

The Valentini H-Function as an Objective for Quantitative Portfolio Management

Maksim Aleksandrovich Andreev, Joseph Howden
Quark Quantitative Research

Abstract

Classical portfolio-management objectives—expected utility (von Neumann & Morgenstern, 1944), Sharpe ratio (Sharpe, 1966), Kelly criterion (Kelly, 1956)—are scalar functionals of the return distribution’s first two moments. Peters (2019) showed that under the multiplicative wealth dynamics that actually govern repeated investment, these objectives are ergodic approximations that systematically misrepresent long-run performance whenever the return distribution is non-Gaussian. We propose the Valentini (1991) H -function $H[\rho, |\psi|^2] = \int \rho(x) \cdot \ln(\rho(x) / |\psi(x)|^2) dx$ as a non-equilibrium alternative. H is the Kullback–Leibler divergence of the realised return density ρ from the options-implied risk-neutral density $|\psi|^2$; it is zero in the efficient-markets limit ($\rho = |\psi|^2$) and strictly positive whenever the realised distribution deviates from the implied one. The portfolio objective is $\int_0^T (-dH/dt)dt$: the rate at which an active strategy harvests sub-quantum non-equilibrium before the market relaxes back to equilibrium under Valentini’s H -theorem. The objective generalises the Kelly criterion to arbitrary realised distributions, is dimensionally consistent with the Bohmian-finance framework (Andreev & Howden, 2026a), and admits a calibration channel via the Breeden–Litzenberger (1978) construction. We provide the derivation, the relation to Kelly, the Peters ergodicity gap, and a proposed prospective validation protocol; the default-form extraction-rate hypothesis (a Gaussian reference in place of the options-implied density) was settled null on 2026-07-23 with its planned live wiring retired, so the empirical contribution is re-scoped to the untested options-implied form (§6).

1. Introduction

A quantitative portfolio strategy is a map from market state to position weights. To evaluate such a map we require an objective functional. The classical answers—expected utility, Sharpe ratio, Kelly fraction—differ in their treatment of risk and time-aggregation but share a common form: each is a scalar function of the first two moments of the return distribution. This is a powerful simplification when the distribution is close to Gaussian, but the historical record makes clear that financial return distributions are not close to Gaussian, particularly in the tails (Mandelbrot, 1963; Cont, 2001).

Peters (2019) sharpened the diagnosis: under multiplicative wealth dynamics ($W_{t+1} = W_t(1 + r)$), the ergodic and time-average growth rates differ by a term involving the variance of log-returns. Optimising the expectation (the ensemble average) systematically over-allocates relative to the time-average optimum whenever the distribution has heavy tails or skew. The

Kelly criterion partially repairs this—Kelly is the time-average-optimal allocation under log-utility—but Kelly is itself derived under the assumption that the realised return distribution is known and stationary. In practice it is neither.

Contribution. We propose the Valentini (1991) H -function as a portfolio objective that is (i) defined on the full density rather than scalar moments; (ii) time-asymmetric, and so dimensionally consistent with the irreversibility of investment returns; (iii) reducible to the Kelly criterion in the equilibrium Gaussian limit; (iv) calibrated through the implied risk-neutral density extracted from options; and (v) active even when the realised distribution is non-stationary, because H quantifies precisely the deviation between the realised path-density and the option-implied prior at the current calibration instant.

2. Background: Ergodicity, Kelly, and Quantum Non-Equilibrium

2.1 The ergodicity gap (Peters 2019)

For a multiplicative wealth process $W_{t+1} = W_t \cdot R_t$ with $R_t = 1 + r_t$ and i.i.d. returns r_t , the ensemble-average growth rate is

$$g_{\text{ens}} = \ln E[R_t] = \ln(1 + E[r_t])$$

while the time-average growth rate is

$$g_{\text{time}} = E[\ln R_t] = E[\ln(1 + r_t)]$$

By Jensen's inequality $g_{\text{time}} \leq g_{\text{ens}}$; the gap is $-(1/2)\text{Var}[r] + O(r^3)$ for small returns, and is larger for distributions with heavier tails. An investor maximising g_{ens} over-allocates relative to the optimal long-horizon strategy.

2.2 Kelly criterion as the time-average optimum

The Kelly criterion (Kelly, 1956; Thorp, 2006) addresses this by maximising g_{time} directly:

$$f^* = \arg \max_f E[\ln(1 + f \cdot r)]$$

For a single asset with mean μ and variance σ^2 , the closed-form solution is $f^* = \mu/\sigma^2$. With higher moments accounted for the result generalises to the standardised-skew-and-kurtosis-corrected fraction:

$$f^* \approx (\mu / \sigma^2) \cdot [1 - (1/2) S f - (1/6) K f^2]$$

The Kelly criterion is optimal under three assumptions: (i) log-utility; (ii) known stationary return distribution; (iii) independent rounds. Real markets violate all three. The distribution is unknown (only options give us a forward density); it is non-stationary (parameters drift); and rounds are not independent (correlations cluster in time).




Figure 1. The ergodicity gap on realised SPY daily returns (2000–2024, $N = 6,288$ trading days). For each allocation fraction f (values above 1 denote leverage), the ensemble-average growth rate $\ln E[1 + fr]$ (amber) rises monotonically, so an investor optimising the expectation over-allocates without bound. The time-average growth rate $E[\ln(1 + fr)]$ (blue) is concave and peaks at the Kelly fraction $f^* \approx 2.4$ (dashed line); beyond it, realised compound growth falls even as the expectation keeps rising. The widening vertical gap between the two curves is the Peters (2019) ergodicity penalty that the H-objective is designed to circumvent by optimising a time-average rate directly. Computed from the cached SPY daily-close series.

2.3 Valentini non-equilibrium

In the Bohmian formulation of quantum mechanics (Bohm, 1952), the particle position has a definite value at each instant; the “wavefunction” $|\psi|^2$ describes a distribution over an ensemble of trajectories. Valentini (1991) showed that the assumption $\rho = |\psi|^2$, traditionally an axiom of standard quantum mechanics, is dynamically derivable as the equilibrium limit of a more general framework in which ρ can deviate from $|\psi|^2$. The deviation is quantified by the H -function:

$$H[\rho, |\psi|^2] = \int \rho(x) \cdot \ln(\rho(x) / |\psi(x)|^2) dx$$

Valentini’s H -theorem states that under the guidance equation dynamics, $dH/dt \leq 0$ —the realised density ρ relaxes monotonically to $|\psi|^2$. In the equilibrium limit ($H = 0$) the system is indistinguishable from standard quantum mechanics: probability of a measurement outcome equals $|\psi|^2$. In the non-equilibrium regime ($H > 0$) the system encodes more information about its underlying trajectory than the wavefunction prescribes; sub-quantum signal locality permits, in principle, faster-than-light communication and other phenomena ruled out by standard quantum mechanics.

The financial application. In stochastic-mechanical finance (Andreev & Howden, 2026a), the options market constructs the implied density $|\psi|^2$ via Breeden–Litzenberger; the realised path determines ρ . In an efficient market $\rho = |\psi|^2$: there is no exploitable deviation between the options-implied prior and the realised distribution. In a non-equilibrium market $\rho \neq |\psi|^2$: $H > 0$ quantifies the deviation, and — *conditional on the reference density being the genuinely options-implied $|\psi|^2$* — that deviation may encode an extractable information advantage. Whether any such advantage is in fact extractable is an empirical question, and for the only estimator implemented to date (which substitutes a Gaussian default reference for the options-implied density) the answer was negative: the deviation scalar carried no marginal predictive content (see the settlement disclosure in §6). The options-implied form remains untested.

3. The H-Function as Portfolio Objective

3.1 Definition

Given the realised path-density $\rho_{\text{realised}}(x, t)$ and the options-implied density $|\psi|^2(x, t)$ at calibration time t , the instantaneous H -value is

$$H(t) = \int \rho_{\text{realised}}(x, t) \cdot \ln[\rho_{\text{realised}}(x, t) / |\psi(x, t)|^2] dx$$

and the H -extraction rate is $-dH/dt$. The proposed portfolio objective over horizon T is

$$\begin{aligned} J[\text{strategy}] &= \int_0^T (-dH/dt) dt \\ &= H(0) - H(T) \end{aligned}$$

That is, an effective strategy is one that *extracts* non-equilibrium H over its holding period—converting deviation between realised and implied densities into portfolio gains—rather than passively allowing the H to dissipate through the Valentini relaxation.

3.2 Kelly as the equilibrium-limit special case

To recover the Kelly criterion as a special case, consider the equilibrium reference $|\psi|^2$ equal to the Gaussian density with mean μ and variance σ^2 , and the realised ρ concentrated on the deterministic outcome $r = f \cdot x$ for portfolio fraction f . Direct substitution yields

$$\begin{aligned} H(f) &= -(1/2) \ln(2\pi\sigma^2) - (f \mu - f^2\sigma^2/2) + \\ &\quad O(\text{higher moments}) \end{aligned}$$

Maximising $-H$ over f recovers $f^* = \mu/\sigma^2$, the Kelly fraction. The H -objective reduces to Kelly in the limit where (i) the implied density is Gaussian and (ii) only the second moment of the realised distribution matters. Outside that limit, the H -objective uses the full density and is genuinely more general.

Higher-moment generalisation. The H-functional naturally incorporates skew, kurtosis, and tail features of the realised distribution because it integrates against the log-ratio of two full densities. The corresponding Kelly approximation $f^* \approx (\mu/\sigma^2)[1 - (S/2)f - (K/6)f^2]$ is the second-order Taylor expansion of the H-extraction maximum around the Gaussian point.

3.3 Connection to the second law of thermodynamics

Valentini's H-theorem is the quantum analogue of Boltzmann's classical H-theorem: both state that entropy increases over time. In Boltzmann, H is the negative Shannon entropy of a classical phase-space distribution; in Valentini, H is the Kullback–Leibler divergence between the actual distribution and the equilibrium one prescribed by the wavefunction. In both cases the relaxation is irreversible: positive H decays toward zero.

The financial reading of this relaxation is a theoretical idealisation, and we state it as such. *Under the idealised guidance-equation dynamics*, trading flows push ρ toward the implied density—arbitrageurs trade the mispricings until they vanish—so H would decay monotonically and alpha would be the rate of H-

extraction *before* that relaxation completes. The structural limit on extractable alpha would then be the rate at which fresh non-equilibrium (news, flows, regime changes) is injected into the system.

The empirical relaxation is not established—and the one market-specific test points the other way.

Whether financial markets actually relax

monotonically toward $\rho = |\psi|^2$ is an open empirical question, not a demonstrated rate. The single market-specific test of relaxation dynamics in this research programme measures how the aggregate osmotic velocity u (Nelson, 1966) scales with the symmetry-breaking condensate magnitude $|\Phi|^2$: a companion study (Andreev & Howden, 2026h) finds a *positive* log-log slope $v \approx +0.5$ across four independent specifications ($|t| > 3.0$ in every one; Figure 2). The canonical Nelson/Valentini Goldstone-mode reading predicts $v < 0$ —a stabilising condensate that relaxes—whereas the realised positive sign indicates a *destabilising* condensate in which crowded positioning widens the density gradient rather than dissipating it. The consequence for the H-objective is that regimes of H-growth (widening deviation) are expected, not excluded: the objective harvests non-equilibrium where it exists, but must not lean on a guaranteed monotone relaxation.




Figure 2. The market-specific empirical counterpart to the H-theorem relaxation. Four independent specifications of the log-log slope ν relating the aggregate osmotic velocity to the symmetry-breaking condensate magnitude $|\Phi|^2$ (SPY, QQQ, IWM, TLT, GLD, 2005–2024) all return $\nu \approx +0.5$ with $|t|$ well above the Harvey–Liu–Zhu (2016) threshold of 3.0. The canonical Nelson/Valentini Goldstone-mode reading predicts $\nu < 0$ (the relaxation region, below the axis); the realised sign is positive, indicating a destabilising rather than a stabilising financial condensate. The H-theorem relaxation is therefore a theoretical idealisation, not an established empirical rate. Source: the companion osmotic-velocity scaling study (Andreev & Howden, 2026h).

4. Computational Forms

4.1 Implied-density calibration

The wavefunction-implied density $|\psi|^2$ is obtained per Breeden & Litzenberger (1978) from the option price surface. For a single underlying at horizon T ,

$$|\psi(K, T)|^2 \equiv \rho_Q(K, T) = \exp(rT) \cdot \partial^2 C(K, T) / \partial K^2$$

For multi-asset portfolios the per-asset density is projected through the historical or implied co-cumulant tensor (Baaquie, 2004). Cross-asset implied correlation can be extracted from variance dispersion (Dispersion Trading; Driessen *et al.*, 2009).

4.2 Realised density estimation

The realised density ρ_{realised} is the density of recent realised returns over the calibration window W . Standard methods—kernel density estimation, Gram-Charlier expansion (matching the implied-density

skew and kurtosis), or empirical cumulant expansion—are appropriate. Care is required to align the time-scaling so that ρ and $|\psi|^2$ refer to the same horizon.

4.3 Practical H estimator

Given the two densities, H is computed by numerical integration over a discrete grid:

$$\hat{H}(t) \approx \sum_i \rho_{\text{realised}}(x_i) \cdot \ln[\rho_{\text{realised}}(x_i) / |\psi(x_i)|^2] \cdot \Delta x$$

Bias correction (Beirlant *et al.*, 2001) is required when the empirical density is estimated from a finite sample. For the prospective validation protocol described in §6 we recommend a kernel density estimator with bandwidth chosen by cross-validation and a clip on the ratio inside the logarithm (regularisation to prevent division-by-zero where the implied density is small).

4.4 Rate of extraction

The H-extraction rate over a window $[t, t + \Delta t]$ is the finite difference

$$-dH/dt \approx -[\hat{H}(t + \Delta t) - \hat{H}(t)] / \Delta t$$

For optimisation under the proposed objective the strategy is one that maximises the cumulative extraction over the holding period. In the limit of small Δt this reduces to a local optimisation problem;

in practice we recommend horizon-scale optimisation (daily or weekly Δt) consistent with the implied-density calibration cadence.

5. Relation to Existing Objectives

Table 1 summarises the relation between the H-objective and standard alternatives.

Table 1: H-objective vs classical alternatives

OBJECTIVE	FUNCTIONAL	REDUCES TO H WHEN
Expected utility	$E[u(W)]$	$u = \log$, density = empirical, $ \psi ^2 = \text{uniform}$
Sharpe ratio	$(\mu - r_f) / \sigma$	Gaussian limit + small-deviation expansion
Kelly criterion	$E[\ln(1 + f r)]$	Gaussian $ \psi ^2$ + Taylor to 2nd order
Sortino	$(\mu - r_f) / \sigma_{\text{down}}$	Half-normal $ \psi ^2$ + log-utility
Omega ratio	$\int_{-L}^{\infty} (1 - F) dx / \int_{-\infty}^L F dx$	Threshold-cut H-divergence
Valentini H	$\int \rho \ln(\rho / \psi ^2) dx$	(reference point—all others are limits of it)

The structural advantages of the H-objective are: (a) it is parametrised by both the realised and the implied density, automatically calibrating against an external (options-implied) prior; (b) it incorporates all moments of the realised distribution, not just first two; (c) it is explicitly time-dependent (the implied density evolves with option chains), so non-stationary markets pose no special complication; and (d) it is dimensionless, eliminating the need for a risk-free benchmark.

6. Proposed Empirical Validation

Settlement disclosure (2026-07-23): the default-form extraction-rate hypothesis is settled null, and the planned live wiring is retired. The research programme's registered master hypothesis for this objective — the extraction-rate hypothesis, testing whether the instantaneous H computed by the one estimator implemented to date predicts forward risk-adjusted return — was settled as a well-powered null on 2026-07-23. That estimator substitutes a Gaussian reference fitted to the empirical mean and variance for $|\psi|^2$, so its H is a non-Gaussianity scalar of the rolling return density, not a realised-versus-options-implied divergence. Under the frozen gates: controlled Newey–West $t = +0.814$ (gate > 3.0), univariate $t = -0.068$ (inert even before controls), the permutation gate failed, the decile long-short deflated Sharpe ratio was 0.230, and the sign was unstable across specifications. The instrument itself was validated before market data (it detects a planted signal and collapses a disguised-volatility proxy), so the null is informative rather than scorer-blindness. Consequently the planned live wiring of that estimator was retired rather than deployed. What survives — and what this paper's contribution is hereby re-scoped to — is the *distinct, untested* hypothesis in which the reference density is the genuinely options-implied $|\psi|^2$ of §4.1 (Breedén–Litzenberger). That variant requires its own fresh registration, under its own identifier, before any empirical claim is made; the protocol below is the proposed design for exactly that registration.

Per Harvey, Liu & Zhu (2016), we set out the prospective validation protocol we propose for the options-implied form; running it is contingent on the

fresh registration described above.

Proposed validation protocol (contingent on fresh registration of the options-implied variant).

(P1) Universe: fixed top-30 liquid US equities plus SPY, QQQ, TLT, GLD.

(P2) Implied-density source: option chains with >90 days to expiry, weekly recalibration.

(P3) Realised-density estimator: kernel density on last 252 trading days.

(P4) Allocation rule: maximise expected H-extraction over forward 1-week horizon subject to maximum position size and Markowitz long-only constraints.

(P5) Outcome metric: forward 4-week realised return per unit H-extracted in the calibration window.

Hypothesis: positive correlation, slope > 0 at $t > 3.0$.

(P6) Minimum observation count: 100 weeks of forward-out-of-sample tracking before the t-statistic is computed (per Validation-Limits doctrine; Andreev & Howden 2026e).

(P7) Comparison baselines, declared with the registration: Sharpe-optimal, Kelly-optimal, equal-weight.

Caveat. The H-objective requires reliable estimation of both the realised and the implied density. Both are subject to substantial finite-sample noise. We expect at least 12–18 months of forward tracking before a robust signal-vs-noise determination is possible. Until that protocol is complete, we make no empirical claims about the H-objective relative to Kelly or Sharpe.

7. Discussion

The proposed H-objective is a generalisation of Kelly to (i) arbitrary realised distributions and (ii) options-calibrated implied reference distributions. Its structural advantages over Kelly are pedagogical: H is defined on the full density, not the first two moments; H is time-dependent through the option-chain implied density; and H is dimensionless. Its operational advantages are practical: it gives an automated calibration channel via options prices, and it gives a natural stopping criterion (stop trading the strategy when $H \rightarrow 0$).

The connection to Peters' ergodicity gap is straightforward. Peters showed that ensemble and time-averages differ for non-Gaussian distributions;

the H-objective directly optimises a time-average rate (the H-extraction rate) over the realised path, side-stepping the ergodicity gap entirely. The implied-density calibration further protects against distributional shift: as the option chain evolves, $|\psi|^2$ updates, so the H-objective tracks the market's own forward expectation.

We emphasise the methodological status: this is a derivation paper, and its empirical record is now partly settled. The default-form extraction-rate hypothesis — instantaneous H from the implemented Gaussian-reference estimator predicting forward risk-adjusted return — was settled null on 2026-07-23 and that estimator's planned live wiring was retired (§6). The H-function with a genuinely options-implied reference density remains a candidate objective whose superiority to Sharpe and Kelly is structurally plausible but wholly untested; it is a distinct hypothesis that must receive its own fresh registration before the protocol of §6 is run, and we will report the empirical result only after that registration and its minimum observation count are met.

References

- [1] A. Valentini. Signal-locality, uncertainty, and the subquantum H-theorem. *Physics Letters A*, 158(1–2):1–8, 1991.
- [2] D. Bohm. A suggested interpretation of the quantum theory in terms of “hidden” variables. *Physical Review*, 85(2):166–193, 1952.
- [3] E. Nelson. Derivation of the Schrödinger equation from Newtonian mechanics. *Physical Review*, 150(4):1079–1085, 1966.
- [4] D.T. Breeden & R.H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978.
- [5] O. Peters. The ergodicity problem in economics. *Nature Physics*, 15(12):1216–1221, 2019.
- [6] J.L. Kelly. A new interpretation of information rate. *Bell System Technical Journal*, 35(4):917–926, 1956.
- [7] E.O. Thorp. The Kelly criterion in blackjack, sports betting, and the stock market. *Handbook of Asset and Liability Management*, 1:385–428, 2006.
- [8] W.F. Sharpe. Mutual fund performance. *Journal of Business*, 39(1):119–138, 1966.
- [9] H. Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.
- [10] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.
- [11] B. Mandelbrot. The variation of certain speculative prices. *Journal of Business*, 36(4):394–419, 1963.
- [12] B.E. Baaquie. *Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates*. Cambridge University Press, 2004.
- [13] J. Driessen, P.J. Maenhout & G. Vilkov. The price of correlation risk: evidence from equity options. *Journal of Finance*, 64(3):1377–1406, 2009.
- [14] J. Beirlant *et al.*. Nonparametric entropy estimation: An overview. *International Journal of Mathematical and Statistical Sciences*, 6(1):17–39, 2001.
- [15] C.R. Harvey, Y. Liu & H. Zhu. ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.
- [16] M.A. Andreev & J. Howden. Stochastic mechanical methods for quantitative portfolio management. *Quark Quantitative Research Working Paper*, 2026a.
- [17] M.A. Andreev & J. Howden. Ontological limits of historical validation for wavefunction-calibrated financial systems. *Quark Quantitative Research Working Paper*, 2026e.
- [18] M.A. Andreev & J. Howden. Osmotic-velocity scaling at symmetry-breaking transitions in multi-asset markets. *Quark Quantitative Research Working Paper*, 2026h.

© 2026 Maksim Aleksandrovich Andreev & Joseph Howden.

All rights reserved.

Quark Quantitative Research Platform — Working Paper

Generated 2026-08-25 21:59 UTC

This paper proposes the Valentini H-function as a portfolio objective generalising Kelly to non-equilibrium distributions and resolving the Peters ergodicity gap. Correspondence: M.A. Andreev, Quark Quantitative Research.