state space.

If $0 < \psi_1 < \psi_2$, then at the ψ_2 root, $F_{\psi_1}(x_{\psi_2}) = (\psi_2 - \psi_1)x_{\psi_2} > 0$, so $x_{\psi_1} > x_{\psi_2}$. Hence x_ψ increases as $\psi \downarrow 0$ and has a limit $\bar{x} \leq x_*$. If $\bar{x} < x_*$, continuity gives $C^*(G(\bar{x})) > 0$, whereas the equilibrium identity $C^*(G(x_\psi)) = \psi x_\psi$ tends to zero, a contradiction. Therefore $x_\psi \uparrow x_*$ and $\mathcal{I}_\psi = G(x_\psi) \downarrow \mathcal{I}_*$. If $\mathcal{I}_0 \leq \mathcal{I}_*$, then $G(x) \leq \mathcal{I}_*$ for all $x \geq 0$, so the dynamic reduces to $\dot{x} = -\eta \psi x$ and $x = 0$ is globally stable. $\square$

The equilibrium divides opportunities at the certification frontier. Below the frontier, reliable post-certification value is zero. Above the frontier, rents finance entry and crowding reduces the remaining information budget. As research cost falls, the equilibrium residual information approaches the reliable-information threshold.

The lower envelope C_0 also yields a useful nonasymptotic bound. Since G is strictly decreasing, write G^{-1} for its inverse on its range. The optimal equilibrium satisfies

$$\mathcal{I}_\psi - \mathcal{I}_* \leq \frac{\psi G^{-1}(\mathcal{I}_*)}{v(1 - \beta)}.$$

For $G(x) = \mathcal{I}_0/(1 + \chi x)$,

$$\mathcal{I}_\psi - \mathcal{I}_* \leq \frac{\psi}{\chi v(1 - \beta)} \left(\frac{\mathcal{I}_0}{\mathcal{I}_*} - 1 \right).$$

Hence the residual excess information above the CAC frontier vanishes at least linearly with research cost, without requiring differentiability of the stopping value at the threshold.

6.1. A finite number of competing arbitrageurs

The proportional allocation in the large-number competitive model can itself be obtained as the limit of a strategic game with finitely many arbitrageurs. Let $N \geq 2$ arbitrage desks choose efforts $e_i \geq 0$, with aggregate $E := \sum_{i=1}^N e_i$; write e_{-i} for the vector of efforts of desks other than i . Let

$$R(E) := C^*(G(E)).$$

For $E > 0$, desk i receives proportional share e_i/E of the endogenous post-certification value pool and pays ψe_i :

$$u_i(e_i, e_{-i}) = \frac{e_i}{E} R(E) - \psi e_i.$$

At a symmetric interior equilibrium, write $E_N > 0$ for aggregate equilibrium effort. Differentiability of R at E_N gives

$$\psi = \left(1 - \frac{1}{N}\right) \frac{R(E_N)}{E_N} + \frac{1}{N} R'(E_N).$$

The second term is a strategic prize-destruction effect. Since $R'(E_N) < 0$, a finite arbitrageur internalizes that additional effort reduces the total opportunity available to all participants. Under the usual single-crossing condition for the symmetric first-order equation, E_N increases with N and converges to the large-number competitive solution $R(E)/E = \psi$.

A closed-form benchmark is obtained by combining, on $0 \leq E \leq \mathcal{I}_0/\kappa$, the linear information impact

$$G(E) = \mathcal{I}_0 - \kappa E, \quad \kappa > 0,$$

with the constructive post-certification value

$$C_0(\mathcal{I}) = a(\mathcal{I} - \mathcal{I}_*)_+, \quad a = v(1 - \beta).$$

Proposition 6.2 (Exact finite- N linear benchmark). *Let $N \geq 2$, $\mathcal{I}_0 > \mathcal{I}_*$, $a, \kappa > 0$, and $\psi \geq 0$. Writing $\Delta\mathcal{I} = \mathcal{I}_0 - \mathcal{I}_*$, the proportional-share game induced by the linear information map and C_0 has a unique symmetric positive equilibrium with aggregate effort*

$$E_N = \frac{a\Delta\mathcal{I}(N-1)}{N(a\kappa + \psi)}.$$

Define the equilibrium residual information by $\mathcal{I}_N^{\text{eq}} := \mathcal{I}_0 - \kappa E_N$. The equilibrium remains in the positive-rent region, and its residual information satisfies

$$\mathcal{I}_N^{\text{eq}} - \mathcal{I}_* = \Delta \mathcal{I} \frac{a\kappa + N\psi}{N(a\kappa + \psi)}.$$

Equivalently,

$$\frac{\mathcal{I}_N^{\text{eq}} - \mathcal{I}_*}{\mathcal{I}_0 - \mathcal{I}_*} = \frac{\psi}{a\kappa + \psi} + \frac{a\kappa}{N(a\kappa + \psi)}.$$

In particular, if $\psi = 0$, then

$$\mathcal{I}_N^{\text{eq}} = \mathcal{I}_* + \frac{\mathcal{I}_0 - \mathcal{I}_*}{N}.$$

Proof. In the positive-rent region the payoff is

$$u_i = \frac{e_i}{E} a \Delta \mathcal{I} - (a\kappa + \psi) e_i.$$

Set $\mathcal{V}_0 := a \Delta \mathcal{I}$ and $k := a\kappa + \psi$. Define the aggregate rival effort by $E_{-i} := \sum_{j \neq i} e_j$. Against $E_{-i} > 0$, the first and second derivatives are

$$\frac{\partial u_i}{\partial e_i} = \frac{\mathcal{V}_0 E_{-i}}{(e_i + E_{-i})^2} - k, \quad \frac{\partial^2 u_i}{\partial e_i^2} = -\frac{2\mathcal{V}_0 E_{-i}}{(e_i + E_{-i})^3} < 0.$$

At a symmetric equilibrium $e_i = e$ and $E = Ne$, the first-order condition gives $e = \mathcal{V}_0(N-1)/(N^2k)$ and hence the stated E_N . Moreover,

$$\frac{\kappa E_N}{\Delta \mathcal{I}} = \frac{a\kappa(N-1)}{N(a\kappa + \psi)} < 1,$$

so the solution lies strictly above the CAC frontier. Substituting $\mathcal{I}_N^{\text{eq}} = \mathcal{I}_0 - \kappa E_N$ yields the residual-information identities. A unilateral deviation that pushes aggregate effort beyond the positive-rent region obtains zero gross rent and nonpositive net payoff, whereas the stated interior equilibrium gives a strictly positive payoff. Strict concavity in each player's own effort within the positive-rent region then gives uniqueness of the symmetric best-response fixed point. $\square$

The first term in the residual-information fraction comes from research cost; the second comes from having only finitely many competitors. Thus a finite set of arbitrageurs leaves positive excess information even when research itself is costless; only the large-number competitive limit reaches the CAC frontier. In the symmetric linear benchmark, $1/N$ is also the Herfindahl index of effort shares, so the zero-cost residual-information fraction equals market concentration. With nonlinear crowding or asymmetric payoffs the relationship changes with the market technology. In the zero-hurdle exponential Gaussian model, $\mathcal{I} = S^2 h_{1/2}/(4 \log 2)$. Consequently, for opportunities that are actively compressed by many low-cost arbitrageurs,

$$S^2 h_{1/2} \longrightarrow 4 \log 2 \mathcal{I}_*.$$

This limit describes opportunities whose lifetime is actively compressed by many low-cost arbitrageurs.




Figure 2: Residual excess lifetime information in the exact linear-information strategic benchmark. With zero research cost the residual fraction is exactly $1/N$; positive research cost creates an additional limiting gap even as the number of competitors becomes large.

6.2. Heterogeneous implementation costs and capacity

These comparative statics are conditional on the way arbitrage activity shortens signal lifetime. In a pooled cross-section, their sign depends on how research cost, price impact, and competitive breadth vary with the opportunity's information budget. The exact linear strategic benchmark makes this distinction transparent. For each opportunity i , let $\mathcal{I}_i^0 > 0$ be its no-crowding lifetime information, let $N_i \geq 2$ be the number of competing arbitrageurs, let $a_i > 0$ be its gross value-per-unit coefficient, let $\kappa_i > 0$ be its linear crowding sensitivity, and let $\psi_i \geq 0$ be its marginal implementation/research cost. Define

$$J_i := (\mathcal{I}_i^0 - \mathcal{I}_*)_+$$

as excess lifetime information above the CAC frontier. The amount destroyed by equilibrium arbitrage is

$$D_i = m_i J_i, \quad m_i = \frac{a_i \kappa_i (1 - 1/N_i)}{a_i \kappa_i + \psi_i}.$$

Here $m_i \in [0, 1)$ measures how easily arbitrage activity can compress the opportunity: it rises with the effectiveness of crowding and competitive breadth and falls with research and implementation cost. For the pooled cross-sectional argument below, let (J, M, D) denote a generic population draw of excess information, implementation coefficient, and information destruction, respectively. We write $\mathbb{P}$ for their joint population law and $\mathbb{E}$, Cov , and Var for the corresponding moments.

Proposition 6.3 (Pooled regressions with heterogeneous implementation technology). *Let $J \geq 0$, $M \geq 0$, and suppose $D = MJ + \varepsilon$ with $\mathbb{E}[\varepsilon \mid J, M] = 0$ and finite second moments sufficient for all covariances below; for example, assume $\mathbb{E}[J^2] + \mathbb{E}[(MJ)^2] + \mathbb{E}[\varepsilon^2] < \infty$. Writing $\mu(j) = \mathbb{E}[M \mid J = j]$,*

$$\mathbb{E}[D \mid J = j] = j\mu(j).$$

Hence the conditional relation between excess information and subsequent decay is increasing exactly when $j\mu(j)$ is increasing. If μ is differentiable and positive, then for $j > 0$ the strict local condition is

$$\frac{d \log \mu(j)}{d \log j} > -1.$$

Moreover, when $\text{Var}(J) > 0$, the pooled population ordinary-least-squares (OLS) slope of D on J satisfies

$$\beta_{\text{pool}} = \mathbb{E}[M] + \frac{\text{Cov}(M, J^2) - \mathbb{E}[J] \text{Cov}(M, J)}{\text{Var}(J)}.$$

In particular, independence of M and J gives $\beta_{\text{pool}} = \mathbb{E}[M] \geq 0$, with strict positivity whenever $\mathbb{P}(M > 0) > 0$; sufficiently negative dependence can instead make the pooled slope zero or negative.

Proof. The conditional mean follows from iterated expectations: $\mathbb{E}[D | J] = J\mathbb{E}[M | J]$. For $j > 0$, $d(j\mu(j))/dj = \mu(j) + j\mu'(j)$, which gives the logarithmic condition. For the population OLS slope,

$$\beta_{\text{pool}} = \frac{\text{Cov}(J, D)}{\text{Var}(J)} = \frac{\mathbb{E}[MJ^2] - \mathbb{E}[J]\mathbb{E}[MJ]}{\text{Var}(J)},$$

because $\mathbb{E}[J\varepsilon] = 0$. Writing $\mathbb{E}[MJ^2] = \mathbb{E}[M]\mathbb{E}[J^2] + \text{Cov}(M, J^2)$ and $\mathbb{E}[MJ] = \mathbb{E}[M]\mathbb{E}[J] + \text{Cov}(M, J)$ yields the displayed identity. Independence eliminates both covariance terms. Since $M \geq 0$, the resulting slope is nonnegative, and it is strictly positive whenever $\mathbb{P}(M > 0) > 0$. $\square$

The proposition identifies

$$m_i(\mathcal{I}_i^0 - \mathcal{I}_*)_+ = m_i J_i$$

as the structural predictor of information destroyed by arbitrage. The implementation coefficient matters empirically because turnover, short-sale costs, capacity, and price impact vary substantially across anomaly strategies. Recent evidence that stock-borrow fees can eliminate average anomaly profits provides a direct example of economically large heterogeneity in implementation cost [27].

6.3. Cross-sectional empirical check

The equilibrium mechanism predicts subsequent information destruction through both excess lifetime information and implementation technology. In a pooled cross-section, the relevant state is therefore the product $m_i(\mathcal{I}_i^0 - \mathcal{I}_*)_+ = m_i J_i$.

Using the Open Source Asset Pricing data set [1], let t_i^{op} denote the originally reported in-sample t -statistic for predictor i , and define the pre-publication feasibility proxy

$$\text{CR}_i^{\text{proxy}} := \frac{(t_i^{\text{op}})^2/2}{\mathcal{I}_{\text{crit}}(0.05, 0.10)}.$$

Let $\bar{r}_{i,\text{pre}}$ be the mean long-short return between the end of the original sample and the publication year, $\bar{r}_{i,\text{post}}$ the mean over the first 60 available post-publication months, and $\hat{\sigma}_{i,\text{IS}}$ the original-sample return standard deviation. The normalized attenuation outcome is

$$\Delta_i^{\text{pub}} := \frac{\bar{r}_{i,\text{pre}} - \bar{r}_{i,\text{post}}}{\hat{\sigma}_{i,\text{IS}}}.$$

Requiring at least 24 observations in the original, pre-publication, and post-publication windows leaves 205 eligible predictors. The Spearman association between $\text{CR}_i^{\text{proxy}}$ and Δ_i^{pub} is 0.059, with a 95% bootstrap interval of $[-0.086, 0.199]$ obtained by resampling author $\times$ publication-year clusters. An ordinary least-squares regression of Δ_i^{pub} on $\log(1 + \text{CR}_i^{\text{proxy}})$, using the heteroskedasticity-consistent HC3 covariance estimator, gives slope 0.0058 with standard error 0.0240 ($t = 0.24$). The available pooled benchmark therefore contains little information about the equilibrium comparative static when the implementation coefficient m_i is omitted.

A separate historical BTCUSDT holdout tests the two-stage entry-and-destruction mechanism within a single liquid market. The holdout contains 57 events. Let R_A^{hold} denote the pre-specified

Stage-A event regressor labeled *excess certifiability* $\times$ *log depth* in the released replication record and let R_B^{hold} denote the pre-specified Stage-B regressor labeled *arbitrage-aligned entry*. Let $\hat{\beta}_A$ and $\hat{\beta}_B$ denote the regression coefficients attached to R_A^{hold} and R_B^{hold} , respectively. The raw third-party microstructure inputs are not redistributed, so these two event-level regressors are treated here as the fixed observables produced by the documented holdout construction. The estimates are $\hat{\beta}_A = -1.741 \times 10^{-5}$ with one-sided 97.5% week-block bootstrap lower bound -3.002×10^{-5} , and $\hat{\beta}_B = 2817.99$ with one-sided 97.5% week-block bootstrap lower bound -9857.24 . The joint criterion requires both lower bounds to be positive (equivalently, both stages to clear their pre-specified one-sided 2.5% thresholds) and is not met in this sample. Together with Proposition 6.3, these benchmarks are consistent with implementation technology being a central state variable for empirical identification.

The funding-rate study below studies a separate implication of the framework: whether a pre-outcome estimate of signal lifetime contains information about subsequent persistence beyond current signal magnitude.

7. Funding-rate persistence illustration

Recent work documents predictability, cross-venue structure, and persistence in perpetual-futures funding rates [36–39]. The funding benchmark asks whether a pre-outcome estimate of signal lifetime adds information about future persistence beyond contemporaneous signal magnitude. The exercise is retrospective and targets the persistence channel developed in the survival-frontier analysis.

7.1. Data and design

The data contain BTC-USDT and ETH-USDT perpetual-futures funding histories from Binance, Bybit, Gate, HTX, KuCoin, and MEXC, observed at eight-hour settlements. An event requires a continuous 180-day pre-event history, current absolute funding at or above its trailing 95th percentile, a seven-day cooldown, and a complete seven-day outcome window. The resulting sample contains 545 events, 543 of which have a finite persistence-based lifetime score.

For each event time t , let f_j denote the observed funding rate at settlement index j , let $W = 540$ be the 180-day pre-event window in eight-hour settlements, and let $K = 21$ be the seven-day outcome horizon. The pre-event funding path is fitted with a zero-intercept autoregressive model of order one (AR(1))

$$f_{j+1} = \phi f_j + \varepsilon_{j+1}, \quad j = t - W, \dots, t - 2,$$

where ε_{j+1} is the one-step AR(1) regression residual. Let $\hat{\phi}$ be the least-squares coefficient and $\hat{\sigma}_\varepsilon$ the sample standard deviation of the fitted AR(1) residuals. When $0 < \hat{\phi} < 1$, define

$$\hat{\lambda} := -\log \hat{\phi}, \quad \hat{\mathcal{I}}_{\text{life}} := \frac{f_t^2}{4\hat{\sigma}_\varepsilon^2 \hat{\lambda}}, \quad L := \frac{\hat{\mathcal{I}}_{\text{life}}}{\mathcal{I}_{\text{crit}}(0.05, 0.10)}.$$

Here $\mathcal{I}_{\text{crit}}(0.05, 0.10) = 4.2819236753 \dots$ nats is the canonical Gaussian threshold for level 0.05 and power 0.90. We call L the *lifetime score* in this empirical section. The normalization is monotone and leaves all rank statistics unchanged; it does not itself constitute a statistical certificate.

For comparison, let $s_{f,t}$ be the sample standard deviation of the W pre-event funding observations and define the contemporaneous edge magnitude

$$z_{\text{edge},t} := \frac{|f_t|}{s_{f,t}}.$$

Let $g_j := \text{sgn}(f_t)f_j$ orient the pre-event history by the current funding sign, and define its historical t -statistic

$$t_{\text{hist},t} := \frac{\bar{g}_t}{s_{g,t}/\sqrt{W}},$$

where $\bar{g}_t$ and $s_{g,t}$ are the sample mean and sample standard deviation of $(g_{t-W}, \dots, g_{t-1})$. Finally, define the seven-day sign-oriented outcome

$$Y_{7d,t} := \sum_{k=1}^K \text{sgn}(f_t)f_{t+k}.$$

With $K = 21$, define the same-sign survival length S_t^{sign} as the number of consecutive post-event settlements for which funding retains the event-time sign, capped at K ; equivalently, the count stops at the first k for which $\text{sgn}(f_t)f_{t+k} \leq 0$. Every event-level regressor uses only information available at time t ; the exercise is retrospective. Events entering the finite-score analysis are required to have finite values for all displayed statistics, so the sample standard deviations and other denominators used above are positive and finite on that analysis sample.

7.2. Results

For event-level variables U and V , write $\rho_S(U, V)$ for their sample Spearman rank correlation. In this subsection we suppress the event-time subscript t , so Y_{7d} , z_{edge} , and t_{hist} denote the vectors with entries $Y_{7d,t}$, $z_{\text{edge},t}$, and $t_{\text{hist},t}$ across the 543 finite-score events, while L denotes the corresponding lifetime-score vector. The notation $\text{rank}(U)$ means the sample midrank vector of U ; “residualizing $\text{rank}(U)$ on controls” means taking residuals from an ordinary least-squares projection with an intercept onto the midrank-transformed controls. Across the 543 finite-score events, the rank associations with next-seven-day signed funding are

$$\rho_S(L, Y_{7d}) = 0.4756, \quad \rho_S(z_{\text{edge}}, Y_{7d}) = 0.3833, \quad \rho_S(t_{\text{hist}}, Y_{7d}) = 0.3606.$$

Calendar-month cluster bootstrap 95% intervals (5,000 resamples) are $[0.293, 0.615]$, $[0.171, 0.545]$, and $[0.174, 0.524]$, respectively. The lifetime score and current edge magnitude are correlated, with $\rho_S(L, z_{\text{edge}}) = 0.7673$. Residualizing $\text{rank}(L)$ on $\text{rank}(z_{\text{edge}})$ leaves a Spearman association of 0.2824 with future signed funding, with a calendar-month cluster bootstrap 95% interval of $[0.128, 0.412]$. Deleting any one of the twelve asset–venue series leaves this incremental statistic between 0.2529 and 0.3092.

The top quintile by lifetime score contains 109 events: 100 have positive future signed funding, and the mean same-sign survival length S_t^{sign} is 16.97 settlements. The top quintile by instantaneous edge magnitude also contains 109 events: 93 have positive future signed funding, and the mean same-sign survival length is 14.97 settlements.

As a secondary robustness calculation, residualizing $\text{rank}(L)$ simultaneously on $\text{rank}(z_{\text{edge}})$ and $\text{rank}(t_{\text{hist}})$ leaves an association of 0.1278 with future signed funding. A three-month block bootstrap gives a 95% interval of $[0.012, 0.234]$, and the corresponding association is positive in both BTC and ETH subsamples.

The outcome is signed funding cashflow and persistence. Net basis-trade returns additionally depend on execution fees, spot/perpetual hedge drift, borrow, liquidation, collateral funding, slippage, and venue credit risk. The reported quantities measure persistence.




Figure 3: Association with next-seven-day sign-oriented funding for the same 543 events. Points report Spearman associations; horizontal bars report calendar-month cluster bootstrap 95% intervals from 5,000 resamples. The final row is the association of future funding with the lifetime-score rank after residualizing it on current-edge rank.




Figure 4: Average same-sign survival length S_t^{sign} over the 21-settlement outcome window for the top quintile by lifetime score and by instantaneous edge magnitude. The vertical axis starts at zero. Bars are descriptive sample means. Each group contains 109 events; 100/109 lifetime-score events and 93/109 current-edge events have positive seven-day signed funding.

8. Discussion

A finite information budget changes when an alpha-research project should be evaluated. Statistical significance describes evidence already observed. The survival calculation asks a prior question: whether the remaining opportunity can still supply the evidence required by the deployment rule. This suggests a screening stage before a long validation program begins. The inputs are the current signal strength, a model for persistence, the search breadth, and the desired error and power levels. If these inputs place a candidate below the survival frontier, the declared certification target is infeasible under the model even if observation continues until the opportunity is exhausted.

The same framework separates statistical scarcity from market implementation. Lifetime information determines how much can still be learned. The implementation coefficient in Proposition 6.3 determines how strongly research and arbitrage activity translate exploitable information into subsequent decay. Empirical tests of crowding therefore require variables for both channels. Capacity, borrow cost, turnover, market impact, and the number of active competitors are natural candidates for measuring implementation technology. This distinction also clarifies why a pooled relation between a statistical feasibility score and realized decay can be weak even when the equilibrium mechanism is present.

The multivariate value-information ratio provides a second allocation criterion for research. Two directions can require the same amount of statistical information and still differ economically because return covariance and trading-cost curvature assign different values to learning them. A research desk can therefore compare experiments by expected economic value per nat in addition to conventional quantities such as Sharpe ratio, power, or forecast error. In settings with scarce research capacity, this gives a natural way to rank which uncertainties are worth resolving first. The static frontier also points to a dynamic problem. Persistence may itself be learned while capital is deployed; the economic boundary may move with market conditions; and crowding may respond stochastically to observed performance. In such settings the remaining information budget becomes a state variable jointly determined by filtering, control, and strategic interaction. The resulting problem is an endogenous information-allocation problem in which the researcher chooses both exposure and the rate at which evidence is acquired.

9. Conclusion

The practical bottleneck in validating a decaying alpha is the finite supply of future evidence. Once the opportunity has a finite information budget, statistical validation becomes a feasibility constraint: the deployment rule specifies how much evidence is required, while the signal's strength and persistence determine how much evidence can still arrive. For the declared reliability contract, write $\mathcal{I}_{\text{req}} := \mathcal{I}_{\text{crit}}(\alpha, \beta)$. The quantity

$$\mathcal{I}_{\text{life}} - \mathcal{I}_{\text{req}}$$

therefore has an immediate decision interpretation. A negative value rules out the declared reliability target under the model; a positive value measures the information margin available after that target is met.

This viewpoint changes the research workflow. The reliability contract can be fixed first, search multiplicity can be charged explicitly, and the remaining information budget can be estimated before committing to a long validation exercise. Candidates with insufficient capacity can be discarded on feasibility grounds. For candidates with positive capacity, the residual information margin is a resource that can be compared with implementation cost, expected economic value, and the effect of competition on future lifetime.

The broader principle is that statistical evidence has an economic horizon. A validation rule is useful only when its information requirements are compatible with the lifetime of the opportunity it is meant to validate. Certified Alpha Capacity makes that compatibility explicit.

10. Code and data availability

The AlphaValue release accompanying this paper contains the numerical routines, command-line interface (CLI), synthetic examples, regression tests, and reproducibility material used for the reported calculations. The software archive is versioned at Zenodo under the concept DOI [10.5281/zenodo.23070742](https://doi.org/10.5281/zenodo.23070742). Third-party market data are redistributed only when their licenses permit it; source references and reproduction code are included. A clean editable installation uses `python -m pip install -e '.[test]'` followed by `python -m pytest -q`.

References

- [1] A. Y. Chen and T. Zimmermann. Open Source Cross-Sectional Asset Pricing. *Critical Finance Review*, 11(2):207–264, 2022. doi:10.1561/104.00000112.
- [2] A. Wald. *Sequential Analysis*. Wiley, New York, 1947.
- [3] D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer, New York, 1985.
- [4] P. V. Gapeev and G. Peskir. The Wiener sequential testing problem with finite horizon. *Stochastics and Stochastics Reports*, 76(1):59–75, 2004. doi:10.1080/10451120410001663753.
- [5] E. Ekström and J. Vaicenavicius. Bayesian sequential testing of the drift of a Brownian motion. *ESAIM: Probability and Statistics*, 19:626–648, 2015. doi:10.1051/ps/2015012.
- [6] E. Ekström, I. Karatzas, and J. Vaicenavicius. Bayesian sequential least-squares estimation for the drift of a Wiener process. *Stochastic Processes and their Applications*, 145:335–352, 2022. doi:10.1016/j.spa.2019.09.006.
- [7] H. Wang and A. Ramdas. Anytime-valid t -tests and confidence sequences for Gaussian means with unknown variance. *Sequential Analysis*, 44(1):56–110, 2025. doi:10.1080/07474946.2024.2428245.
- [8] N. W. Koning and S. van Meer. Anytime validity is free: inducing sequential tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 88(4):1366–1384, 2026. doi:10.1093/jrsssb/qkag050.
- [9] C. R. Harvey, Y. Liu, and H. Zhu. ... and the Cross-Section of Expected Returns. *Review of Financial Studies*, 29(1):5–68, 2016. doi:10.1093/rfs/hhv059.
- [10] D. H. Bailey and M. López de Prado. The Deflated Sharpe Ratio: Correcting for Selection Bias, Backtest Overfitting, and Non-Normality. *Journal of Portfolio Management*, 40(5):94–107, 2014. doi:10.3905/jpm.2014.40.5.094.
- [11] D. H. Bailey, J. Borwein, M. López de Prado, and Q. J. Zhu. The Probability of Backtest Overfitting. *Journal of Computational Finance*, 20(4):39–69, 2017.
- [12] D. C. Nachman. Optimal Stopping with a Horizon Constraint. *Mathematics of Operations Research*, 5(1):126–134, 1980. doi:10.1287/moor.5.1.126.

- [13] D. P. Kennedy. On a constrained optimal stopping problem. *Journal of Applied Probability*, 19(3):631–641, 1982.
- [14] P. I. Frazier and A. J. Yu. Sequential Hypothesis Testing under Stochastic Deadlines. In *Advances in Neural Information Processing Systems 20*, 2007.
- [15] S. Dayanik and A. J. Yu. Reward-Rate Maximization in Sequential Identification under a Stochastic Deadline. *SIAM Journal on Control and Optimization*, 51(4):2922–2948, 2013. doi:10.1137/100818005.
- [16] E. Ekström and Y. Wang. Stopping Problems with an Unknown State. *Journal of Applied Probability*, 61(2):515–528, 2024. doi:10.1017/jpr.2023.52.
- [17] S. Campbell, G. Gaitsgori, R. Groenewald, and I. Karatzas. Grab It Before It’s Gone: Testing Uncertain Rewards under a Stochastic Deadline. *Stochastic Processes and their Applications*, 201:105069, 2026. doi:10.1016/j.spa.2026.105069.
- [18] S. J. Grossman and J. E. Stiglitz. On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70(3):393–408, 1980.
- [19] R. Da, S. Nagel, and D. Xiu. The Statistical Limit of Arbitrage. NBER Working Paper No. 33070, 2024. doi:10.3386/w33070.
- [20] S. Basak and B. Croitoru. On the Role of Arbitrageurs in Rational Markets. *Journal of Financial Economics*, 81(1):143–173, 2006. doi:10.1016/j.jfineco.2004.11.004.
- [21] J.-P. Zigrand. A General Equilibrium Analysis of Strategic Arbitrage. *Journal of Mathematical Economics*, 40(8):923–952, 2004. doi:10.1016/j.jmateco.2003.09.002.
- [22] J.-P. Zigrand. Endogenous Market Integration, Manipulation and Limits to Arbitrage. *Journal of Mathematical Economics*, 42(3):301–314, 2006. doi:10.1016/j.jmateco.2004.12.010.
- [23] V. DeMiguel, A. Martín-Utrera, and R. Uppal. Can Competition Increase Profits in Factor Investing? *Management Science*, 71(7):5552–5571, 2025. doi:10.1287/mnsc.2022.02684.
- [24] R. D. McLean and J. Pontiff. Does Academic Research Destroy Stock Return Predictability? *Journal of Finance*, 71(1):5–32, 2016. doi:10.1111/jofi.12365.
- [25] H. Jacobs and S. Müller. Anomalies across the Globe: Once Public, No Longer Existent? *Journal of Financial Economics*, 135(1):213–230, 2020. doi:10.1016/j.jfineco.2019.06.004.
- [26] G. Kaplanski. The Race to Exploit Anomalies and the Cost of Slow Trading. *Journal of Financial Markets*, 62:100754, 2023. doi:10.1016/j.finmar.2022.100754.
- [27] D. Muravyev, N. D. Pearson, and J. M. Pollet. Anomalies and Their Short-Sale Costs. *Journal of Finance*, 80(6):3639–3694, 2025. doi:10.1111/jofi.13501.
- [28] S. Jaimungal and X. Shi. Short Communication: The Price of Information. *SIAM Journal on Financial Mathematics*, 15(3):SC54–SC67, 2024. doi:10.1137/24M1644791.
- [29] E. Lehrer and T. Wang. The Value of Information in Stopping Problems. *Economic Theory*, 78(2):619–648, 2024. doi:10.1007/s00199-023-01543-8.
- [30] F. S. Anand, J. H. Lee, and M. J. Realff. Optimal decision-oriented Bayesian design of experiments. *Journal of Process Control*, 20(9):1084–1091, 2010. doi:10.1016/j.jprocont.2010.06.011.

- [31] S. Zhong, W. Shen, T. Catanach, and X. Huan. Goal-Oriented Bayesian Optimal Experimental Design for Nonlinear Models Using Markov Chain Monte Carlo. *SIAM/ASA Journal on Uncertainty Quantification*, 14(1):19–47, 2026. doi:10.1137/24M1649344.
- [32] J. Pénasse. Understanding Alpha Decay. *Management Science*, 68(5):3966–3973, 2022. doi:10.1287/mnsc.2022.4353.
- [33] A. Varma. The Public-Signal Alpha Half-Life Hypothesis. SSRN Working Paper 7395678, 2026. doi:10.2139/ssrn.7395678.
- [34] M. S. Zulfiqar. When Alpha Dies: A Signal Autopsy Approach to Predicting Strategy Decay. SSRN Working Paper 7376818, 2026. doi:10.2139/ssrn.7376818.
- [35] C. Ma and P. Smith. On the Effect of Alpha Decay and Transaction Costs on the Multi-period Optimal Trading Strategy. arXiv:2502.04284, 2025.
- [36] E. Inan. Predictability of Funding Rates. SSRN Working Paper 5576424, 2025. doi:10.2139/ssrn.5576424.
- [37] S. He, S. Wang, and T. Zhang. A Shared Template Without Shared Feedback: Funding Rates in Cryptocurrency Perpetual Futures. SSRN Working Paper 6185958, 2026. doi:10.2139/ssrn.6185958.
- [38] T. Lau. The Funding Carry and a Cross-Venue Spread on Perpetual Futures: A Significance-Tested Study of Hyperliquid and Centralized Venues. SSRN Working Paper 6993978, 2026. doi:10.2139/ssrn.6993978.
- [39] P. Zhivkov, V. Todorov, and S. Georgiev. Temporal Dynamics of Market Microstructure in Cryptocurrency Perpetual Futures: Econometric Evidence from Centralized and Decentralized Exchanges. *International Journal of Financial Studies*, 14(5):103, 2026. doi:10.3390/ijfs14050103.
- [40] N. Alexander and F. J. Fabozzi. Measuring Strategy-Decay Risk: Minimum Regime Performance and the Durability of Systematic Investing. *Journal of Portfolio Management*, 52(4):198–219, 2026. doi:10.3905/jpm.2025.1.807.
- [41] R. Stephan. Sequential Tradeability Testing for Alpha Signals. SSRN Working Paper 6922558, 2026. doi:10.2139/ssrn.6922558.
- [42] J. R. Baxter and R. V. Chacon. Compactness of stopping times. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 40(3):169–181, 1977. doi:10.1007/BF00736045.

A. Capacity with a type-I error constraint

For comparison, let $c_\alpha(A)$ denote the maximal expected post-certification information time when only a type-I constraint is imposed:

$$c_\alpha(A) := \sup \left\{ \mathbb{E}_1[(A - S)\delta] : S \in \mathcal{T}_A, \delta \in \{0, 1\} \text{ is } \mathcal{G}_S\text{-measurable, } \mathbb{P}_0(\delta = 1) \leq \alpha \right\}.$$

Independent randomization may be adjoined at time zero when convexification is needed. The corresponding economic value is

$$\text{CAC}_\alpha(A) := \frac{\sigma^2}{2\gamma} c_\alpha(A).$$

This constrained-stopping problem includes the α -lottery floor. Reliable certification adds a power constraint, which requires a prescribed probability of deployment under the alternative.

Theorem A.1 (Nonasymptotic converse without a fixed power target). *Let*

$$r_\alpha(A) = \frac{c_\alpha(A)}{A}.$$

Then

$$\frac{A}{2}(1 - r_\alpha(A)) \geq \text{kl}(r_\alpha(A), \alpha).$$

Equivalently $r_\alpha(A)$ is bounded above by the unique root $\bar{r} \in [\alpha, 1)$ of

$$A(1 - \bar{r}) = 2 \text{kl}(\bar{r}, \alpha).$$

Proof. An independent immediate level- α lottery gives $c_\alpha(A) \geq \alpha A$, so $r_\alpha(A) \geq \alpha$. Let (S_n, δ_n) be a sequence of feasible rules with rewards $R_n \uparrow c_\alpha(A)$, and set $r_n = R_n/A$, $p_n = \mathbb{P}_1(\delta_n = 1)$, and $q_n = \mathbb{P}_0(\delta_n = 1)$. Because the reward is at most A on deployment, $p_n \geq r_n$; also $q_n \leq \alpha$. With $S_n = A$ on nondeployment,

$$\frac{1}{2}(A - R_n) = D_{\text{KL}}(\mathbb{P}_1^{S_n} \parallel \mathbb{P}_0^{S_n}) \geq \text{kl}(p_n, q_n).$$

For all sufficiently large n , either $r_n \geq \alpha$, in which case $\text{kl}(p_n, q_n) \geq \text{kl}(r_n, \alpha)$, or $r_\alpha(A) = \alpha$ and the desired inequality follows directly by continuity at $r = \alpha$. Letting $n \rightarrow \infty$ gives

$$\frac{A}{2}(1 - r_\alpha(A)) \geq \text{kl}(r_\alpha(A), \alpha).$$

The root is unique because $A(1 - r) - 2 \text{kl}(r, \alpha)$ is positive at $r = \alpha$, negative as $r \uparrow 1$, and strictly decreasing for $r > \alpha$. $\square$

As $A \downarrow 0$, the capacity fraction tends to α . The limit comes from the permitted false-deployment probability and contains no learned evidence about the alpha. One may therefore report either reliable CAC with an explicit power target or the evidence-added quantity $c_\alpha(A) - \alpha A$.

A.1. Constructive likelihood-threshold lower bound

For $h > 0$ let

$$\tau_h = \inf\{s \geq 0 : Z_s \geq h\}.$$

For Brownian motion with drift $\mu \in \mathbb{R}$, the first-passage cumulative distribution function (CDF), for $t > 0$, is

$$F_\mu(t; h) = \Phi\left(\frac{\mu t - h}{\sqrt{t}}\right) + e^{2\mu h} \Phi\left(\frac{-\mu t - h}{\sqrt{t}}\right).$$

Choose the finite-horizon boundary h_A as the unique solution of

$$F_{-1/2}(A; h_A) = \alpha.$$

Then the stopping rule that deploys at $\tau_{h_A} < A$ has exact type-I probability α and value

$$L_\alpha(A) = \int_0^A F_{1/2}(t; h_A) dt \leq c_\alpha(A).$$

The simpler anytime-valid boundary $h_\infty = \log(1/\alpha)$ satisfies $\mathbb{P}_0(\tau_{h_\infty} < \infty) = \alpha$.

Theorem A.2 (Asymptotic information cost of certification). *For fixed $\alpha \in (0, 1)$,*

$$A - c_\alpha(A) \longrightarrow 2 \log(1/\alpha) \quad (A \rightarrow \infty).$$

Consequently

$$\frac{\text{CAC}_\alpha(A)}{V^{PI}} = 1 - \frac{\log(1/\alpha)}{\mathcal{I}_{\text{life}}} + o(\mathcal{I}_{\text{life}}^{-1}),$$

and, in economic units,

$$V^{PI} - \text{CAC}_\alpha(A) \longrightarrow \frac{\sigma^2}{\gamma} \log(1/\alpha).$$

Proof. Under $\mathbb{P}_1$, $Z_s = \tilde{B}_s^{(1)} + s/2$, so the untruncated first-passage time to $h = \log(1/\alpha)$ is almost surely finite and has mean $2h$. Thus the anytime-valid threshold rule satisfies

$$A - c_\alpha(A) \leq \mathbb{E}_1(\tau_h \wedge A) \rightarrow 2h.$$

The nonasymptotic converse gives, writing $g_A = A - c_\alpha(A)$,

$$\frac{g_A}{2} \geq \text{kl}(1 - g_A/A, \alpha).$$

The preceding achievability makes g_A bounded, so the right side converges to $\log(1/\alpha)$. Hence $\liminf g_A \geq 2h$, proving equality. $\square$

For $\alpha = 0.05$, the asymptotic information-time tax is $2 \log 20 = 5.9914645471 \dots$ units of A , i.e. $\log 20 = 2.9957322736 \dots$ KL nats.

B. Constrained stopping representation

Constrained stopping with horizon restrictions and Lagrange multipliers has a classical foundation [12, 13]; stochastic-deadline testing and identification provide closely related time-constrained formulations [14–17]. Exact finite- A capacity can be written as a convexified finite-horizon stopping problem. Here a randomized rule consists of a $\mathbb{G}$ -stopping time $S \leq A$ together with a $\mathcal{G}_S$ -measurable deployment probability $\rho \in [0, 1]$. For $0 \leq s \leq A$, define

$$\mathcal{T}_{s,A} := \{S \in \mathcal{T}_A : S \geq s \text{ a.s.}\}.$$

For $z \in \mathbb{R}$, let $\mathbb{P}_{s,z}^{(1)}$ denote the law under which the information-time likelihood process starts from $Z_s = z$ and evolves as $dZ_u = d\tilde{B}_u^{(1)} + \frac{1}{2}du$ for $u \geq s$, and let $\mathbb{E}_{s,z}^{(1)}$ be its expectation. Convexification by randomized stopping rules yields a standard Lagrange dual.

Theorem B.1 (Lagrange representation for the type-I constrained problem). *For the randomized stopping class,*

$$c_\alpha(A) = \inf_{\lambda \geq 0} \{\lambda \alpha + V_\lambda(0, 0)\},$$

where, under $\mathbb{P}_1$,

$$V_\lambda(s, z) = \sup_{S \in \mathcal{T}_{s,A}} \mathbb{E}_{s,z}^{(1)} \left[(A - S - \lambda e^{-Z_S})_+ \right].$$

The positive part encodes the option to decline deployment.

Proof. For a randomized rule, let $\rho \in [0, 1]$ be the $\mathcal{G}_S$ -measurable conditional probability of deployment at the stopping time S . The type-I deployment probability is $\mathbb{E}_0[\rho]$. Since $d\mathbb{P}_0/d\mathbb{P}_1 = e^{-Z_S}$ on $\mathcal{G}_S$,

$$\mathbb{E}_0[\rho] = \mathbb{E}_1[e^{-Z_S}\rho].$$

For a multiplier $\lambda \geq 0$, the Lagrangian is

$$\lambda\alpha + \mathbb{E}_1\left[(A - S - \lambda e^{-Z_S})\rho\right].$$

For fixed S , pointwise maximization over $0 \leq \rho \leq 1$ gives $\rho = 1$ when the bracket is positive and $\rho = 0$ when it is negative. The resulting value is

$$\lambda\alpha + \mathbb{E}_1\left[(A - S - \lambda e^{-Z_S})_+\right].$$

Taking the supremum over stopping times gives $\lambda\alpha + V_\lambda(0, 0)$, and then the infimum over $\lambda \geq 0$ gives the dual upper bound. To obtain equality, use the standard randomized-stopping formulation. On a bounded horizon, randomized stopping times are compact in the Baxter–Chacon topology [42]. In the present Brownian model the reward is bounded and the stopped likelihood ratios are uniformly integrable, so the expected reward and type-I probability are continuous under this convergence. Randomization makes the attainable set of (type-I probability, reward) pairs convex. Since the no-deploy rule has type-I probability zero, the constraint is strictly feasible whenever $\alpha > 0$. The separating-hyperplane theorem applied to this compact convex set therefore produces a nonnegative multiplier and no duality gap. $\square$

Randomization is part of the admissible decision rule in the economic model. It also convexifies the constrained-stopping problem, which gives the exact primal–dual representation above under the stated compactness and continuity conditions.

Let

$$\mathcal{L}f(z) := \frac{1}{2}f''(z) + \frac{1}{2}f'(z)$$

be the $\mathbb{P}_1$ -generator of the information-time likelihood process. In the continuation region of the one-multiplier problem,

$$\partial_s V_\lambda + \mathcal{L}V_\lambda = 0,$$

with obstacle

$$\mathcal{O}_\lambda(s, z) := (A - s - \lambda e^{-z})_+.$$

In the positive-obstacle region,

$$(\partial_s + \mathcal{L})(A - s - \lambda e^{-z}) = -1,$$

because e^{-Z_s} is a $\mathbb{P}_1$ -martingale. Standard one-dimensional stopping arguments lead to an upper boundary $b_\lambda(s)$: deploy when $Z_s \geq b_\lambda(s)$.

For $\lambda > 0$, two structural simplifications are exact:

$$V_\lambda(s, z) = V_1(s, z - \log \lambda), \quad b_\lambda(s) = \log \lambda + b_1(s).$$

If smooth fit and the standard boundary regularity hold, the early-exercise premium formula gives

$$V_\lambda(s, z) = \int_s^A \mathbb{P}_{s,z}^{(1)}\{Z_u \geq b_\lambda(u)\} du,$$

and the boundary solves

$$A - s - \lambda e^{-b_\lambda(s)} = \int_s^A \Phi\left(\frac{b_\lambda(s) + (u - s)/2 - b_\lambda(u)}{\sqrt{u - s}}\right) du.$$

This integral equation characterizes the free boundary. Related free-boundary formulations arise in finite-horizon Wiener testing and stochastic-deadline learning problems.

For reliable CAC with a power constraint, introduce a second multiplier $\omega \geq 0$ and define

$$V_{\lambda,\omega}(s, z) := \sup_{S \in \mathcal{T}_{s,A}} \mathbb{E}_{s,z}^{(1)} \left[(A - S + \omega - \lambda e^{-Zs})_+ \right].$$

The corresponding two-multiplier obstacle is

$$(A - s + \omega - \lambda e^{-z})_+,$$

and the dual objective is

$$\lambda \alpha - \omega(1 - \beta) + V_{\lambda,\omega}(0, 0),$$

whenever the standard strong-duality hypotheses hold. The exact Gaussian feasibility threshold above should be checked before solving this free-boundary problem: below A_{crit} the desired reliable certificate is impossible regardless of stopping optimization.

C. Unknown scale and predictable regressors

Consider an AR(1) predictor with a Gaussian return equation in which the return loading, persistence, and innovation variances are unknown. The predictor generates a stochastic regression design. We combine a universal KL converse with two constructive benchmarks: an exact fixed-design t frontier in an exogenous submodel and a predictable-design e-process bound under a declared variance and effect-size envelope.

C.1. A pairwise KL necessary condition

Let $\mathfrak{P}_0$ denote an economic-null family of probability laws on the observed sample space and $\mathfrak{P}_1$ an alternative family. Suppose a binary deployment decision $\delta \in \{0, 1\}$ must satisfy

$$\sup_{\mathbb{Q}_0 \in \mathfrak{P}_0} \mathbb{Q}_0(\delta = 1) \leq \alpha, \quad \inf_{\mathbb{Q}_1 \in \mathfrak{P}_1} \mathbb{Q}_1(\delta = 1) \geq 1 - \beta,$$

with $1 - \beta > \alpha$.

Theorem C.1 (Pairwise KL necessary condition). *For every admissible alternative/boundary pair $(\mathbb{Q}_1, \mathbb{Q}_0) \in \mathfrak{P}_1 \times \mathfrak{P}_0$ with $\mathbb{Q}_1 \ll \mathbb{Q}_0$,*

$$D_{\text{KL}}(\mathbb{Q}_1 \| \mathbb{Q}_0) \geq \text{kl}(1 - \beta, \alpha).$$

Thus a single compatible pair below this threshold rules out a uniformly reliable certificate over the whole composite class.

Proof. Data processing from the observed path to δ gives

$$D_{\text{KL}}(\mathbb{Q}_1 \| \mathbb{Q}_0) \geq \text{kl}(\mathbb{Q}_1(\delta = 1), \mathbb{Q}_0(\delta = 1)).$$

On $p \geq 1 - \beta > \alpha \geq q$, binary relative entropy is minimized at $(p, q) = (1 - \beta, \alpha)$. $\square$

C.2. Stationary information rate in the AR(1) predictor model

Fix an integer horizon $L \in \mathbb{N}$. Write the model as

$$X_{t+1} = \phi X_t + \xi_{t+1}, \quad Y_{t+1} = \theta X_t + \varepsilon_{t+1},$$

where, in the orthogonal Gaussian benchmark, (ξ_t) and (ε_t) are independent centered Gaussian innovation sequences with variances v_X and v_R , respectively. When $|\phi| < 1$, X has a unique stationary Gaussian law, denoted by $\mathbb{P}_{\phi, v_X}^X$, with expectation $\mathbb{E}_{\phi, v_X}^X$. Compare a true loading θ to an economic-boundary companion b with the same nuisance parameters. In the orthogonal-innovation submodel the conditional return-channel KL is

$$\mathcal{I}_L(\theta : b \mid X) = \frac{(\theta - b)^2}{2v_R} \sum_{t=0}^{L-1} X_t^2.$$

Under the stationary signal law,

$$\mathbb{E}_{\phi, v_X}^X[X_t^2] = \frac{v_X}{1 - \phi^2},$$

so the expected information rate is

$$\kappa(\theta, \phi, v_R, v_X; b) = \frac{(\theta - b)^2 v_X}{2v_R(1 - \phi^2)}.$$

For any declared compact rectangular parameter set $\mathcal{R}$ contained in $\{(\theta, \phi, v_R, v_X) : |\phi| < 1, v_R > 0, v_X > 0\}$, define the robust lower information rate

$$\underline{\kappa}(b; \mathcal{R}) := \inf_{(\theta, \phi, v_R, v_X) \in \mathcal{R}} \kappa(\theta, \phi, v_R, v_X; b).$$

Because the displayed rate is monotone in the coefficient gap, variances, and $|\phi|$ in the relevant directions, its extrema over a rectangle occur at endpoints. If $\underline{\kappa}(b; \mathcal{R}) > 0$, every uniformly reliable rule must obey

$$L \geq L_{\text{KL}}^{\text{nec}} := \left\lceil \frac{\text{kl}(1 - \beta, \alpha)}{\underline{\kappa}(b; \mathcal{R})} \right\rceil.$$

If the declared uncertainty rectangle touches the economic boundary, the robust rate floor is zero; later data may shrink the uncertainty set.

C.3. Fixed-design benchmark with unknown variance

For the exogenous fixed-design regression with $n \geq 2$ and $Q > 0$,

$$Y_i = \theta X_i + \varepsilon_i, \quad \varepsilon_i \stackrel{iid}{\sim} N(0, v_R), \quad Q := \sum_{i=1}^n X_i^2,$$

define

$$\hat{\theta}_n := \frac{\sum_{i=1}^n X_i Y_i}{Q}, \quad \hat{v}_n := \frac{1}{n-1} \sum_{i=1}^n (Y_i - \hat{\theta}_n X_i)^2, \quad T_n := \frac{(\hat{\theta}_n - b)\sqrt{Q}}{\sqrt{\hat{v}_n}}.$$

Under an alternative $\theta > b$, this one-sided no-intercept statistic has a noncentral- t law

$$T_n \sim t_\nu(v), \quad \nu = n - 1, \quad v = (\theta - b)\sqrt{Q/v_R}.$$

Let $F_{\nu, v}$ denote the cdf of the noncentral- t law $t_\nu(v)$, and let $t_{\nu, 1-\alpha}$ denote the $(1 - \alpha)$ -quantile of the central t_ν law. Define $v_\nu^*(\alpha, \beta)$ by

$$1 - F_{\nu, v_\nu^*}(t_{\nu, 1-\alpha}) = 1 - \beta,$$

and put

$$\mathcal{I}_t(\alpha, \beta; \nu) := \frac{1}{2}(v_\nu^*)^2.$$

For a declared uncertainty set, define its minimum alternative gap and maximum return variance by $d_- := \inf(\theta - b) > 0$ and $\bar{v}_R := \sup v_R < \infty$, with the infimum and supremum taken over that set, and put

$$\underline{\mathcal{I}}(Q) = \frac{d_-^2 Q}{2\bar{v}_R}.$$

Then the binary-KL converse yields impossibility when

$$\underline{\mathcal{I}}(Q) < \text{kl}(1 - \beta, \alpha),$$

whereas the ordinary terminal t test has power at least $1 - \beta$ uniformly over the declared set when

$$\underline{\mathcal{I}}(Q) \geq \mathcal{I}_t(\alpha, \beta; n - 1).$$

This produces an explicit finite-sample unresolved region between a universal necessary threshold and a concrete sufficient threshold.

At $\alpha = 0.05$, power 0.90, the information-equivalent t thresholds are 4.888842, 4.551104, 4.382653 and 4.331221 nats for $n = 12, 24, 60, 120$ respectively, converging to the known-scale Gaussian value 4.281924.

The unknown-scale penalty here is the finite-sample cost of Studentization. Koning and van Meer [8] show that fixed- n tests, including the traditional t test, can be sequentialized while matching their terminal power, so this penalty is distinct from anytime validity itself.

C.4. Sequential certification with predictable regressors

Anytime-valid inference with unknown Gaussian variance provides a useful reference point for the sequential construction below [7]. The exact t calculation conditions on an exogenous full design. For a predictable regression, a natural alternative is to index the certificate by the realized design clock, avoiding assumptions on the future signal path.

Let $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ be the discrete-time observation filtration and let X_t be $\mathcal{F}_t$ -measurable before Y_{t+1} is observed. After a fixed calibration time, assume

$$Y_{t+1} = \theta X_t + \varepsilon_{t+1}, \quad \varepsilon_{t+1} \mid \mathcal{F}_t \sim N(0, v), \quad v \leq \bar{v}.$$

For each fixed parameter pair (θ, v) , denote the resulting law by $\mathbb{P}_{\theta, v}$. Test $H_0 : \theta \leq b$ against the separated class $H_1 : \theta \geq b + d$, $d > 0$. For a declared remaining design-energy budget $q > 0$, let

$$\ell_\alpha := \log(1/\alpha), \quad \lambda_q := \sqrt{\frac{2\ell_\alpha}{\bar{v}q}},$$

$$Q_t := \sum_{s < t} X_s^2, \quad U_t := \sum_{s < t} X_s(Y_{s+1} - bX_s),$$

and define the e-process

$$E_t := \exp \left\{ \lambda_q U_t - \frac{1}{2} \lambda_q^2 \bar{v} Q_t \right\}.$$

Let $\tau_q := \inf\{t : Q_t \geq q\}$ be the design-clock hitting time and let $\tau_E := \inf\{t : E_t \geq 1/\alpha\}$ be the certification crossing time.

Theorem C.2 (Predictable-design e-process certificate). *Under every null parameter value $\theta \leq b$, $v \leq \bar{v}$, (E_t) is a nonnegative supermartingale and therefore*

$$\sup_{\theta \leq b, v \leq \bar{v}} \mathbb{P}_{\theta, v} \left\{ \sup_t E_t \geq 1/\alpha \right\} \leq \alpha.$$

Under every alternative $\theta \geq b + d$, $v \leq \bar{v}$, if the design clock reaches q then

$$\mathbb{P}_{\theta, v} \{ \tau_E > \tau_q, \tau_q < \infty \} \leq \exp \left[-(\sqrt{\mathcal{I}_q} - \sqrt{\ell_\alpha})^2 \right], \quad \mathcal{I}_q := \frac{d^2 q}{2\bar{v}},$$

for $\mathcal{I}_q > \ell_\alpha$. Consequently, if $\mathbb{P}_{\theta, v}(\tau_q < \infty) = 1$ throughout the alternative class, the sufficient information condition

$$\mathcal{I}_q \geq \mathcal{I}_e^{\text{sup}}(\alpha, \beta) := \left[\sqrt{\log(1/\alpha)} + \sqrt{\log(1/\beta)} \right]^2$$

guarantees $\mathbb{P}_{\theta, v}(\tau_E \leq \tau_q) \geq 1 - \beta$ throughout the alternative class. Without almost-sure reachability of q , the preceding displayed bound remains a valid bound on failure among paths for which $\tau_q < \infty$.

Proof. Under the null, the one-step conditional moment generating function is at most one because the drift contribution is nonpositive and $v \leq \bar{v}$. This proves the supermartingale claim. Under an alternative, write $N_t := \sum_{s < t} X_s \varepsilon_{s+1}$. Then

$$\log E_t \geq (\lambda_q d - \tfrac{1}{2} \lambda_q^2 \bar{v}) Q_t + \lambda_q N_t.$$

On the event $\tau_q < \infty$, failure to cross by τ_q implies $N_{\tau_q} \leq -c_q Q_{\tau_q}$ for

$$c_q := d - \frac{1}{2} \lambda_q \bar{v} - \frac{\ell_\alpha}{\lambda_q q}.$$

The process $\exp\{-\zeta N_t - \zeta^2 \bar{v} Q_t/2\}$ is a nonnegative supermartingale for every $\zeta > 0$. Hence, on the event $N_{\tau_q} \leq -c_q Q_{\tau_q}$,

$$-\zeta N_{\tau_q} - \frac{1}{2} \zeta^2 \bar{v} Q_{\tau_q} \geq \left(\zeta c_q - \frac{1}{2} \zeta^2 \bar{v} \right) q.$$

Markov's inequality and optional sampling therefore give

$$\mathbb{P}_{\theta, v} \{ N_{\tau_q} \leq -c_q Q_{\tau_q} \} \leq \exp \left[-q \sup_{\zeta > 0} \left(\zeta c_q - \frac{1}{2} \zeta^2 \bar{v} \right) \right] = \exp \left[-\frac{q c_q^2}{2\bar{v}} \right].$$

For the displayed λ_q , $c_q = d - \sqrt{2\ell_\alpha \bar{v}/q}$, so the exponent is $(\sqrt{\mathcal{I}_q} - \sqrt{\ell_\alpha})^2$. Setting this exponent at least $\log(1/\beta)$ proves the result. $\square$

For $\alpha = 0.05$ and $\beta = 0.10$,

$$\text{kl}(0.9, 0.05) = 2.376205, \quad \mathcal{I}_e^{\text{sup}} = 10.551099 \text{ nats}.$$

The first is a portable converse; the second is a conservative constructive guarantee under a variance envelope and a random predictable design.

If d or $\bar{v}$ comes from a previously learned confidence set, that set's coverage failure probability enters the total error budget separately.

C.5. Discrete information–value identity

To make the economic comparison explicit, fix an integer trading horizon $L \in \mathbb{N}$ and consider positions $q_t \in \mathbb{R}$, initial inventory $q_{-1} = 0$, instantaneous reward

$$\theta X_t q_t - \frac{\gamma}{2} q_t^2 - \frac{\lambda_{\text{tr}}}{2} (q_t - q_{t-1})^2, \quad \gamma > 0, \quad \lambda_{\text{tr}} \geq 0,$$

and terminal inventory penalty $\eta_T q_{L-1}^2/2$, $\eta_T \geq 0$. For a fixed signal persistence $\phi \in (-1, 1)$, introduce auxiliary Bellman coefficients K_t and Λ_t^q by setting $K_L := \eta_T$ and $\Lambda_L^q := 0$, and define backward for $t = L-1, \dots, 0$

$$D_t := \gamma + \lambda_{\text{tr}} + K_{t+1}, \quad a_t := \frac{\lambda_{\text{tr}}}{D_t}, \quad g_t(\phi) := \frac{1 + \phi \Lambda_{t+1}^q}{D_t},$$

$$K_t := \lambda_{\text{tr}} \frac{\gamma + K_{t+1}}{D_t}, \quad \Lambda_t^q := \lambda_{\text{tr}} g_t(\phi).$$

The full-information optimal position under loading θ has the linear feedback form

$$q_t^\theta = a_t q_{t-1}^\theta + \theta g_t(\phi) X_t,$$

and the boundary-world oracle is $q_t^b = a_t q_{t-1}^b + b g_t(\phi) X_t$. Thus D_t is the Bellman curvature and $g_t(\phi)$ is the unit-loading signal gain. Under stationarity, define the boundary-relative expected regret by

$$R_L^{\theta:b} := \frac{v_X}{2(1-\phi^2)} (\theta - b)^2 \sum_{t=0}^{L-1} D_t g_t(\phi)^2.$$

Define the corresponding orthogonal-innovation return-channel information by

$$\mathcal{I}_{L,\text{return}}^{\theta:b} := \frac{L v_X}{2(1-\phi^2)} \frac{(\theta - b)^2}{v_R}.$$

Therefore:

Theorem C.3 (Discrete information–value identity).

$$R_L^{\theta:b} = \Pi_L(\phi, v_R) \mathcal{I}_{L,\text{return}}^{\theta:b}, \quad \Pi_L(\phi, v_R) = v_R \frac{1}{L} \sum_{t=0}^{L-1} D_t g_t(\phi)^2.$$

The signal variance v_X and the coefficient separation $\theta - b$ cancel exactly. With zero trading cost and zero terminal penalty, $D_t g_t^2 = 1/\gamma$ and $\Pi_L = v_R/\gamma$, the discrete counterpart of the scalar continuous-time factor σ^2/γ .

Proof. Divide the displayed regret by the displayed return-channel information:

$$\frac{R_L^{\theta:b}}{\mathcal{I}_{L,\text{return}}^{\theta:b}} = v_R \frac{1}{L} \sum_{t=0}^{L-1} D_t g_t(\phi)^2.$$

All factors involving v_X , $1 - \phi^2$, and $(\theta - b)^2$ cancel. The zero-cost specialization follows by substituting $D_t g_t^2 = 1/\gamma$. $\square$

If the contemporaneous Gaussian innovations satisfy $\rho := \text{Corr}(\xi_{t+1}, \varepsilon_{t+1})$ with $|\rho| < 1$, while retaining variances v_X and v_R , the full-joint Gaussian KL is

$$\mathcal{I}_{L,\text{full}}^{\theta:b} = \frac{1}{1-\rho^2} \mathcal{I}_{L,\text{return}}^{\theta:b},$$

and the corresponding full-information price per nat is

$$\Pi_{L,\text{full}} = (1-\rho^2)v_R \frac{1}{L} \sum_t D_t g_t^2.$$

Thus zero correlation is the least-informative member of this Gaussian family and is a legitimate least-informative parameter configuration for a uniform converse when it belongs to the model class.

C.6. Normal–Inverse-Gamma confidence sequences

The software also includes a proper Normal–Inverse-Gamma (NIG) likelihood-mixture confidence sequence for horizon-free joint control of a regression coefficient and an unknown Gaussian variance. To define the object explicitly, fix a confidence level $1-\alpha$ and predeclared hyperparameters

$$m_0 \in \mathbb{R}, \quad \kappa_0 > 0, \quad a_0 > 0, \quad b_0 > 0,$$

corresponding to the proper mixture

$$v \sim \text{InvGamma}(a_0, b_0), \quad \theta \mid v \sim N(m_0, v/\kappa_0),$$

where $\text{InvGamma}(a_0, b_0)$ uses the shape–scale density proportional to $v^{-a_0-1}e^{-b_0/v}$ on $(0, \infty)$. We work under the predictable no-intercept Gaussian regression $Y_{t+1} = \theta X_t + \varepsilon_{t+1}$ with $\varepsilon_{t+1} \mid \mathcal{F}_t \sim N(0, v)$. For any $n \geq 1$ with positive design energy, define

$$\begin{aligned} q_n &:= \sum_{t=0}^{n-1} X_t^2, & s_{xy,n} &:= \sum_{t=0}^{n-1} X_t Y_{t+1}, & s_{yy,n} &:= \sum_{t=0}^{n-1} Y_{t+1}^2, \\ \hat{\theta}_n &:= \frac{s_{xy,n}}{q_n}, & RSS_n(\theta) &:= \sum_{t=0}^{n-1} (Y_{t+1} - \theta X_t)^2, & RSS_n &:= RSS_n(\hat{\theta}_n). \end{aligned}$$

Set

$$\begin{aligned} \kappa_n &:= \kappa_0 + q_n, & m_n &:= \frac{\kappa_0 m_0 + s_{xy,n}}{\kappa_n}, & a_n &:= a_0 + \frac{n}{2}, \\ b_n &:= b_0 + \frac{1}{2} \left(s_{yy,n} + \kappa_0 m_0^2 - \kappa_n m_n^2 \right) \end{aligned}$$

and, writing $\Gamma(\cdot)$ here for Euler’s gamma function,

$$K_n := \frac{1}{2} \log \frac{\kappa_0}{\kappa_n} + a_0 \log b_0 - \log \Gamma(a_0) + \log \Gamma(a_n) - a_n \log b_n.$$

For a point null (θ, v) with $v > 0$, the corresponding log likelihood-mixture e-value is

$$\log E_n^{\text{NIG}}(\theta, v) := K_n + \frac{n}{2} \log v + \frac{RSS_n(\theta)}{2v}.$$

Inverting this e-process gives the exact time-uniform two-parameter confidence set

$$\mathcal{C}_n^{\text{NIG}}(\alpha) := \left\{ (\theta, v) \in \mathbb{R} \times (0, \infty) : \log E_n^{\text{NIG}}(\theta, v) < \log(1/\alpha) \right\}.$$

Let R_n denote the half-width of the θ -coordinate projection of the closure of $\mathcal{C}_n^{\text{NIG}}(\alpha)$ around $\hat{\theta}_n$. If

$$\frac{q_n}{n} \rightarrow \kappa_\infty > 0, \quad \frac{RSS_n}{n} \rightarrow v_\infty > 0,$$

$\hat{\theta}_n$ remains bounded, and $(m_0, \kappa_0, a_0, b_0)$ are fixed, then a direct expansion of the displayed closed-form quantities gives

$$R_n^2 \sim \frac{2v_\infty \log n}{\kappa_\infty n}.$$

The logarithmic factor comes from this horizon-free mixture construction. For a finite-lived opportunity, the confidence sequence can monitor nuisance parameters while a horizon-matched test targets certification against the declared economic boundary.

C.7. Finite-horizon certification with bounded nuisance parameters

For a declared integer economic horizon $L \in \mathbb{N}$, certification can also be written directly as a valid test against a fixed economic boundary $b \in \mathbb{R}$. Throughout this subsection, retain the orthogonal Gaussian innovation assumptions from Section C.2: (ξ_t) and (ε_t) are independent centered Gaussian innovation sequences with variances v_X and v_R . Consider

$$X_{t+1} = \phi X_t + \xi_{t+1}, \quad Y_{t+1} = \theta X_t + \varepsilon_{t+1},$$

with the compact nuisance rectangle

$$\begin{aligned} -1 < \phi_- < \phi_+ < 1, \quad \phi \in [\phi_-, \phi_+], \\ 0 < v_X^- \leq v_X \leq v_X^+ < \infty, \quad 0 < v_R^- \leq v_R \leq v_R^+ < \infty. \end{aligned}$$

For an admissible parameter vector $\vartheta = (\theta, \phi, v_R, v_X)$, write $\mathbb{P}_\vartheta$ for the induced law and define

$$\Theta_0 := \{\vartheta : \theta \leq b\}, \quad \Theta_1(d) := \{\vartheta : \theta \geq b + d\}, \quad d > 0,$$

where both sets are understood to satisfy the nuisance rectangle above. The compactness is substantive. If the return variance is unbounded, then for any fixed calendar horizon the null and alternative KL divergence can be driven to zero by sending $v_R \rightarrow \infty$; if the signal innovation variance can vanish from $X_0 = 0$, the loading models can become observationally identical. Therefore no nontrivial finite uniform certification horizon exists without a positive information-rate floor.

Fix a target design energy $q > 0$. Let predictable weights $0 \leq c_t \leq 1$ be chosen sequentially, with each c_t measurable with respect to the predictor information available before Y_{t+1} is observed; allow a fractional final weight so that the target energy can be hit exactly. Define the weighted design clock

$$Q_n^{(c)} := \sum_{t=0}^{n-1} c_t^2 X_t^2, \quad \tau_q^{(c)} := \inf\{n \geq 1 : Q_n^{(c)} = q\}.$$

On $\{\tau_q^{(c)} < \infty\}$ define the weighted score

$$M_q := \sum_{t=0}^{\tau_q^{(c)}-1} c_t X_t (Y_{t+1} - b X_t).$$

Define the certification time

$$\tau_{\text{cert}}^{(q)} := \begin{cases} \tau_q^{(c)}, & \tau_q^{(c)} < \infty \text{ and } M_q \geq z_{1-\alpha} \sqrt{v_R^+ q}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Equivalently, certification occurs on the event $\{\tau_{\text{cert}}^{(q)} < \infty\}$. The rule satisfies

$$\sup_{\vartheta \in \Theta_0} \mathbb{P}_{\vartheta} \{\tau_{\text{cert}}^{(q)} < \infty\} \leq \alpha.$$

If

$$d\sqrt{q/v_R^+} \geq z_{1-\alpha},$$

then, for every $\vartheta \in \Theta_1(d)$,

$$\mathbb{P}_{\vartheta} \{\tau_{\text{cert}}^{(q)} < \infty\} \geq \Phi \left(d\sqrt{\frac{q}{v_R^+}} - z_{1-\alpha} \right).$$

The condition places the rule in the regime where the worst-case variance is v_R^+ ; it is automatically satisfied whenever this lower bound is at least $1/2$. Writing $\mathcal{I}(q) := d^2 q / (2v_R^+)$ then gives the canonical power curve $\Phi(\sqrt{2\mathcal{I}(q)} - z_{1-\alpha})$.

For $\eta \in (0, 1)$, let $q_L(\eta)$ be a uniform lower $(1 - \eta)$ -quantile of $Q_L := \sum_{t=0}^{L-1} X_t^2$, meaning that, under the stationary predictor law $\mathbb{P}_{\phi, v_X}^X$,

$$\inf_{\phi \in [\phi_-, \phi_+], v_X \in [v_X^-, v_X^+]} \mathbb{P}_{\phi, v_X}^X \{Q_L \geq q_L(\eta)\} \geq 1 - \eta.$$

Set the target weighted design energy to $q = q_L(\eta)$. If $d\sqrt{q_L(\eta)/v_R^+} \geq z_{1-\alpha}$, then

$$\inf_{\vartheta \in \Theta_1(d)} \mathbb{P}_{\vartheta} \{\tau_{\text{cert}}^{(q)} \leq L\} \geq \left[\Phi \left(d\sqrt{\frac{q_L(\eta)}{v_R^+}} - z_{1-\alpha} \right) - \eta \right]_+.$$

Together with a least-informative pairwise-KL obstruction this gives a finite-horizon minimax sandwich.

The finite-sample sandwich separates a universal information obstruction from a concrete sufficient certification rule on the declared nuisance-parameter set.

On a positive-persistence linear-quadratic (LQ) parameter rectangle, define $K_{\min} > 0$ as the minimum coefficient multiplying θ^2 in the full-information value over the nuisance rectangle, and let $h_{\text{econ}} > 0$ be the required economic-value hurdle. The control-defined statistical boundary is

$$b_{\text{econ}} = \sqrt{h_{\text{econ}}/K_{\min}},$$

so the certification gap is $d = \theta - b_{\text{econ}}$.

D. Random economic lifetime

Let $H = A_T$ be the end of economic lifetime measured in information time. Consider an admissible deployment pair (S, δ) with $S \leq H$, where S is a $\mathbb{G}$ -stopping time and $\delta \in \{0, 1\}$ is $\mathcal{G}_S$ -measurable, with $\delta = 0$ if the deadline has already occurred. Under the reliable constraints

$$\mathbb{P}_0(\delta = 1) \leq \alpha, \quad \mathbb{P}_1(\delta = 1) \geq 1 - \beta,$$

optional sampling gives the mean-information obstruction

$$\frac{1}{2} \mathbb{E}_1[H] \geq \text{kl}(1 - \beta, \alpha).$$

This condition is only necessary: deadline laws with the same mean can have very different lower tails and therefore very different probabilities of successful certification.

For any fixed certification rule, let τ_{cert} denote its deployment time, with $\tau_{\text{cert}} = +\infty$ on nondeployment. For $h \geq 0$, let $F(h)$ be any available lower bound on $\mathbb{P}_1\{\tau_{\text{cert}} \leq h\}$ and let $S_H(h)$ be any available lower bound on $\mathbb{P}_1\{H > h\}$. Without assuming independence,

$$\mathbb{P}_1(\tau_{\text{cert}} < H) \geq \sup_h [F(h) + S_H(h) - 1]_+.$$

The canonical deterministic-horizon envelope $\pi_\alpha(h) = \Phi(\sqrt{h} - z_{1-\alpha})$ gives the complementary upper bound

$$\mathbb{P}_1(\delta = 1, S < H) \leq \inf_h \{\pi_\alpha(h) + \mathbb{P}_1(H > h)\} \wedge 1.$$

Define the residual alive information value

$$m_s := \mathbb{E}_1[(H - s)_+ \mid \mathcal{G}_s],$$

and let $\mathcal{T}_{\mathbb{G}}(H)$ denote the class of $\mathbb{G}$ -stopping times S satisfying $S \leq H$ almost surely. For the dual representation, use the convexified deployment variable $\varpi \in [0, 1]$, measurable with respect to $\mathcal{G}_S$, and require $\varpi = 0$ on $\{S = H\}$. Define the reliable random-horizon capacity by

$$c_{\alpha, \beta}^H := \sup \left\{ \mathbb{E}_1 \left[\mathbf{1}_{\{S < H\}} (H - S) \varpi \right] : S \in \mathcal{T}_{\mathbb{G}}(H), 0 \leq \varpi \leq 1, \right. \\ \left. \varpi \text{ } \mathcal{G}_S\text{-measurable, } \varpi \mathbf{1}_{\{S=H\}} = 0, \mathbb{E}_0[\varpi] \leq \alpha, \mathbb{E}_1[\varpi] \geq 1 - \beta \right\}.$$

Under the same compactness, continuity, and randomized-stopping assumptions used in the deterministic problem, the same Lagrange argument gives the random-horizon dual representation

$$c_{\alpha, \beta}^H = \inf_{\lambda, \omega \geq 0} \left\{ \lambda \alpha - \omega(1 - \beta) + \sup_{S \in \mathcal{T}_{\mathbb{G}}(H)} \mathbb{E}_1 \left[\mathbf{1}_{\{S < H\}} (m_S + \omega - \lambda e^{-Z_S})_+ \right] \right\}.$$

For deterministic $H = A$, $m_s = A - s$ and the deterministic-horizon stopping obstacle is recovered. A useful closed-form benchmark is available when the information-time lifetime has the same law under both hypotheses, is independent of the likelihood process under both laws, and is exponentially distributed with hazard $\varrho > 0$, $H \sim \text{Exp}(\varrho)$. For $h > 0$, consider the constant likelihood-ratio rule $\tau_h = \inf\{s : Z_s \geq h\}$, deployed only if $\tau_h < H$. Put

$$s_\varrho = \sqrt{1/4 + 2\varrho}, \quad a_1 = s_\varrho - 1/2, \quad a_0 = s_\varrho + 1/2.$$

The Laplace transform of the Brownian first-passage time gives

$$\mathbb{P}_0(\tau_h < H) = e^{-a_0 h}, \quad \mathbb{P}_1(\tau_h < H) = e^{-a_1 h}.$$

Choosing $h = -\log \alpha / a_0$ makes the rule exactly level α , with pre-death power

$$\pi_{\alpha, \varrho}^{\text{exp}} = \alpha^{\kappa_\varrho}, \quad \kappa_\varrho = \frac{a_1}{a_0}.$$

Thus the constant-threshold rule attains the target power whenever $\alpha^{\kappa_\varrho} \geq 1 - \beta$. Solving equality, with

$$r = \frac{\log(1 - \beta)}{\log \alpha},$$

gives the benchmark hazard

$$\varrho_{\text{bench}} = \frac{r}{2(1 - r)^2}$$

and mean lifetime KL $\mathbb{E}_1[H]/2 = (1 - r)^2/r$. For this constant-threshold rule, at $\alpha = 0.05$ and power 0.90, the benchmark mean lifetime is about 26.47 nats, compared with 4.28 nats under a deterministic lifetime. Random deadlines depend on their lower-tail behavior as well as their mean. Early-deadline risk can materially increase the information required for certification.