SVVVL
/Research

Nature Abhors an Undiversified Bet

An Evolutionary Foundation for Continuous-Time Asset Allocation

J. E. Whitehead III[1]

May 30, 2026

> *"Nature abhors an undiversified bet."*
> — Brennan and Lo (2012)

We extend the evolutionary binary-choice framework of Brennan and Lo [3] to continuous-time asset allocation with one risky and one risk-free asset. The key move is a change in the unit of selection: individual dollars, not investors. Every dollar faces the same systematic Brownian shock, so the Brennan-Lo logic applies directly: the growth-optimal strategy is randomization. Maximizing instantaneous log-portfolio growth yields $\ell^* = \mu_e/\sigma^2 + (1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t$. The deterministic part is the Kelly criterion [7] and, in the continuous-time limit, Merton's rule [9,10]. The stochastic part is the investor's belief about market randomness, characterized by its correlation ρ with the true Brownian innovation, the continuous-time analog of the Brennan-Lo intelligence parameter. Substituting ℓ^* into the log-portfolio dynamics produces an intelligence correction of $\rho - \frac{1}{2} \in [-\frac{3}{2}, +\frac{1}{2}]$: misaligned beliefs destroy three times what aligned beliefs create. The stochastic term also resolves a long-standing puzzle. The instantaneous growth-optimal rule ℓ^* and the long-run rule $\ell^{**} = \mu_e/\sigma^2$ are distinct objects. The variance of the randomization term, $1/(\mathrm{d}t \cdot \sigma^2)$, is the hedge. The expected allocation remains Kelly. The spread around it widens as the horizon shortens. Because no investor lives forever, that spread is always positive.

[1] SVVVL, Towson, MD, United States. Contact: james.whitehead@svvvl.com. See Disclosures at the end of this document.

Introduction

The Kelly criterion [7] and Merton's portfolio rules [9,10] are derived from optimization principles that, while mathematically clean, leave open the question of why evolution would produce investors who behave consistently with them. They tell us what such investors would do, not why selection should favor them.

Brennan and Lo [3] asked a harder question: what does natural selection actually produce? In their binary-choice model, a population of organisms picks between two actions. Each action yields a random number of offspring. The punchline is quiet but consequential. When environmental shocks are systematic, when the whole population is in the same ecological boat, rational deterministic behavior loses. Randomization wins. Probability matching, loss aversion, and other phenomena that look irrational from the individual's perspective are optimal from the population's. Rational behavior emerges only when shocks are idiosyncratic.

This paper carries that logic into continuous-time portfolio allocation. One risky asset, one risk-free asset. The risky asset exposes every investor to the same Brownian shock simultaneously. That is complete systematic correlation, and it appears at first to rule out any deterministic portfolio rule. A single reframing dissolves the difficulty.

We make four contributions.

Contribution 1: The population-of-dollars reframing. At first glance, the Brennan-Lo logic seems to apply to investors. The original framework describes creatures that face a binary choice, and investors are the natural analog. But squint, tilt the head, and the logic admits a different reading. It applies just as cleanly to the dollars inside an investor's portfolio. The investors are no longer center stage; their dollars are. Dollars are the population for the investor, just as investors are the population for the market. With this reframing, each dollar faces the same systematic Brownian shock $\mathrm{d}W_t$, the Brennan-Lo logic for systematic environments applies directly, and the growth-optimal strategy is randomization. Each dollar draws its allocation from a Gaussian distribution, a biased coin flip, and the aggregate Kelly weight emerges from the law of large numbers.

Contribution 2: Kelly and Merton from selection. Maximizing instantaneous log-portfolio growth yields

$$\ell^* \;=\; \frac{\mu_e}{\sigma^2} + \frac{1}{\sigma}\frac{\mathrm{d}W_t}{\mathrm{d}t},$$

where $\mu_e = \mu_s - \mu_b$ is the expected excess return. The deterministic part is the Kelly weight, which is Merton's fraction in the continuous-time limit. Both results fall out of evolutionary selection.

Contribution 3: Intelligence, defined. Write the investor's belief about market randomness as M_t', decomposed as

$$M_t' \;=\; \sqrt{1-\rho^2}\,\widetilde{W}_t' + \rho\, W_t',$$

where $\rho = \mathrm{Corr}[M_t', W_t']$ measures how well the investor's model of market randomness tracks the truth. Substitute ℓ^* into the log-portfolio dynamics. The drift picks up an intelligence correction of $\rho - \frac{1}{2} \in [-\frac{3}{2}, +\frac{1}{2}]$. Bad beliefs destroy three times what good beliefs create.

Contribution 4: The Kelly short-run puzzle, resolved. $\ell^* = \mu_e/\sigma^2 + (1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t$ and $\ell^{**} = \mu_e/\sigma^2$ are different objects. The first maximizes instantaneous stochastic growth rate of the portfolio. The second maximizes long-run average growth rate. The stochastic term is a hedge: the expected dollar allocation is always Kelly, but each dollar draws from a Gaussian centered at Kelly with variance $1/(\mathrm{d}t \cdot \sigma^2)$. The shorter the horizon, the wider the spread, and the greater the hedge required. The pure Kelly weight ℓ^{**} is the limiting case as $\mathrm{d}t \to \infty$, a horizon no investor reaches.

The paper proceeds as follows. Section 2 surveys the literature. Section 3 sets up the model. Section 4 derives the optimal allocation. Section 5 develops the population-of-dollars reframing. Section 6 maps the framework to Brennan-Lo. Section 7 defines intelligence. Section 8 treats bounded rationality. Section 9 resolves the Kelly short-run puzzle. Section 10 discusses implications. Section 11 concludes.

Related Literature

Evolutionary foundations of economic behavior

Brennan and Lo [3] build a binary-choice model in which a population of organisms picks between two actions, each yielding a random number of offspring. The growth-optimal behavior f^* maximizes the geometric growth rate $\alpha(f) = \mathbb{E}[\log(f x_a + (1-f)x_b)]$. The central result is a clean dichotomy: systematic shocks favor randomization; idiosyncratic shocks favor individually-optimal deterministic behavior. Intelligence is $\rho = \mathrm{Corr}[I_{it}^f, (x_a - x_b)]$: how well the organism's choice tracks which action nature rewards. When intelligence carries a cost $c(\rho)$, bounded rationality is the outcome. The earlier model is in Brennan and Lo [2]. The adaptive coin-flipping interpretation of random individual variation under natural selection originates with Cooper and Kaplan [4], whose decision-theoretic framework provides the conceptual foundation for the randomization results developed here. Related work on the natural selection of utility functions [12,13,14,15] and the evolution of trading strategies [5] informs the broader context.

The Kelly criterion and growth-optimal portfolios

Kelly [7] found the growth-optimal betting strategy for an investor with an edge. Breiman [1] proved it wins almost surely in the long run. Thorp [16] and Cover [6] extended the result to portfolios. Every paper in this line assumes growth-optimality as the goal. We derive it.

Merton's continuous-time framework

Merton [9,10] solved the dynamic portfolio problem in continuous time under general utility. Under log utility, the solution collapses to a clean myopic rule identical to Kelly in the continuous-time limit. We supply the evolutionary argument for why the market produces investors whose behavior is consistent with this solution.

Ergodicity economics

Peters [11] makes the case that time-average growth, not ensemble-average growth, is the right objective. The time-average growth rate of a leveraged portfolio is

$$\bar{g}(\ell) \;=\; \mu_b + \ell\mu_e - \frac{\ell^2\sigma^2}{2},$$

with optimal leverage $\ell_{\text{opt}} = \mu_e/\sigma^2$, derived from the time-average growth rate directly. We embed this result in the Brennan-Lo evolutionary framework. The variance drag $\ell^2\sigma^2/2$ is the continuous-time face of the same evolutionary penalty that makes over-commitment fatal in Brennan-Lo.

The Adaptive Markets Hypothesis

Lo [8] reads financial markets as adaptive systems in which strategies compete and evolve. The intelligence parameter ρ and its cost $c(\rho)$ developed here give that intuition a quantitative backbone: market efficiency is time-varying because the cost-benefit calculus of belief accuracy shifts with market conditions.

The Continuous-Time Portfolio Model

Asset dynamics

Two assets. The risky asset S follows a geometric Brownian motion,

$$\frac{\mathrm{d}S}{S} \;=\; \mu_s\,\mathrm{d}t + \sigma\,\mathrm{d}W_t, \tag{1}$$

where $\mu_s > 0$ is the drift, $\sigma > 0$ is the volatility, and W_t is a standard Brownian motion[2] on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, \mathbb{P})$. The risk-free asset B grows at a known rate,

$$\frac{\mathrm{d}B}{B} \;=\; \mu_b\,\mathrm{d}t. \tag{2}$$

The excess drift is

$$\mu_e \;=\; \mu_s - \mu_b \;>\; 0, \tag{3}$$

and we assume the risky asset earns a positive premium throughout.

[2] Throughout the paper, we adopt continuous-time notation for derivational convenience. The formal expression $\mathrm{d}W_t/\mathrm{d}t$ that appears later in the derivation is not a classical derivative; W_t is nowhere differentiable in the ordinary sense. We use it as shorthand for the finite increment $\Delta W_t/\Delta t \sim \mathcal{N}(0, 1/\Delta t)$ at a small but strictly positive rebalancing interval $\Delta t > 0$. The continuous-time framing serves as the idealized limit useful for derivation; the result is meant to be applied at finite Δt. This is consistent with the long-standing practice in physics, control theory, and quantitative finance of treating white noise as the formal derivative of a Wiener process, with all expressions interpreted at a finite sampling frequency when implemented. It is also consistent with the argument developed in Section 9: $\mathrm{d}t \to \infty$ is never reached in practice, and neither is $\mathrm{d}t \to 0$.

Portfolio dynamics

Let $\ell \in \mathbb{R}$ be the portfolio weight on the risky asset. We allow ℓ to range over the real line. Short positions ($\ell < 0$), full investment ($\ell = 1$), and leverage ($\ell > 1$) are all permitted. This generality is required: the optimal allocation ℓ^* derived in Section 4 is drawn from a distribution supported on $\mathbb{R}$. The portfolio return is

$$\frac{\mathrm{d}P}{P} \;=\; \ell\,\frac{\mathrm{d}S}{S} + (1-\ell)\,\frac{\mathrm{d}B}{B}. \tag{4}$$

Substituting (1) and (2),

$$\mathrm{d}P \;=\; P\big(\mu_b + \mu_e\,\ell\big)\mathrm{d}t + P\,\ell\,\sigma\,\mathrm{d}W_t. \tag{5}$$

Itô's lemma applied to $\ln P(t)$ gives

$$\mathrm{d}\ln P(t) \;=\; \Big(\mu_b + \mu_e\,\ell - \frac{\sigma^2}{2}\,\ell^2\Big)\mathrm{d}t + \sigma\,\ell\,\mathrm{d}W_t. \tag{6}$$

The term $\sigma^2\ell^2/2$ is the Itô variance drag. It is the continuous-time price of holding the risky asset, and it cannot be avoided. Integrating (6),

$$\ln P(t) \;=\; \ln P(0) + \Big(\mu_b + \mu_e\,\ell - \frac{\sigma^2}{2}\,\ell^2\Big)t + \sigma\,\ell\,W_t, \tag{7}$$

so the time-average growth rate is

$$\bar{g}(\ell) \;=\; \mu_b + \mu_e\,\ell - \frac{\sigma^2}{2}\,\ell^2. \tag{8}$$

A concave quadratic in ℓ. Peters [11] arrives at the same expression from ergodicity arguments.

The Optimal Allocation Rule

Maximizing instantaneous log-portfolio growth

Differentiate $\mathrm{d}\ln P(t)$ in (6) with respect to ℓ and set to zero,

$$\frac{\partial\,\mathrm{d}\ln P}{\partial\,\ell} \;=\; \mu_e\,\mathrm{d}t - \sigma^2\,\ell\,\mathrm{d}t + \sigma\,\mathrm{d}W_t \;=\; 0. \tag{9}$$

Solving for ℓ,

$$\boxed{\ell^* \;=\; \frac{\mu_e}{\sigma^2} + \frac{1}{\sigma}\,\frac{\mathrm{d}W_t}{\mathrm{d}t}.} \tag{10}$$

Two terms. The first, μ_e/σ^2, is the Kelly weight: excess return over variance. The second, $(1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t$, is stochastic. The optimal allocation moves with the Brownian innovation at every instant. It is a process, not a number.

The belief decomposition

The stochastic term is the investor's response to market randomness. Call it her belief. Define

$$M'_t \;\equiv\; \frac{\mathrm{d}W_t}{\mathrm{d}t}, \tag{11}$$

the investor's view of the market's instantaneous innovation. A perfectly informed investor sets $M'_t = W'_t$, where $W'_t \equiv \mathrm{d}W_t/\mathrm{d}t$. Most investors do not. We write

$$M'_t \;=\; \sqrt{1-\rho^2}\,\widetilde{W}'_t + \rho\, W'_t, \tag{12}$$

where $\widetilde{W}'_t \sim \mathcal{N}(0, 1/\mathrm{d}t)$ may be regarded as the investor's view on the non-systematic part of the risky asset's return $W'_t \sim \mathcal{N}(0, 1/\mathrm{d}t)$. The correlation

$$\rho \;=\; \mathrm{Corr}[M'_t, W'_t] \;\in\; [-1, 1] \tag{13}$$

measures how well the investor reads the market. The optimal allocation becomes

$$\ell^* \;=\; \frac{\mu_e}{\sigma^2} + \frac{1}{\sigma}\, M'_t. \tag{14}$$

The Population of Dollars

Why investors are the wrong population

With one risky asset, every investor sees the same $\mathrm{d}W_t$ at the same time. That is complete systematic correlation. The Brennan-Lo result [3] is unambiguous: under systematic shocks, no fixed deterministic allocation survives. A bad enough realization of $\mathrm{d}W_t$ will eventually ruin any investor holding a fixed ℓ. The answer is to find the right population, not to escape the systematic risk.

The Gaussian coin flip

The population is individual dollars, not investors. Each dollar is the unit of selection. Every dollar faces the same systematic $\mathrm{d}W_t$. The Brennan-Lo logic applies to the dollar population directly, and the conclusion is the same as it always is under systematic risk: randomize. Each dollar i draws its allocation independently,

$$\ell^*_i \;\sim\; \mathcal{N}\!\left(\frac{\mu_e}{\sigma^2},\; \frac{1}{\mathrm{d}t\cdot\sigma^2}\right). \tag{15}$$

The mean is the Kelly weight μ_e/σ^2; the variance is $1/(\mathrm{d}t \cdot \sigma^2)$, the Gaussian coin flip. Probability matching in Brennan-Lo and this coin flip are the same thing, expressed in different notation. Both are instances of the adaptive coin-flipping principle of Cooper and Kaplan [4].

The aggregate result

The draws are independent across dollars. The law of large numbers takes over,

$$\mathbb{E}[\ell_i^*] \;=\; \frac{\mu_e}{\sigma^2} \;=\; \ell^{**}. \tag{16}$$

The Kelly weight appears at the portfolio level as the aggregate of optimal dollar-level randomization. In Brennan-Lo, probability matching at the individual level produces the fastest-growing population. Here, Gaussian randomization at the dollar level produces the Kelly portfolio. Same logic, continuous time.

Mapping to the Brennan-Lo Framework

The correspondence between the continuous-time framework and the Brennan-Lo binary-choice model is exact. Table 1 lays it out.

Table 1: Mapping between the Brennan-Lo framework and the continuous-time portfolio model.

Brennan-Lo (2012)	**This Paper**
Individual organism	Individual dollar
Population of organisms	Portfolio of dollars
Behavioral phenotype $f \in [0, 1]$	Portfolio weight $\ell \in \mathbb{R}$
Offspring count x_a, x_b	Gross portfolio return
Geometric growth rate $\alpha(f)$	Log-portfolio growth $\mathrm{d}\ln P(t)$
Systematic shock (single niche)	Shared Brownian increment $\mathrm{d}W_t$
Randomization under systematic risk	Gaussian coin flip (15)
Growth-optimal phenotype f^*	Kelly weight $\ell^{**} = \mu_e/\sigma^2$
Intelligence $\rho = \mathrm{Corr}[I_{it}^f, (x_a - x_b)]$	$\rho = \mathrm{Corr}[M_t', W_t']$
Volatility drag (implicit)	Itô drag $\sigma^2\ell^2/2$

The key correspondence: $\alpha(f)$ in Brennan-Lo and $\mathrm{d}\ln P(t)$ here are both the objects selection maximizes. The Itô variance drag $\sigma^2\ell^2/2$ is the continuous-time face of the same penalty that makes overcommitment lethal in Brennan-Lo. Nature penalizes the undiversified bet. In Brennan-Lo the penalty is extinction. Here it is the volatility drag, smooth and analytic, but no less real.

Intelligence: A Continuous-Time Definition

The Brennan-Lo definition

Brennan and Lo [3] define intelligence as the correlation between an organism's choice and the difference in reproductive outcomes,

$$\rho \;=\; \mathrm{Corr}\big[I^f_{it},\; (x_a - x_b)\big].$$

Populations with $\rho > 0$ grow faster. Populations with $\rho < 0$ do not survive. Intelligence, in this framework, is alignment between behavior and what nature rewards.

The continuous-time analog

Here, intelligence is the correlation between the investor's belief process and the true market innovation,

$$\rho \;=\; \mathrm{Corr}[M'_t, W'_t].$$

An investor is intelligent to the degree that her model of market randomness tracks the truth. The definition carries the precise evolutionary meaning of Brennan-Lo, behavior correlated with good outcomes, expressed in continuous time.

The optimal log-portfolio dynamics under ρ

Substitute ℓ^* from (14) into (6) using the decomposition (12),

$$\mathrm{d}\ln P(t) \;=\; \left(\mu_b + \frac{1}{2}\left(\frac{\mu_e}{\sigma}\right)^2\right)\mathrm{d}t + \left(\rho - \frac{1}{2}\right) + \frac{\mu_e}{\sigma}\,\mathrm{d}W_t. \tag{17}$$

Three terms. The first scales with time: the risk-free rate and the squared Sharpe ratio. The second is the intelligence correction $\rho - \frac{1}{2}$. The third is the diffusion. Algebraically, the correction enters as a level shift outside $\mathrm{d}t$. We will return to the question of whether ρ itself can be treated as independent of $\mathrm{d}t$ shortly; physically, it cannot.

A conjecture: ρ as a function of the clock

Treating ρ as a free parameter on $[-1, 1]$, independent of $\mathrm{d}t$, is mathematically clean but physically suspect. Any real strategy that converts information into edge, into positive ρ, faces hard limits. There is only so much information any agent can process per unit time, and only so much of that information that translates into alignment with W'_t. As the horizon $\mathrm{d}t$ shrinks, the maximum achievable ρ should shrink with it. The shorter the window, the less time there is to gather, process, and act on signal.

We therefore conjecture an asymmetric, horizon-dependent bound on ρ.

Conjecture 1 (Horizon-dependent bounds on intelligence)**.** *The achievable correlation ρ between the in-*

vestor's belief process M_t' and the true innovation W_t' is bounded by

$$\rho_{\min} \;\leq\; \rho \;\leq\; \rho_{\max}(\mathrm{d}t),$$

where

$$\rho_{\min} \;=\; -1 \quad \textit{(constant)}, \qquad \rho_{\max}(\mathrm{d}t) \;\in\; [0,1] \textit{ is monotonically increasing in } \mathrm{d}t.$$

The asymmetry is the point. Building positive ρ requires information, processing, and execution, all rate-limited. Negative ρ requires none of these. An investor can achieve perfect misalignment instantaneously, by holding beliefs that are exactly backward. Stupid knows no speed limit.

The conjecture sharpens (17) into a horizon-dependent statement. Substituting $\rho_{\max}(\mathrm{d}t)$ for the upper end of the range, the achievable intelligence correction runs from $-\frac{3}{2}$ up to $\rho_{\max}(\mathrm{d}t) - \frac{1}{2}$, which itself shrinks toward $-\frac{1}{2}$ as $\mathrm{d}t \to 0$. At short horizons, the upside from intelligence narrows; the downside does not. We treat (17) as the reduced-form expression and Conjecture 1 as the structural constraint that must eventually be respected when implementing the framework against real markets.

The full range of $\rho \in [-1, 1]$

Setting the conjecture aside for the moment and reading ρ purely as a parameter, the intelligence correction $\rho - \frac{1}{2}$ runs from $-\frac{3}{2}$ to $+\frac{1}{2}$. Table 2 gives the full picture. The reader should keep Conjecture 1 in mind: the upper end of the table is not unconditionally available at every horizon.

Table 2: Intelligence correction $\rho - \frac{1}{2}$ for $\rho \in [-1, 1]$.

ρ	M_t'	Correction	Interpretation
1	W_t'	$+\frac{1}{2}$	Perfect alignment; maximum growth premium
$\left(\frac{1}{2}, 1\right)$	Partial signal	$\left(0, +\frac{1}{2}\right)$	Bounded intelligence; partial signal and noise
$\frac{1}{2}$		0	Signal exactly offsets noise; zero net correction
$\left(-1, \frac{1}{2}\right)$	Perverse	$\left(-\frac{3}{2}, 0\right)$	Perverse beliefs; investor leans against true innovation
-1	$-W_t'$	$-\frac{3}{2}$	Perfect misalignment; maximum growth destruction

Remark 1 (The asymmetry of intelligence)**.** *The interval $[-\frac{3}{2}, +\frac{1}{2}]$ is not symmetric around zero, and the asymmetry is worth sitting with. A perfectly aligned investor ($\rho = 1$) gains $+\frac{1}{2}$. A perfectly contra-indicator investor ($\rho = -1$) loses $\frac{3}{2}$. Bad beliefs are three times as costly as good beliefs are valuable. The continuous-time framework quantifies what Brennan-Lo states qualitatively: subpopulations with $\rho < 0$ do not survive. Here we see exactly why. The perfectly misaligned investor allocates against every market movement, compounding her losses at each instant. Extinction pressure, made precise.*

Bounded Rationality

The cost of intelligence

In Brennan and Lo [3], intelligence carries a biological cost that caps the evolutionarily stable ρ. In markets, accurate beliefs cost money: data, research, time, and attention. Let $c(\rho)$ be the cost per unit time of achieving belief accuracy ρ. The net growth rate of the dollar population is

$$\bar{g}_{\mathrm{net}}(\rho) \;=\; \mu_b + \frac{1}{2}\Big(\frac{\mu_e}{\sigma}\Big)^2 + \frac{\rho - \frac{1}{2}}{\mathrm{d}t} - c(\rho). \tag{18}$$

The evolutionarily stable ρ^* maximizes $\bar{g}_{\mathrm{net}}(\rho)$. When $c(\rho)$ is increasing and convex, the solution sits in the interior: $\rho^* < 1$. Bounded rationality is the outcome, not an assumption.

The evolutionarily stable portfolio

Conjecture 1 introduces a horizon dependence: ρ is properly written $\rho(\mathrm{d}t)$, with feasibility set $[-1, \rho_{\max}(\mathrm{d}t)]$. The investor chooses $\rho(\mathrm{d}t)$ to maximize $\bar{g}_{\mathrm{net}}(\rho)$ subject to that feasibility constraint. The first-order condition becomes

$$c'\big(\rho^*(\mathrm{d}t)\big) \;=\; \frac{1}{\mathrm{d}t}, \qquad \rho^*(\mathrm{d}t) \;\le\; \rho_{\max}(\mathrm{d}t). \tag{19}$$

The interior solution $c'(\rho^*) = 1/\mathrm{d}t$ holds whenever the unconstrained optimum lies below the feasibility ceiling. When it does not, when the marginal benefit $1/\mathrm{d}t$ would call for a ρ beyond what the horizon allows, the solution is the corner $\rho^*(\mathrm{d}t) = \rho_{\max}(\mathrm{d}t)$, and the investor is intelligence-bound by physics rather than by cost.

Investors with low intelligence costs (quantitative funds, well-resourced institutions) reach $\rho^*(\mathrm{d}t)$ near the upper boundary and hold portfolios close to Kelly. Investors with high intelligence costs (retail investors, infrequent rebalancers) land at lower $\rho^*(\mathrm{d}t)$ and drift from Kelly in predictable directions.

Coexistence of strategies

Different investors are good at different horizons. A high-frequency strategy and a value strategy do not compete on the same axis. Each carries an implicit best operating point in $\mathrm{d}t$. We capture this by writing each investor's intelligence as a ρ-$\mathrm{d}t$ *curve*: a function $\rho_j(\mathrm{d}t)$ that gives investor j's achievable correlation as a function of horizon.

The curves have characteristic shapes. The HFT investor's $\rho_{\mathrm{HFT}}(\mathrm{d}t)$ peaks at small $\mathrm{d}t$ and decays as $\mathrm{d}t$ grows. Her edge dissolves at long horizons. The value investor's $\rho_{\text{value}}(\mathrm{d}t)$ does the opposite, climbing with $\mathrm{d}t$ as fundamentals reassert themselves over noise. Each investor's curve traces the horizons where her strategy translates information into alignment, and the horizons where it does not.

The cross-section of investors generates a population-average curve $\bar{\rho}(\mathrm{d}t) = \mathbb{E}_j[\rho_j(\mathrm{d}t)]$, which traces the typical level of intelligence required to be successful at horizon $\mathrm{d}t$. The relationship between an individual investor's $\rho_j(\mathrm{d}t)$ and the population average $\bar{\rho}(\mathrm{d}t)$ at her chosen horizon is what determines whether she earns the intelligence premium or pays the penalty. Investors do not compete with one another for a single ρ^*. They compete within the horizon they have selected, against the average investor at that horizon.

When intelligence costs vary across investors, different investors will optimally choose different levels of ρ^* and will therefore achieve different net growth rates $\bar{g}_{\text{net}}(\rho^*)$. The portfolio of an investor who has a lower cost of intelligence, achieving ρ^* close to her ceiling, will in general grow faster than the portfolio of an investor who has a high cost of intelligence, constrained to a lower ρ^*. The result is heterogeneous optima, each investor growing at the rate her cost structure and her ρ-$\mathrm{d}t$ curve permit.

Each investor's ρ^* is the best achievable given her cost structure and her chosen horizon. No investor can unilaterally improve her net growth rate by changing her strategy without also changing her cost structure. The distribution of ρ^* values across investors therefore reflects the joint distribution of intelligence costs and horizon-specific capabilities in the population, and these distributions persist as long as their heterogeneity persists.

This provides a measured evolutionary interpretation of the empirical coexistence of passive indexers, active managers, and quantitative strategies in financial markets. Each strategy reflects a different point on the cost-benefit frontier of intelligence and a different shape of $\rho_j(\mathrm{d}t)$. Higher-cost strategies are suboptimal only relative to a lower cost structure that may not be available to the investor pursuing them. The mix of strategies observed in the market reflects the underlying distribution of intelligence costs and horizon-specific edges across investors, with no single strategy positioned to drive the others to extinction.

The Investment Horizon and the Short-Run Conundrum

The long-standing conundrum

Kelly dominates in the long run [1]. In the short run, Kelly can underperform and does so with considerable volatility. The profession has generally treated this as a fact of life. The framework here gives it a precise resolution.

Two distinct optimal allocations

There are two optimal allocations, not one:

$$\ell^* \;=\; \frac{\mu_e}{\sigma^2} + \frac{1}{\sigma}\,\frac{\mathrm{d}W_t}{\mathrm{d}t}, \tag{20}$$

$$\ell^{**} \;=\; \frac{\mu_e}{\sigma^2}. \tag{21}$$

ℓ^* maximizes instantaneous stochastic growth. ℓ^{**} maximizes the ensemble growth rate and the long-run time-average. They are related by

$$\lim_{\mathrm{d}t\to\infty} \ell^* \;=\; \mathbb{E}[\ell^*] \;=\; \ell^{**}. \tag{22}$$

The equality is precise and carries two distinct meanings. The first limit eliminates uncertainty by letting time run out: as $\mathrm{d}t \to \infty$, the stochastic term $(1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t \sim \mathcal{N}(0, 1/\mathrm{d}t)$ vanishes, and the instantaneous growth-optimal rule converges to ℓ^{**} pathwise. The second equality eliminates uncertainty by aggregation: taking the expectation over an infinite number of dollar draws, the law of large numbers collapses the distribution to its mean, which is also ℓ^{**}. Two routes, same destination. The long-run time-average and the ensemble average coincide, and both equal the Kelly weight, a result that Peters [11] identifies as the central insight of ergodicity economics. The conundrum arises because the literature applies ℓ^{**} where ℓ^{*} belongs: in the short run, where the stochastic term is large and neither limiting procedure has had a chance to operate.

The randomization term as a hedge

The variance of the stochastic component of ℓ^* is

$$\mathrm{Var}\left[\frac{1}{\sigma}\,\frac{\mathrm{d}W_t}{\mathrm{d}t}\right] \;=\; \frac{1}{\mathrm{d}t \cdot \sigma^2}. \tag{23}$$

The stochastic term is a hedge. Each dollar draws from a Gaussian centered at the Kelly weight, so the expected allocation is always ℓ^{**}, but the spread around that center is set by the horizon. A shorter horizon means a wider Gaussian, a larger hedge, more variability in how individual dollars are deployed.

Proposition 1 (Horizon-dependent hedging)**.** *The optimal spread of the Gaussian coin flip* (15)*, measured by its variance, is monotonically decreasing in the investment horizon* $\mathrm{d}t$*:*

- *Short horizon* ($\mathrm{d}t$ *small*)*: variance* $1/(\mathrm{d}t \cdot \sigma^2)$ *is large; each dollar hedges aggressively around Kelly; the dollar population fans out widely across allocation values.*
- *Long horizon* ($\mathrm{d}t$ *large*)*: variance shrinks; the hedge tightens; dollar allocations cluster near* ℓ^{**}*.*
- *Infinite horizon* ($\mathrm{d}t \to \infty$)*: variance* $\to 0$*; no hedge required; every dollar allocates exactly at* ℓ^{**}*.*

In all cases, $\mathbb{E}[\ell_i^*] = \ell^{**}$*. The hedge does not shift the center; it widens the spread.*

Because no investor reaches infinite horizon, the hedge is always positive. The horizon $\mathrm{d}t$ is itself a preference parameter, an investor's choice governed by goals, mandate, and life stage. But once it is fixed, the framework determines the magnitude of the hedge as a function of $\mathrm{d}t$ and σ. The decomposition is sharp: preference enters only through the choice of horizon; the optimal hedge for that horizon follows mechanically.

Connection to the Brennan-Lo framework

In Brennan-Lo, f^* is the evolutionary imperative under systematic risk. The term $(1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t$ plays that same role: the growth-optimal response to systematic $\mathrm{d}W_t$, calibrated to the investor's horizon. The Gaussian coin flip $\mathcal{N}(\mu_e/\sigma^2, 1/(\mathrm{d}t \cdot \sigma^2))$ is simultaneously the evolutionary randomization mechanism of Section 5 and the optimal hedge for any investor with a finite horizon. Longer horizons tighten the hedge; shorter horizons widen it. The evolutionary framework and the financial framework agree on exactly how wide.

Discussion

The normative and the evolutionary

Kelly and Merton describe what evolutionary selection produces at the dollar level, in a market with one risky asset and complete systematic correlation. The investor that survives is the one who can avoid injecting perverse intelligence into their allocation.

The population-of-dollars reframing as a general method

The reframing from investors to dollars may travel beyond the single-asset problem. Many questions in financial economics involve systematic shocks that stall the evolutionary analysis at the agent level. Dropping to a sub-agent population (dollars, trades, individual decisions) may open those problems up. The connection to the Brennan-Lo framework is available whenever the right population is found.

Implications for the Adaptive Markets Hypothesis

The bounded rationality result in Section 8 puts numbers on the Adaptive Markets Hypothesis [8]. When excess returns μ_e are high, the payoff to accurate beliefs rises, more investors pay $c(\rho)$ to push ρ^* toward 1, and the market grows more efficient. When μ_e falls or σ climbs, the calculus reverses, investors settle for lower ρ^*, and efficiency erodes. The intelligence correction $\rho - \frac{1}{2}$ in (17) tracks this in real time.

Limitations and future work

Five extensions are immediate. First, multiple risky assets: a vector of weights and a full covariance matrix will enrich both the intelligence mapping and the hedging results. Second, time-varying parameters: letting μ_e and σ move connects the framework to the return predictability literature. Third, price impact: if investor behavior feeds back into asset prices, strategic interactions emerge that the present framework sets aside.

Fourth, leverage constraints: the framework permits $\ell \in \mathbb{R}$, which lets each dollar's allocation ℓ_i^* take values from $-\infty$ to $+\infty$. No real investor can use arbitrary amounts of leverage. Future work would impose realistic bounds on ℓ and examine how those bounds reshape the optimal hedge. In particular, how the truncation of the Gaussian coin flip distorts the convergence to the Kelly weight at finite horizons.

Fifth, the granularity of the dollar population: the framework treats each dollar as an independent allocation decision. In practice, the cost to enter a long or short position is almost always at least two orders of magnitude greater than a single dollar, which makes per-dollar allocation infeasible. Future work would examine how to implement this strategy under a minimum-position-size constraint, in which the effective unit of selection is not a dollar but a block of dollars large enough to clear transaction costs.

The interesting question is how the size of that block, a function of trading costs, account size, and rebalancing frequency, governs how closely the implementable hedge can approximate the theoretical one.

Each of these extensions is a worthwhile next step.

Conclusion

We extended the Brennan-Lo binary-choice framework to continuous-time portfolio allocation with one risky and one risk-free asset. Four results follow.

Contribution 1. For the individual investor, the right population is dollars, not investors. Each dollar faces the same systematic $\mathrm{d}W_t$, so the Brennan-Lo logic applies directly: randomize. The Gaussian coin flip $\ell_i^* \sim \mathcal{N}(\mu_e/\sigma^2, 1/(\mathrm{d}t \cdot \sigma^2))$ is the continuous-time expression of that imperative.

Contribution 2. Kelly and Merton fall out of evolutionary selection. The optimal allocation $\ell^* = \mu_e/\sigma^2 + (1/\sigma)\,\mathrm{d}W_t/\mathrm{d}t$ recovers the Kelly weight in its deterministic part and encodes the investor's beliefs about market randomness in its stochastic part.

Contribution 3. Intelligence is $\rho = \mathrm{Corr}[M_t', W_t']$. The intelligence correction $\rho - \frac{1}{2} \in [-\frac{3}{2}, +\frac{1}{2}]$ is asymmetric. Misaligned beliefs destroy three times what aligned beliefs build. The asymmetry is exact, and it is not gentle.

Contribution 4. ℓ^* and $\ell^{**} = \mu_e/\sigma^2$ are different objects. The stochastic term is a hedge that each dollar executes by drawing from a Gaussian centered at the Kelly weight, with variance $1/(\mathrm{d}t{\cdot}\sigma^2)$. The expected allocation remains Kelly regardless of horizon, but the spread around it depends on $\mathrm{d}t$. Investors with longer horizons hold tighter distributions; investors with shorter horizons hold wider ones. The choice of horizon belongs to the investor, but once that choice is made, the framework specifies the hedge exactly.

Further Reading and References

[1] Breiman, L. (1961). "Optimal gambling systems for favorable games." *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1, 65–78.

[2] Brennan, T. J. and Lo, A. W. (2011). "The origin of behavior." *Quarterly Journal of Finance*, 1(1), 55–108.

[3] Brennan, T. J. and Lo, A. W. (2012). "An evolutionary model of bounded rationality and intelligence." *PLoS ONE*, 7(11), e50310.

[4] Cooper, W. S. and Kaplan, R. H. (1982). "Adaptive 'coin-flipping': a decision-theoretic examination of natural selection for random individual variation." *Journal of Theoretical Biology*, 94(1), 135–151.

[5] Farmer, J. D. (2002). "Market force, ecology and evolution." *Industrial and Corporate Change*, 11(5), 895–953.

[6] Cover, T. M. (1984). "An algorithm for maximizing expected log investment return." *IEEE Transactions on Information Theory*, 30(2), 369–373.

[7] Kelly, J. L. (1956). "A new interpretation of information rate." *Bell System Technical Journal*, 35(4), 917–926.

[8] Lo, A. W. (2004). "The adaptive markets hypothesis: market efficiency from an evolutionary perspective." *Journal of Portfolio Management*, 30(5), 15–29. See also: Lo, A. W. (2017). *Adaptive Markets: Financial Evolution at the Speed of Thought*. Princeton University Press.

[9] Merton, R. C. (1969). "Lifetime portfolio selection under uncertainty: the continuous-time case." *Review of Economics and Statistics*, 51(3), 247–257.

[10] Merton, R. C. (1971). "Optimum consumption and portfolio rules in a continuous-time model." *Journal of Economic Theory*, 3(4), 373–413.

[11] Peters, O. (2019). *Ergodicity Economics*. London Mathematical Laboratory.

[12] Robson, A. J. (2001). "The biological basis of economic behavior." *Journal of Economic Literature*, 39(1), 11–33.

[13] Robson, A. J. (2001). "Why would nature give individuals utility functions?" *Journal of Political Economy*, 109(4), 900–914.

[14] Rogers, A. R. (1994). "Evolution of time preference by natural selection." *American Economic Review*, 84(3), 460–481.

[15] Samuelson, L. (2001). "Introduction to the evolution of preferences." *Journal of Economic Theory*, 97(2), 225–230.

[16] Thorp, E. O. (1971). "Portfolio choice and the Kelly criterion." *Proceedings of the Business and Economics Section of the American Statistical Association*, 215–224.

Disclosures

The views expressed in this document are those of the author and SVVVL alone and do not constitute investment, legal, tax, or accounting advice. This document is provided for informational and research purposes only and is not an offer or solicitation to buy or sell any security or financial product.

The framework presented is theoretical and relies on idealized assumptions that do not hold in practice. The results should not be implemented in a real portfolio without additional constraints and considerations not addressed in this paper.

While reasonable care has been taken in the preparation of this document, no representation or warranty, express or implied, is made as to its accuracy or completeness.

Reading this document does not create an advisory, fiduciary, or client relationship of any kind between the reader and SVVVL or the author.