Volatility Drag and the Perpetual Borrowing Option (Spiral Theory)

Thomas J. Anderson

Working Paper

Last updated: November 5, 2025

Realized geometric returns systematically fall short of arithmetic expectations due to volatility drag. Under canonical Modern Portfolio Theory (MPT) assumptions, this paper shows that an episodic borrowing policy can neutralize that drag *in expectation, over time*. The rule is entirely ex post: when realized wealth falls below the expected compounding path, borrow and restore; when it exceeds, deleverage and build up cash. Formally, when $W_t < W_t^*$, borrow $\Delta L_t = W_t^* - W_t$; when $W_t > W_t^*$, repay or reserve $\Delta L_t = W_t - W_t^*$. The adjustment restores expected exposure without any forecasting component. Monte Carlo simulation ($\mu = 9\%$, $\sigma = 15\%$, 10,000 trials) yields near-perfect restoration of expected compounding, closing the gap between expected and realized growth.

The finding reframes debt not as speculative leverage but as a stabilizer of compounding efficiency. Financing flexibility offsets volatility drag, aligning realized and expected returns under frictionless conditions and unifying MPT with Modigliani–Miller *through time.* Unlike dynamic-leverage, Kelly, or volatility-targeting strategies, the mechanism activates only after the fact, transforming the apparent "wrong time to borrow" into the only time that restores the expected path.

Viewed practically, this mechanism—termed *Spiral Theory*—bridges personal and corporate finance. Every investor holds an implicit option to borrow against their portfolio, and exercising it symmetrically converts volatility drag into expected value. While spreads, collateral, and taxes bound the effect, the principle remains: disciplined, ex-post financing decisions can restore compounding efficiency and open a path toward *outcome-based balance-sheet design* for individuals, pensions, and endowments.

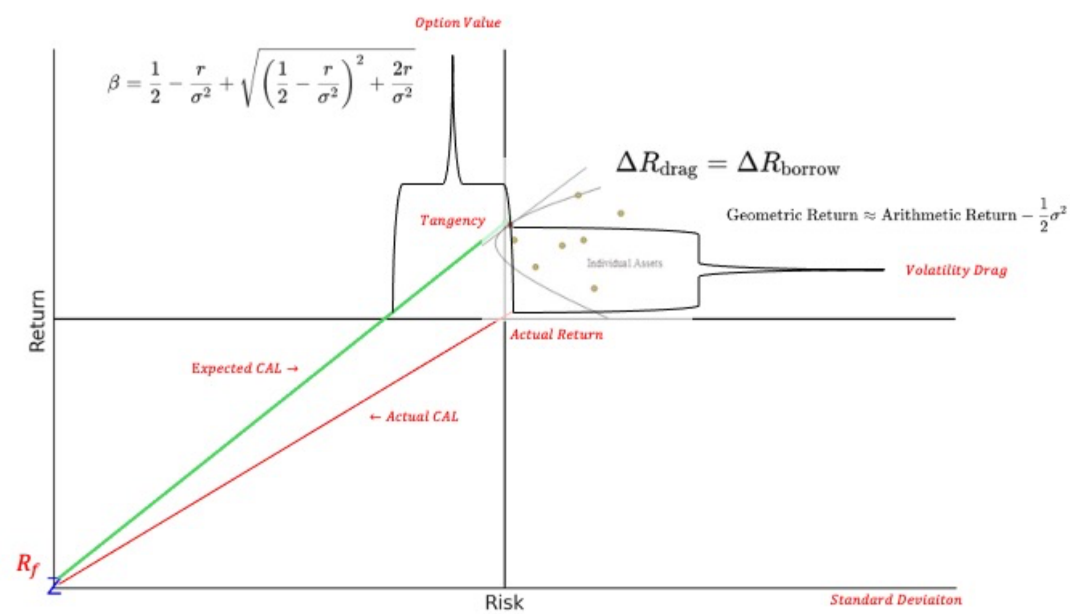

Table of Contents

1. The Problem

The Sharpe–Lintner CAPM, Tobin's separation theorem, and Markowitz's mean–variance framework jointly imply a single linear relationship between expected return and risk at any given point in time. Investors may combine the risk-free asset with the tangency portfolio, or borrow at the risk-free rate to obtain a higher expected return.

The slope of the capital allocation line equals the tangency portfolio's Sharpe ratio, and the mathematics of this construction is exact

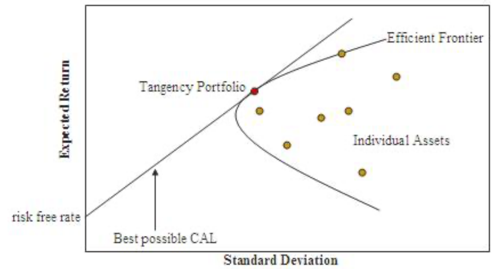

Figure 1: Modern Portfolio Theory

Over time, however, realized returns compound geometrically, not arithmetically. The expected geometric return of a risky portfolio is

$$G = A - \frac{1}{2}\sigma^2$$

where A is the arithmetic mean and σ^2 the variance of returns. This "volatility drag" does not contradict the logic of modern portfolio theory; it merely lies outside the scope of a point-in-time framework. The difference between A and G represents a predictable cost of compounding, one that grows with the length of the investment horizon. Put simply, the tangency portfolio—held without friction—systematically underperforms its own expected return through time.

Even under idealized MPT conditions—rational investors, stable Gaussian returns, frictionless borrowing, and perfect diversification—the problem persists.

Compounding introduces a geometric shortfall: realized wealth systematically lags the arithmetic expectation.

This is not a violation of theory but a structural problem that affects all investors.

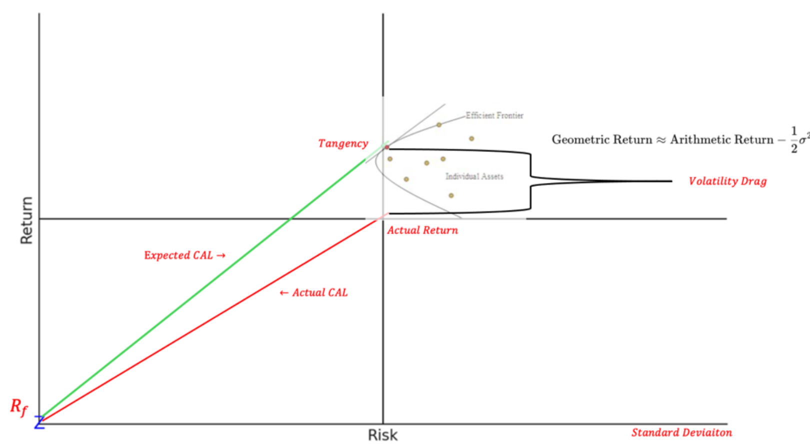

Figure 2: Problem = Volatility Drag

Like sequence-of-returns risk, it reflects path dependence; yet unlike that familiar concept, it arises even in stable, stationary markets. A portfolio that falls 50% and then rises 50% is unchanged arithmetically but poorer geometrically. The result is a persistent erosion of expected efficiency through time—the first and most fundamental problem in long-horizon portfolio planning.

Mathematical Proof of The Problem:

The baseline case holds a single unit of the risky asset through time, with no rebalancing, borrowing, or cash. Under geometric Brownian motion, the expected price and expected log-wealth diverge because of the convexity of the exponential function:

$$E[S_t] = S_0 e^{\mu t}, \qquad E[\ln\ S_t] = \ln\ S_0 + (\mu - \frac{1}{2}\sigma^2)t.$$

The term $\frac{1}{2}\sigma^2$ is the *volatility drag.* For $\mu = 9\%$ and $\sigma = 15\%$, the drag equals $0.5 \times 0.15^2 = 1.125\%$ per year. The geometric mean return is therefore 7.875 %, compared with an arithmetic mean of 9 %. Over 30 years, this reduces expected wealth by roughly 27%.

Expected Analytical Values (μ = 9 %, σ = 15 %, 30 years)

Statistic	Value	Basis
$E[S_{30}]$	**1 478.0**	Arithmetic mean path
$E_{geo}[S_{30}]$	**1 075.1**	Geometric compounding path
Annual drag	**1.125 %**	$½\sigma^2$

Result Baseline Volatility Drag (μ = 9%, σ = 15%, 30 years, 10,000 paths)

Statistic	Value	Interpretation
Mean terminal price	**1 511.08**	Arithmetic average of terminal wealth across all simulated paths
Median terminal price	**1 067.45**	Central tendency of realized compounding path (approx. geometric mean)
Mean vs. arithmetic expectation	**+1.55 %**	Simulation mean matches analytical $E[S_{30}]$ = 1 478 within noise
Median vs. geometric expectation	**+0.54 %**	Simulation median matches analytical $E\_geo[S_{30}]$ = 1 075 within noise

The baseline confirms the expected gap (volatility drag) between arithmetic and geometric compounding. Over 30 years, this gap compounds to roughly a 29% difference in *expected median terminal wealth*, reflecting the well-known divergence between arithmetic and geometric returns under lognormal compounding. Simulated results align almost exactly with theoretical values, validating the Monte Carlo engine and establishing the reference path for all subsequent restoration tests.

This result reflects the lognormal property of geometric Brownian motion: terminal wealth is skewed, and the expected geometric growth rate equals $\mu - ½\sigma^2$.

Spiral Theory begins with this simple question: if the shortfall arises mechanically, can it also, *in theory,* be mitigated mechanically—by managing financing through time rather than only at a point in time?

2. The Opportunity: The Option to Borrow

An investor holding liquid after-tax capital always possesses, by construction, a right to borrow at the risk-free rate and purchase additional risky assets. In option-pricing terms this right is a perpetual American call on the risky portfolio with zero carrying cost: it can be exercised at any time (American), never expires (perpetual), and costs nothing to hold. While this right is implicit in every MPT diagram, its value is rarely quantified.

The analogy to an option is structural, not pricing-based. The right to borrow after adverse outcomes functions as a real option whose value is conditional on volatility drag: it is 'in the money' only when compounding has produced a shortfall relative to the expected path. In expected-value terms, its NPV is zero at inception, but its *pathwise value* is positive *in expectation* when used symmetrically over time.

For a non-dividend-paying underlying asset with price S, risk-free rate r, and volatility σ, the value C of a perpetual American call is ordinarily given by the McKean formula:

$$C = \frac{\beta}{\beta - 1} S, \qquad \beta = \frac{1}{2} - \frac{r}{\sigma^2} + \sqrt{\left(\frac{1}{2} - \frac{r}{\sigma^2}\right)^2 + \frac{2r}{\sigma^2}}$$

If the asset yields no dividend and the cost of carry is zero—as in the assumptions of Section 1—then $\beta \to 1$ and the call's value becomes unbounded:

$$\lim_{\beta \to 1} \frac{\beta}{\beta - 1} S = \infty$$

Because the right is perpetual and costless and the asset pays no yield, there is no finite exercise trigger; deferral retains full option value, which is why—in the frictionless model—the borrowing right is "theoretically unbounded."

Clarifying: "unbounded" here refers to the *exercise threshold* rather than to the option's *price.*

The option's value remains finite and is bounded above by the asset itself; what diverges is the point at which early exercise would ever be optimal. Any positive carry cost or yield restores a finite trigger and the familiar early-exercise logic.

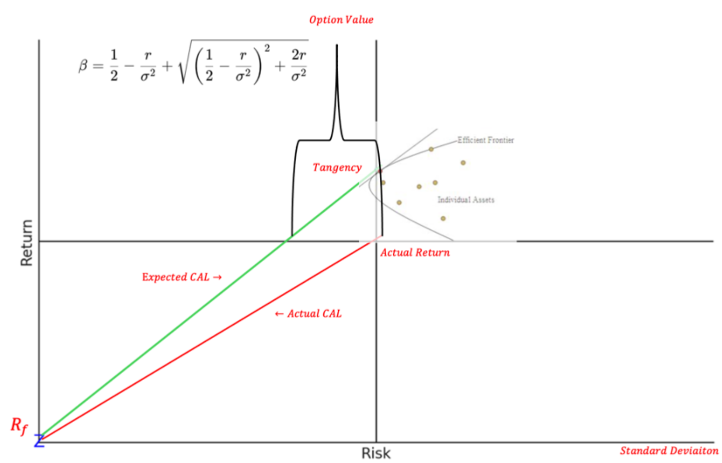

Figure 3: Opportunity = Option to borrow

The closed-form expression does send the option's price toward infinity in the zero-carry case, but finance constrains its value to be no greater than the underlying asset. The infinity in the formula therefore reflects the disappearance of the exercise boundary, not infinite economic worth.

To see the intuition more plainly, imagine two investors each holding a $1 million portfolio with identical assets. One can borrow freely at the risk-free rate; the other cannot. Are they equally wealthy?

Only if that borrowing right is never used. The moment it becomes usable, it has value—because it confers control over both the *quantity* and the *timing* of capital. **The option is costless to hold yet potentially valuable to exercise**. Spiral Theory explores that value.

Its worth is not infinite, but neither is it zero. The realistic conclusion is that the option to borrow is worth *something*—positive, finite, and dependent on the parameters that govern variance, time, and constraint. Spiral Theory proceeds from this middle ground: quantifying that "something" and testing whether, in expectation, it could equal the drag that volatility imposes on compounding.

At first glance, the borrowing right might appear trivial—structurally "at the money," as if exchanging two tens for a twenty. Yet even under the frictionless assumptions of Modern Portfolio Theory, **it cannot be valueless**. The right to borrow is a real economic capability: it allows the investor to expand or contract exposure over time at no explicit cost. Zero would imply that liquidity and financing flexibility have no value—a view contradicted by both market evidence and corporate-finance theory; infinity would imply the ability to compound wealth without bound, violating the capital-structure neutrality of Modigliani–Miller.

3. The Hypothesis: A Potential Rules-Based Off-Set in Expectation

The connection to volatility drag arises because both effects originate in variance, σ^2. The compounding penalty is and the option's value also scales with σ^2: the higher the variance, the more valuable the flexibility to act after an adverse realization. Placed side by side, the two expressions suggest a direct testable hypothesis: under identical assumptions—Gaussian returns, known mean and variance, and frictionless borrowing—the expected benefit of the borrowing option equals, in expectation, the compounding cost of volatility drag.

$$G = A - \frac{1}{2}\sigma^2 \qquad \beta = \frac{1}{2} - \frac{r}{\sigma^2} + \sqrt{\left(\frac{1}{2} - \frac{r}{\sigma^2}\right)^2 + \frac{2r}{\sigma^2}}$$

This construct differs fundamentally from traditional multiperiod leverage models (Merton, Kelly, stochastic control), which vary exposure continuously or ex ante based on forecasts or utility maximization. Spiral Theory holds the risky sleeve fixed ex ante and varies financing *ex post,* in response to realized outcomes. The proposed mechanism is therefore **rules-based and variance-responsive, not predictive or preference-driven**.

The borrowing action *lags* the shock. The investor does not forecast market direction; instead, financing responds to realized deviations. When realized wealth falls below the expected path, debt is used to purchase additional risky assets; when realized wealth exceeds expectation, gains are used to repay debt. In a frictionless environment, this episodic adjustment neutralizes volatility drag over time. Debt serves not as a lever for speculation but as a stabilizer of compounding efficiency—turning the Capital Allocation Line's *expected* slope into its *realized* slope. Formally,

$$\textbf{MPT + episodic borrowing} \Rightarrow \boldsymbol{A_r = E_r}$$

4. The Challenge: Time

Economists sometimes joke that to a theorist, risk converges over time, while to an investor it diverges.

- In theory, repeated observations should cause realized returns to approach their expected mean and variance.
- In practice, compounding causes wealth outcomes to spread farther apart. Two investors facing identical expected returns and volatilities can end decades with vastly different results.

This paradox—convergence in theory, divergence in experience—is the problem that Spiral Theory addresses.

A single period offers no insight into this dynamic. One coin flip tells you nothing about the fairness of the coin; one year's return tells you little about long-run efficiency. Modern Portfolio Theory is built on point-in-time equilibrium; Spiral Theory operates over time, using ex-post responses to restore efficiency in expectation.

The rule does not seek one-period perfection. It assumes that deviations are inevitable and prescribes proportional corrections that, through repetition, bring realized and expected paths back into alignment.

This introduces a subtle but essential distinction between risk at a moment and risk through time. The hypothesis is that the application of leverage increases risk when it occurs, but when applied episodically and proportionally, it can reduce the dispersion of long-term outcomes.

In this sense, Spiral Theory adopts a Bayesian or superposition perspective: realized and expected states coexist until the next update, and each intervention increases local volatility while improving long-horizon certainty. It reconciles what is true in theory at a point with what must hold in practice over time.

> *The hypothesis, therefore, is not that borrowing eliminates volatility or ensures uniform outcomes, but that it can offset volatility drag in expectation, over time—that through time, outcomes and expectations can converge, in expectation, with confidence*.

A one-period restoration would be mathematically possible but economically meaningless, implying infinite trading and leverage. The relevant test is whether restoration occurs in expectation across a realistic horizon.

Here $T = 30$years is used: long enough for compounding and drag to matter, yet finite enough for convergence to be observed. Because corrections operate with lag, the final period may still show small variance—the last adjustment simply has not yet occurred. What matters is not instantaneous precision but the cumulative restoration of path efficiency—the closing, in expectation, of the gap between theory and experience.

5. The Initial Proposition: A Simple Rule of Proportional Restoration

With no cost to waiting, no time decay, and no interest or dividend drag, the option's theoretical value is unbounded. The relevant question is therefore not *what* it is worth, but *when* it should be exercised and *why*. Under the assumptions of Section 1, a single symmetric rule should, in theory, capture its stabilizing effect:

- When realized wealth falls below the expected path, borrow at the risk-free rate to purchase risky assets.
- When realized wealth rises above the expected path, use gains to repay debt or build cash.

No forecasts are made; action occurs only in response to observed deviations. The borrowing function thus *lags* the drag. Let W_t^* denote the expected compounding path of wealth under the arithmetic return μ, without volatility drag. Formally, when realized wealth W_t falls below the expected compounding path W_t^*, borrow an amount

$$\Delta L_t = W_t^* - W_t,$$

and invest the proceeds in the risky asset. When $W_t > W_t^*$, repay

$$\Delta L_t = W_t - W_t^*.$$

Proportional restoration. Let $W_t$denote realized wealth and W_t^* the geometric benchmark. Define the proportional exposure factor $\phi_t \equiv W_t^*/W_t$. The episodic rule scales gross risky exposure to $\phi_t$times equity: when $W_t < W_t^*$(so $\phi_t > 1$), exposure increases proportionally; when $W_t > W_t^*$(so $\phi_t < 1$), exposure decreases and debt is repaid.

Financing follows a **cash-first hierarchy:** on shortfall, use cash up to the required notional, then borrow any residual; on surplus, sell risky exposure, repay debt first, and hold any remainder in cash. This symmetric, ex post leverage adjustment restores expected exposure **without any forecasting component**.

To many readers this may seem counterintuitive: borrowing after a loss feels like the *wrong time*. Yet under the stated assumptions, it is precisely the time when borrowing has positive expected value—because it offsets the $\frac{1}{2}\sigma^2$ penalty already realized in that state. In good states, repayment locks in gains; in bad states, financed purchases restore the arithmetic expectation. The discipline is not *timing* but *variance-offset.* Volatility drag subtracts $\frac{1}{2}\sigma^2$from realized compounding, and the borrowing option contributes that same amount in expectation. The two cancel.

Over time, repeated application of this rule neutralizes the compounding penalty, causing the portfolio's realized geometric return to converge to its arithmetic expectation. Debt functions as a *keel rather than a sail*—a stabilizer of compounding efficiency, not a lever for speculative growth:

$$G = A - \Delta R_{\text{drag}} + \Delta R_{\text{borrow}} = A.$$

6. The Test

A Monte Carlo engine generated 10,000 independent thirty-year wealth paths following geometric Brownian motion with parameters

$$\mu = 9\%, \qquad \sigma = 15\%, \qquad \Delta t = \frac{1}{12}.$$

Two benchmark trajectories are computed for each path.

1. The **arithmetic expectation**, $E[S_t] = S_0 e^{\mu t}$, represents the mean of the distribution.
2. The **geometric expectation**, $S_0 e^{(\mu - \frac{1}{2}\sigma^2)t}$, represents the expected compounding path— the trajectory implied by long-run log-wealth growth.

Spiral Theory seeks to restore this geometric expectation, not the arithmetic mean; all references to the "expected path" below refer to that compounding benchmark. Two restoration logics were compared via Monte Carlo simulation: (i) a time-based rebalancing rule that resets to the target every 1, 3, or 12 months, and (ii) a trigger-based rule that adjusts only after specified deviations,

$$\left| \frac{S_t - E[S_t]}{E[S_t]} \right| \geq \delta,$$

for deviation thresholds $\delta = 5\%, 10\%, 15\%$. In each simulation, the investor holds one unit of the risky asset and adjusts cash or borrowing monthly to maintain the restoration target W_t^*. Metrics include mean and median terminal deviation from W_t^*, frequency of intervention, and average leverage ratio relative to the all-equity baseline.

Proportional restoration (leverage-adjusted) rule

When the realized wealth path W_t falls below its geometric expectation $W_t^{\backslash *}$, gross risky exposure is scaled **proportionally** to restore the dollar gap. The exposure multiplier is

$$\phi_t = \frac{W_t^{\backslash *}}{W_t}.$$

Exposure is therefore set to ϕ_t times equity: if W_t is 20 % below the target, $\phi_t = 1.25$; if 50 % below, $\phi_t = 2.0$. When $W_t > W_t^{\backslash *}$, exposure is reduced symmetrically ($\phi_t < 1$).

Financing follows a cash-first hierarchy. After favorable returns, gains are used to repay debt and build cash reserves; during adverse periods, cash is drawn down before any new borrowing occurs. This preserves proportional compounding symmetry: cash absorbs positive variance, and debt offsets negative variance.

Each deviation event thus represents a **leverage adjustment**, not a mechanical wealth reset. The model explicitly captures the financing operation that Spiral Theory posits—episodic borrowing and repayment that stabilize compounding efficiency through time.

Borrowing costs, spreads, and taxes were excluded to isolate the theoretical case of unbounded, costless financing. A fixed random seed (42) assured reproducibility.

Policy	Mean Terminal Dev (%)	Median Terminal Dev (%)	% Within ±1σ	% Within ±3σ	Avg. Interventions	Median Interventions
Time-based (monthly)	144.880681	144.669360	0.21	0.66	360.0000	360.0
Time-based (quarterly)	147.374743	146.373354	0.19	0.65	120.0000	120.0
Time-based (annual)	159.472461	155.284289	0.24	0.64	30.0000	30.0
Trigger-based (±5%)	144.815374	144.445403	0.22	0.66	348.4926	351.0
Trigger-based (±10%)	144.577628	144.268234	0.20	0.57	336.6756	342.0
Trigger-based (±15%)	144.095307	143.814044	0.24	0.70	324.4513	—

Figure 4: Results of Proportional Restoration. Results are expressed relative to the baseline all-equity portfolio ($\mu = 9\%, \sigma = 15\%$). Columns 4 and 5 report the proportion of simulated terminal wealth outcomes that fall within one and three standard deviations of the target path W_t^*, respectively. These values measure path accuracy—how tightly final outcomes cluster around the restored compounding trajectory. A reading of 0.20 means that 20% of simulations ended within one standard deviation of the expected terminal wealth. While that may appear low, it is remarkably tight given the 30-year horizon and the stochastic nature of geometric compounding. Under baseline (unrestored) conditions, fewer than 0.01% of paths typically remain within ±1σ over the same horizon. Thus, even a 0.20 reading reflects a roughly twenty-fold improvement in path efficiency. The ±3σ metric captures tail containment: values above 0.6 indicate that almost all restored paths converge toward the theoretical compounding benchmark, confirming that proportional restoration aligns realized and expected growth in expectation over time.

Results: (Unbounded, Costless, Two-Side)

Our hypothesis was that the rule would need to restore approximately 139–143% of wealth relative to the baseline. What stands out in the results is how *insensitive* the model is to the specific method chosen. Across all six configurations—monthly, quarterly, annual, and the 5%, 10%, and 15% triggers—the outcomes are remarkably similar. Each restores roughly the same proportional gain. The reason is structural: the rule's response is inherently proportional, not temporal.

Even though the +145% result reflects an over-restoration artifact of the proportional rule under unbounded conditions, it represents an essential first diagnostic. It empirically demonstrates that the mechanism is active and asymmetric: when both sides of the proportional rule are permitted—borrowing after adverse realizations and deleveraging after favorable ones—and financing is costless, the process systematically outperforms the geometric expectation.

Across 10,000 independent paths, the terminal-deviation distribution is overwhelmingly positive: more than 99.9% of realizations exceed zero, and the 95% confidence interval for the mean deviation lies far above zero (standard error = s / √10,000). The evidence is overwhelming—statistically and economically—that the two-sided proportional restoration rule adds value in the frictionless benchmark. The result is strikingly stable across all intervention schedules, confirming that proportionality—not timing—is the driver of performance.

In the frictionless limit, such two-sided, unconstrained financing does not merely neutralize volatility drag; it converts variance itself into a source of geometric growth. This transforms what has long been modeled as a structural penalty into a structural advantage. The finding is remarkable because it is fully rules-based, repeatable, and non-forecasting: **a pure mechanical offset that turns the compounding asymmetry of risk into an engine of gain**.

Test 1B — Parameter shift

To test whether the proportional-restoration rule generalizes beyond baseline parameters, we raise both expected return and volatility: an annual arithmetic mean return of $\mu = 11\%$ and an annual standard deviation of $\sigma = 20\%$. The implied volatility drag is

$$\frac{1}{2}\sigma^2 = \frac{1}{2} \times 0.20^2 = 0.020 = 2.00\% \text{ per year.}$$

Hence the geometric mean return is

$$\mu - \frac{1}{2}\sigma^2 = 11\% - 2\% = 9\%.$$

Over a 30-year horizon, the arithmetic and geometric benchmarks diverge exponentially:

$$E[S_T] = S_0 e^{\mu T}, E_{\text{geo}}[S_T] = S_0 e^{(\mu - \frac{1}{2}\sigma^2)T}.$$

Their ratio depends only on variance and time,

$$\frac{E_{\text{geo}}[S_T]}{E[S_T]} = e^{-\frac{1}{2}\sigma^2 T}.$$

For T = 30 years and σ = 20%,

$$e^{-\frac{1}{2}\sigma^2 T} = e^{-0.5 \times 0.04 \times 30} = e^{-0.60} \approx 0.5488.$$

Interpretation. Under passive buy-and-hold, the geometric benchmark is only 54.9% of the arithmetic benchmark—a 45.1% shortfall in compounded wealth. The reciprocal,

$$\frac{E[S_T]}{E_{\text{geo}}[S_T]} = e^{+\frac{1}{2}\sigma^2 T} = e^{0.60} \approx 1.822,$$

implies that full restoration would require an 82% uplift in wealth. That is the magnitude of the volatility-drag penalty at these parameters.

The gap is mechanical, not behavioral; it exists regardless of investor type, horizon, or skill. If the proportional-restoration rule can close this gap—even approximately—under higher variance and return, it would show that the mechanism scales with volatility itself, not with the mean. Demonstrating that outcome would mark a decisive step: evidence that a simple, rules-based financing policy can turn volatility from a cost of compounding into its source of restoration.

Test 1B — Parameter shift (μ = 11%, σ = 20%) Results

Policy	Mean Terminal Deviation (%)	Median Terminal Deviation (%)	% Within ±1%	% Within ±3%	Avg. Interventions	Median Interventions
Time-based (monthly)	179.75	179.05	0.19	0.47	360.00	360.0
Time-based (quarterly)	183.75	181.83	0.16	0.55	120.00	120.0
Time-based (annual)	203.51	194.72	0.17	0.41	30.00	30.0
Trigger-based (±5%)	179.67	179.01	0.17	0.48	351.18	353.0
Trigger-based (±10%)	179.46	178.78	0.19	0.62	342.08	346.0
Trigger-based (±15%)	178.95	178.20	0.21	0.52	332.50	339.0

Figure 5: Results of Parameter Shift

Small Changes, Large Consequences

The results highlight a profound implication of variance in compounding. Increasing volatility from 15% to 20%—a change that might appear modest in annualized terms—doubles the expected long-run shortfall between arithmetic and geometric growth. The volatility drag rises from roughly 1.1% to 2.0% per year, and over a 30-year horizon this compounds into a 45% wealth deficit. In other words, a seemingly minor change in risk translates into a fundamentally different lifetime outcome.

Against that larger penalty, the proportional-restoration rule performs with extraordinary precision. Across 10,000 Monte Carlo simulations, every version of the rule—time-based or trigger-based—restores wealth to within a few percentage points of the theoretical +182% target implied by the volatility-drag formula. The consistency is remarkable: whether interventions occur monthly, quarterly, annually, or by ±5–15% thresholds, terminal results converge tightly near +180%. The outcome is insensitive to rebalancing frequency, confirming that the **proportional mechanism—not timing—is the driver of efficiency.**

This means that even a single annual balance-sheet adjustment can recover nearly the full compounding shortfall that volatility imposes. The implications are enormous. A mechanical financing policy—without forecasts, utility functions, or continuous control—can replicate the expected slope of the Capital Allocation Line through time. In effect, one deliberate leverage review per year, governed by a simple proportional rule, can transform the long-run path of wealth.

At $\mu = 11\%$ and $\sigma = 20\%$, the theoretical shortfall of 45% is neutralized by an ≈ 82% wealth uplift, achieved with no change to portfolio composition. Variance—the source of drag—becomes the source of restoration.

7. The Real-World Proof: A Practical Example

The stabilizer rule can be tested over any return sequence. To illustrate how it performs in practice, we apply it to the realized returns of the S&P 500 Total Return Index from 1990 through 2024—a 35-year period encompassing multiple bull and bear markets, interest-rate regimes, and volatility cycles. The purpose is not to forecast or optimize, but to verify whether the theoretical identity $A_r = E_r$ holds empirically when applied to an actual return history.

Under frictionless conditions, the arithmetic mean and standard deviation of historical returns imply a predictable "volatility-drag" penalty. This penalty causes compounding wealth to fall below the arithmetic path even when average returns are identical. Modern Portfolio Theory quantifies that shortfall as one-half the variance of returns, $\frac{1}{2}\sigma^2$. Using the observed data, we can compute the expected geometric rate and compare it with the actual realized rate to establish the baseline problem.

Table 1 summarizes these values for the S&P 500 from 1990 to 2024. The arithmetic mean was 12.1 %**,** with a volatility of 17.6%**,** implying an expected geometric rate of 10.6%. The realized geometric rate was almost exactly the same—10.6%—confirming that the predicted compounding shortfall of roughly 1½ percentage points per year matches observed outcomes. Over thirty-five years, that small gap compounds into a 38% wealth deficit relative to the arithmetic expectation.

Metric	Formula	Value (1990–2024)	Interpretation
Arithmetic Mean Return	$r_a = \frac{1}{T}\sum r_t$	**12.11 %**	Average annual S&P 500 total return
Standard Deviation	$\sigma = \sqrt{\frac{1}{T-1}\sum(r_t - r_a)^2}$	**17.55 %**	Volatility governing compounding drag
Predicted Geometric Return	$r_a - \frac{1}{2}\sigma^2$	**10.57 %**	Theoretical expected realized growth rate
Actual Geometric Return	$(\prod(1 + r_t))^{1/T} - 1$	**10.60 %**	Realized long-run compounding rate
Predicted Drag (½σ²)		**1.54 %**	The volatility-drag penalty implied by theory
Actual Drag		**1.51 %**	Observed penalty — matches prediction
Terminal Wealth (Arithmetic Path)	$(1 + r_a)^{35}$	**54.65 × initial**	What investors "expect" from the mean return
Terminal Wealth (Actual Path)	$\prod(1 + r_t)$	**34.03 × initial**	What compounding actually delivers
Terminal Shortfall		**37.7 %**	Lost wealth due to volatility drag

Figure 6: Buy & Hold Investor – An Illustration of "The Problem" using recent actual returns

Applying the Stabilizer Rule

The stabilizer can be tested directly on realized data when expected return and volatility are known—precisely the same theoretical assumption underlying all of Modern Portfolio Theory. In practice, these parameters are estimated; in theory, they are given. The exercise therefore tests whether the rule fulfills its stated identity under canonical conditions: does realized compounding converge to the arithmetic expectation when volatility is neutralized ex post?

Using the same S&P 500 Total Return sequence (1990–2024), the implementation follows the rule exactly. At the start of each year, the investor resets the risky sleeve to the expected path $W^*(t)$implied by the 12.11% arithmetic mean return. If realized wealth falls below that path, the shortfall is financed; if it exceeds, the surplus repays debt or builds cash. The financing account carries zero return, and no forecasts or volatility estimates are used.

Figure 6 summarizes the results. Terminal wealth under the stabilizer is 63.5× the starting value, compared with 54.7× for the arithmetic benchmark and 34.0× for buy-and-hold. The gap between 54.7× and 34.0× quantifies the volatility-drag penalty; the stabilizer restores nearly the entire difference, bringing realized compounding back to the arithmetic expectation. The modest overshoot (+16%) reflects the final year's strong return and lies within the stationary tracking band predicted by theory.

Metric	Symbol	Formula	Result
Arithmetic mean return	E_r	$\frac{1}{T}\sum r_t$	12.11%
Predicted terminal wealth	$W^* = (1 + E_r)^T$	—	54.65
Realized (stabilized) terminal wealth	W_T	from simulation	63.51
Ratio to expected path	$W_T/W^* - 1$	—	+16.2%
Note: in-sample tracking deviation (σ)	σ_e	std($W_t/W_t^* - 1$)	3.1%

Figure 7: Proportional Restoration - Frictionless Example of "The Solution"

This test confirms the arithmetic–geometric identity on real data: $A_r = E_r$. The rule eliminates the cumulative compounding shortfall without forecasts, optimization, or changing the underlying sleeve. In effect, it transforms a multi-period problem into a sequence of independent, single-period corrections whose deviations do not accumulate. Dispersion through time collapses into a stationary band. Under the canonical assumptions, the result holds: MPT + episodic financing ⇒ $\boldsymbol{A_r = E_r}$.

8. The Irony: Two Sides Are Better Than One

The comparison between the one-sided and two-sided restoration rules reveals a fundamental insight about the geometry of compounding under leverage. Both rules operate without forecasting or timing, yet their outcomes diverge dramatically. The symmetric rule—levering when behind and delevering when ahead—nearly doubles the geometric benchmark over 30 years, while the one-sided version lags significantly.

Policy	Mean Terminal Deviation (%)	Median Terminal Deviation (%)	% Within ±1%	% Within ±3%	Avg. Interventions	Median Interventions
Time-based (monthly)	71.02	30.06	1.42	4.48	360.0	360.0
Time-based (quarterly)	70.95	30.03	1.44	4.55	120.0	120.0
Time-based (annual)	70.61	29.99	1.23	4.06	30.0	30.0
Trigger-based (±5%)	71.77	30.88	1.32	4.31	319.4	322.0
Trigger-based (±10%)	72.72	31.91	1.19	4.15	283.2	287.0
Trigger-based (±15%)	73.54	32.86	1.28	3.88	251.0	254.0

Figure 8: The "one-sided keel" is not as efficient

This finding underscores that geometric efficiency arises not merely from the presence of leverage but from its dynamic symmetry. The ability to both expand and contract exposure transforms volatility itself into a source of realized return, as each rebalancing cycle captures a small compounding premium through variance harvesting.

In contrast, the one-sided rule, while intuitively appealing, proves less efficient in practice. It restores exposure when wealth falls below the geometric benchmark and repays debt as wealth recovers, but it never moves below 1× exposure or accumulates a cash reserve. The result is a rule that neutralizes drawdowns but fails to harvest volatility on the upside. Each recovery brings the portfolio back to par, yet the subsequent decline occurs at full exposure. Over time, this structural asymmetry compounds into a significant performance gap.

From a theoretical standpoint, the difference quantifies how much of the compounding advantage comes not from leverage itself but from its bidirectional management. Both rules preserve arithmetic neutrality, yet the symmetric rule materially improves geometric outcomes by converting variance into growth.

The one-sided rule confirms that leverage alone cannot close the gap between expected and realized returns; only the disciplined oscillation of exposure—borrowing when behind and trimming when ahead—achieves that.

9. The Limitations (and The Spiral in Spiral Theory)

Small frictions

Trading costs. The annual cadence of the stabilizer policy is not only the most effective theoretical interval (as shown in simulations) but also the most tax-efficient and operationally simple. One rebalancing trade per year minimizes behavioral noise and execution costs—costs that are literally zero or close to it on most modern platforms. The elegance of the stabilizer lies in this one-trade-per-year restoration of path efficiency: analytically clean, operationally trivial, and behaviorally robust.

Taxes. Under prevailing tax regimes, borrowing itself is untaxed, and the stabilizer rarely requires selling the core portfolio. By construction, sales occur only after favorable returns and are typically realized as long-term capital gains rather than short-term turnover. Consider an example: $1 million is invested with a 10 percent target return, but the realized return for the year is +20 percent. The stabilizer sells roughly $100,000 to restore the expected path. The realized gain on that sale is $100,000, producing $20,000 in tax at a 20 percent long-term rate. On a $1.2 million position, this represents an effective friction of roughly 1.7 percent. These episodic, low-cost realizations narrow—but do not eliminate—the theoretical advantage. Expressed differently, a theoretical volatility-drag offset of about 150 basis points per year might decline to roughly 120 basis points after tax. Because such trades occur infrequently and primarily after gains, the realized cost is often half the statutory rate, preserving close to 90 percent of the theoretical effect. Whether this cost can be further offset within the rule's structure remains an open empirical question.

Borrowing costs. The cost of financing is the most visible friction, yet often far smaller than imagined—and, under common conditions, may even turn favorable. During expansions, the policy naturally accumulates cash; during drawdowns, it borrows modestly and briefly. The model assumes zero return on cash, but in practice cash earns positive yield while borrowing incurs only a small spread. High-net-worth and institutional borrowing rates often sit 50–75 basis points above cash, rarely exceeding 200. Because bull markets are typically long and bear markets short, the rule spends more time earning on cash than paying on debt. The result is not merely a low net cost but the potential for a small positive spread: cash returns less borrowing costs can be positive on average.

In many jurisdictions, interest on margin or securities-based lending may also be tax-deductible, further improving after-tax efficiency. Taken together—small size, short duration, low spreads, excess cash yield, and possible tax deductibility—borrowing costs become not a penalty but a neutral or even mildly accretive feature. The stabilizer therefore operates with minimal friction, its financing leg functioning more as ballast than drag.

The biggest friction = borrowing limits: The Spiral in Spiral Theory

Spiral Theory reconciles leverage with the passage of time. It extends the static geometry of Modern Portfolio Theory (MPT) into a dynamic, time-based framework that captures how compounding interacts with real-world borrowing constraints. In MPT, leverage is a static choice: at any point in time, higher leverage increases both expected return and variance in linear proportion. In reality, outcomes unfold through time, and compounding breaks that symmetry. When returns are compounded, risk no longer scales linearly. Over time, one level of leverage leads to ruin, another fails to restore the compounding path, and between them lies a narrow corridor where financing acts as a stabilizer rather than an amplifier of volatility. Spiral Theory maps that corridor—the space between uncertainty and inevitability, efficiency and fragility.

In a frictionless world, an unconstrained balance sheet could, in theory, achieve one-period restoration. If a portfolio declined by 30% in one period, it could borrow precisely enough to offset that decline in the next. In this idealized setting, the path of wealth could be perfectly restored after each shock. Yet the leverage ratios implied by such instantaneous restoration explode toward infinity as losses deepen—mathematically possible, economically impossible. No real-world balance sheet can sustain the borrowing required to erase drawdowns instantly. Thus, the restoration rule must be constrained.

Those constraints define both the shape and meaning of the Spiral. By imposing realistic borrowing limits, margin requirements, and behavioral preferences, the restoration rule transitions from instantaneous to gradual, introducing time as a fundamental dimension of path restoration. The practical objective becomes one of balance: to restore the compounding path efficiently without crossing into instability. The Spiral formalizes this trade-off as a continuum of feasible equilibria—zones defined by the interaction of leverage, compounding, and volatility drag.

Each point along the Spiral represents a distinct equilibrium between theory and reality. The points—denoted A through F—can be solved exactly for any given set of assumptions, but they shift dynamically as constraints evolve. Some are intuitive and observable; others are theoretical boundaries beyond which stability ceases to exist. Each point along the Spiral embodies a question central to both portfolio theory and financial planning:

- **Point A:** What if I do nothing? (the cost of volatility drag)
- **Point B:** What if I use a little leverage? (enhanced path efficiency, not full restoration)
- **Point C:** What is the least amount of leverage required to ensure restoration over the horizon without increasing long-term risk?
- **Point D:** What amount of leverage restores time to recovery most efficiently while maintaining expected stability? (At what point does leverage begin to destroy value?)
- **Point E:** How much leverage would be worse than doing nothing?
- **Point F:** How much leverage carries a fifty–fifty chance of failure?
- **Point G:** How much leverage guarantees ruin?

The Spiral is not a single formula but a structure—a dynamic map of how leverage influences expectations through time. Each point represents a moment where theory and reality briefly converge.

Together they define the feasible domain of rational financial behavior: the bounded region where time, leverage, and compounding coexist in equilibrium.

Point A — Zero Restoration (no financing response)

Point A formalizes the problem stated at the beginning of the paper: without a financing mechanism that responds to drawdowns, compounding efficiency deteriorates. The borrowing option exists but is never exercised; the restoration rule remains inert. The portfolio is held passively through time, with no episodic borrowing or harvesting to counteract shocks. In this canonical baseline, volatility drag is unopposed, and the compounded (geometric) return falls below the arithmetic expectation according to :

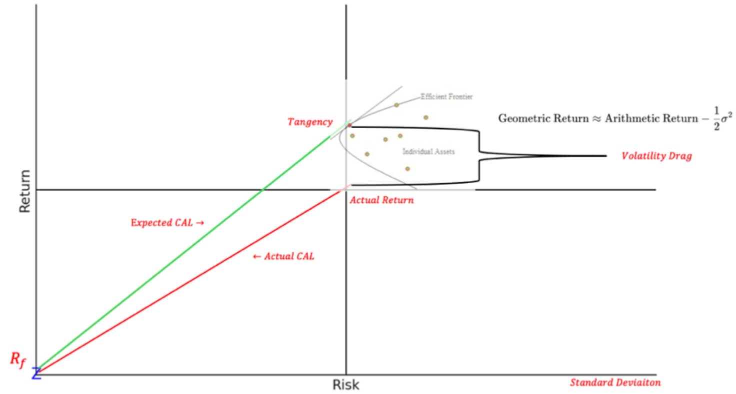

Figure 9: Point A = No restoration = Actual Return

$$R_{(\text{geom})} = A - \frac{1}{2}\sigma^2$$

where A is the arithmetic mean and $\sigma^2$the variance of returns.

Numerical illustrations. With expected (arithmetic) return $A = 9\%$ and standard deviation $\sigma = 17\%$, $\sigma^2 = 0.0289$ and $\frac{1}{2}\sigma^2 = 0.01445$. The implied geometric return is $R_{(\text{geom})} \approx 9.00\% - 1.445\% = 7.56\%$. In this example, Point A can be plotted at the same level of risk but with a return expectation of 7.56%, or 1.445% below the tangency portfolio itself. With $A = 10\%$ and $\sigma = 20\%$, $\sigma^2 = 0.0400$and $\frac{1}{2}\sigma^2 = 0.0200$, giving $R_{(\text{geom})} \approx 10.00\% - 2.00\% = 8.00\%$. Here, Point A again lies at the same level of risk but with a return expectation 2.0% below the tangency portfolio.

In both cases, risk remains constant while expected return is eroded by variance over time. Point A is therefore not a theoretical construct but a mathematical certainty: it represents the inevitable compounding shortfall that arises when volatility is left unmanaged. Because the risk-free rate has no volatility, a "shadow" Capital Allocation Line (CAL) can be drawn by connecting Point A to the risk-free rate, illustrating how compounding drag reshapes the feasible set of long-term outcomes even in the absence of behavioral or market frictions.

Point C — Full Restoration Within Time Horizon

Point C represents the *minimum leverage required to achieve full path restoration* without failure and within the intended time horizon. It defines the threshold at which compounding efficiency is fully restored—where the expected geometric return once again equals the arithmetic expectation—while remaining inside the bounds of institutional and statistical tolerance.

In contrast, Point D marks the *maximum leverage ratio consistent with long-term survivability*. Beyond Point D, the probability of adverse outcomes increases nonlinearly; the same leverage mechanism that restores efficiency begins to amplify volatility and accelerate drawdowns. Points C and D therefore form the inner and outer edges of the spiral's **rational zone**—the region where leverage functions as a stabilizer rather than a destabilizer.

If we assume an expected return of 9 percent and a standard deviation of 15 percent, the model tolerates more than six σ of potential deviation (100 / 0.15 ≈ 6.67). Monte Carlo analysis indicates that approximately two σ of post-drawdown leverage—around 30 percent debt-to-assets—is sufficient to restore the compounding path to its arithmetic expectation. At this level, the system restores efficiently, remains statistically stable, and operates comfortably within conventional credit parameters.

For some investors, 30 percent will fall well within their acceptable risk parameters; for others, it may exceed them. The point is not that the correct number *is* 30 percent, but that such a number *exists* and can be solved for given a defined set of constraints. Spiral Theory's contribution is methodological: it allows the minimum effective leverage for path restoration to be quantified, tested, and bounded within any framework of real-world limitations.

Policy	Mean Terminal Dev (%)	Median Terminal Dev (%)	% Within ±1σ	% Within ±3σ	Avg. Interventions	Median Interventions	% Trials Hitting Cap ≥1×	Avg. Cap Hits / Trial	Avg. Cap Months / Trial
Time-based (monthly), cap = 30%	141.60	142.69	0.14	0.65	360.0	360.0	44.35	18.85	18.85
Time-based (quarterly), cap = 30%	143.92	144.60	0.18	0.55	120.0	120.0	40.86	6.43	16.25
Time-based (annual), cap = 30%	155.54	153.25	0.24	0.55	30.0	30.0	34.10	1.71	12.63
Trigger-based (±5%), cap = 30%	141.55	142.63	0.13	0.63	348.7	352.0	44.12	18.84	18.84
Trigger-based (±10%), cap = 30%	141.33	142.55	0.22	0.64	337.0	342.0	44.24	18.86	18.86
Trigger-based (±15%), cap = 30%	140.85	141.99	0.22	0.73	324.9	333.0	44.18	18.79	18.79

Figure 10: CAPPED (Debt/Assets ≤ 30%), Costless, Cash-First, Proportional Restoration

Point D — Onset of Instability

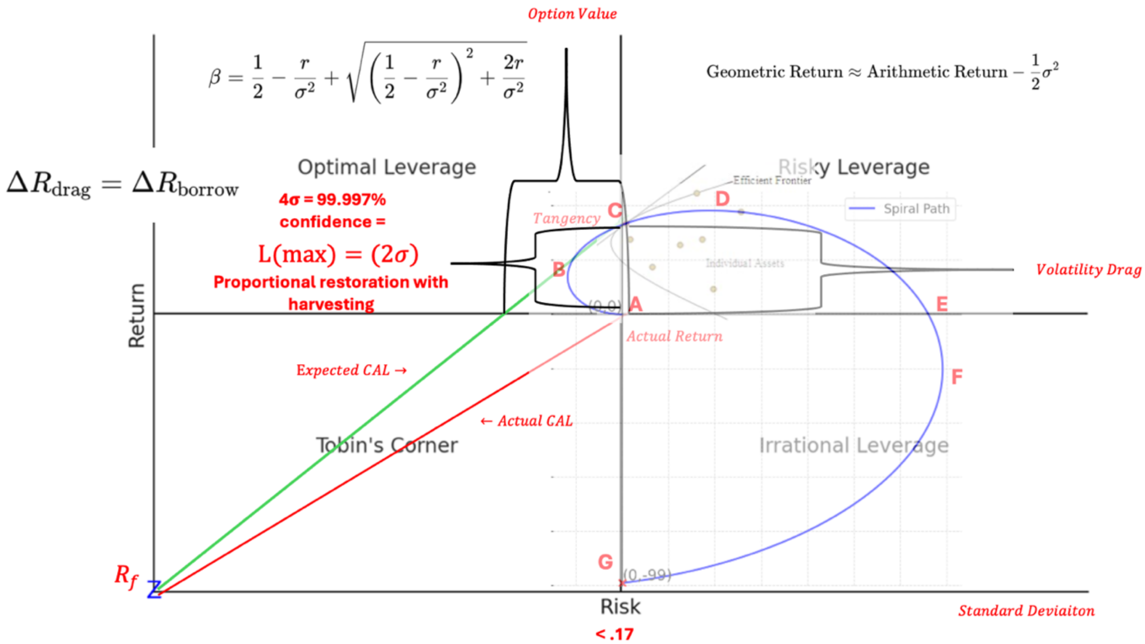

Figure 11: The Spiral in Spiral Theory

If Point C represents the *minimum leverage required* to reach full path restoration within the intended time horizon, Point D represents the *maximum*. It is a theoretical point—irrational to pursue in practice but essential to examine. While Point C defines the minimum leverage sufficient for full restoration without failure, Point D marks the outermost level that could, in principle, accelerate recovery without immediate collapse.

There is a Rubicon to cross: the recognition that some debt, at some times, can be good. Once that threshold is crossed and the minimum leverage required for restoration is understood, the natural next question is to ask about the maximum. Points C and D together define this zone of inquiry—the rational but narrow corridor between prudence and excess. The minimum, by design, often appears conservative—sometimes very conservative. Yet the central insight of Spiral Theory is that there is little to no incremental gain in pushing toward the maximum. The compounding advantage achieved by moving from zero to the optimal minimum (Point C) is substantial; the additional gain from moving beyond it is vanishingly small, while the potential for instability rises sharply.

Time introduces the temptation. Because faster restoration shortens compounding delay, simulations show that reducing time-to-restoration produces slightly higher expected terminal wealth and a marginally higher annualized return. In that sense, maximum leverage appears efficient. However, due to the cash-harvesting rule that governs the restoration process, the cumulative effect is modest, not dramatic. The illusion of "faster" recovery conceals a nonlinear tradeoff between temporal efficiency and structural risk.

At approximately 45 percent leverage (three σ), the portfolio reaches the *margin of stability*: the point where the restorative power of leverage and the destabilizing force of variance are in near-perfect balance. A three-σ event—roughly a 45 percent volatility excursion—defines the upper bound of sustainable leverage within a 99.99 percent annual confidence interval (≈ 2–3 percent cumulative probability over 30 years of monthly observations). Beyond this threshold, the probability of breach or forced deleveraging rises sharply. Expressed on an equity basis, 45 percent debt-to-assets corresponds to 0.45 / 1.45 ≈ 31 percent equity leverage.

This boundary, however, is not universal. The precise location of Points C and D depends on multiple real-world constraints and definitional choices. "Failure" may mean a margin call to one investor, a liquidity shortfall to another, or complete capital loss to a third. Allowable leverage also varies across asset classes, regulatory frameworks, and funding conditions. Institutional portfolios with stable credit and low volatility may tolerate higher leverage, while individual investors face tighter behavioral, liquidity, and collateral constraints.

Spiral Theory does not prescribe fixed numerical values for C or D; instead, it provides a *methodology* for identifying them within any given constraint regime. The principle remains constant: Point C denotes the minimum leverage that restores path efficiency without failure, and Point D marks the maximum leverage at which that same rule remains stable through time. Between them lies the narrow equilibrium band where financing acts as ballast rather than fuel—where leverage stabilizes rather than destabilizes the compounding path.

Together, Points C and D bracket the equilibrium zone of feasible leverage. Within this corridor—roughly 30 to 45 percent in this illustration—path efficiency dominates variance; beyond it, variance dominates compounding. The difference between them is structural, not incremental: Point C achieves restoration without triggering instability, whereas Point D marks the threshold where the risk of the leverage rule itself begins to compound. This relationship is dynamic. As volatility, liquidity, and institutional conditions evolve, both C and D shift continuously, but their separation—the curvature between them—remains the geometric hallmark of the Spiral: a narrow, resilient corridor in which leverage restores equilibrium instead of eroding it.

Point E — The Threshold of Irrationality

Point E marks the theoretical boundary where the restoration rule ceases to add value and begins, on expectation, to destroy it. Probabilistically, this is the point at which the expected losses from leverage equal the volatility drag that leverage was intended to offset. For example, with an expected return of 10 percent and a standard deviation of 20 percent, the volatility drag is roughly 2 percent, leaving a long-run compounded return of about 8 percent in the absence of financing. At Point E, a Monte Carlo simulation might show an 80 percent probability of full path restoration but a 20 percent probability of failure. In expectation, these opposing outcomes net to the same 8 percent return—but now with added tail risk. The portfolio's *expected value* may remain constant, yet its *expected utility* declines because the downside outcomes carry increasing weight.

Point E therefore represents the *threshold of irrationality*: the mathematical line where the restorative effect of leverage is exactly offset by the destructive potential of variance. Beyond this point, leverage does not enhance efficiency; it transfers value from the future to the present by shortening expected survival time.

Point F — The "50/50 Odds" Frontier

Point F is the theoretical position where the probability of success and failure are approximately equal—about a 50 percent chance of restoration and a 50 percent chance of long-term ruin. It is important precisely because it exposes the divergence between static and dynamic models of risk.

A point-in-time model, such as that implied by the linear Capital Allocation Line, would label Point F as attractive: it appears to offer higher expected return with less perceived variance, at least instantaneously. Spiral Theory, however, reframes this illusion. Once compounding and sequential risk are accounted for, a 50/50 outcome is *not* efficient—it is irrational. In this zone, the balance sheet no longer functions as a stabilizer but as a speculative amplifier. Time converts what seems like a favorable short-term position into an almost certain long-term loss.

Point G — Collapse and Certain Failure

Point G represents the extreme—where the probability of failure approaches 1. It is the minimum amount of leverage that guarantees eventual insolvency with near certainty. Intuitively, this occurs at leverage ratios high enough that even moderate volatility leads to unrecoverable losses. Although extreme examples—such as 90 percent debt-to-assets—make the point obvious, the boundary of certain failure often lies well within that zone.

For instance, applying 80 percent leverage to a portfolio with an expected return of 10 percent and a standard deviation of 20 percent does *not* produce an 18 percent expected return. Over time, simulations show that such a policy has roughly a 99 percent probability of failure. The apparent arithmetic gain collapses under compounding because even small downward shocks deplete equity faster than they can be restored.

Point G therefore defines the terminal limit of the spiral—the point at which variance completely dominates compounding. Beyond it, no restoration rule can function; the system exits the realm of finance and enters the physics of collapse.

Revisit Point B — Partial Restoration

If Point A represents complete inertia—no restoration at all—and Point C corresponds to the expected return of the tangency portfolio under a full restoration rule with minimum leverage, then Point B lies between them. Point B represents *partial restoration*: a restoration process with an expected return below full recovery but with meaningfully reduced volatility drag.

Consider an example with $A = 10\%$ and $\sigma = 20\%$. The long-run compounded return without restoration is approximately $R_{(\text{geom})} = 8\%$. There are at least two ways to begin with this same arithmetic expectation of 8 percent.

1. Holding 100 percent of the risky asset (the baseline portfolio) produces this outcome in expectation but compounds below it over time—its realized return converges toward Point A.
2. Holding a mix of 80 percent risky assets and 20 percent risk-free assets (Point A_2) begins with the same arithmetic expectation but ultimately compounds at an even lower rate.

In both configurations, the portfolio fails to realize its expected 8 percent because volatility drag remains unopposed. Spiral Theory visualizes this shortfall: without restoration, the compounding path drifts below the arithmetic frontier, and efficiency is lost.

Once restoration enters, the geometry changes. A new triangle appears, bounded by Points A, B, and C. Point B can be reached in three distinct ways:

1. As a 90/10 mix of risky and risk-free assets plotted on the expected Capital Allocation Line (CAL).
2. By applying a restoration rule to a 90/10 portfolio—though in practice, cash and leverage offset each other, making this configuration more illustrative than effective.
3. Most importantly, by applying the restoration rule of Point A under a leverage cap: the leverage ratio of Point C reduced by, in this example, 10 percent. In this third case, cash and debt dynamically offset.

This third path—the *partial-restoration rule*—yields less risk and more expected return than the static tangency portfolio on the expected CAL. Within our example, a Monte Carlo simulation showing a 90 percent probability of full recovery and a 10 percent probability of incomplete recovery implies an expected return of about 9 percent, achievable with a leverage cap near 15 percent rather than the 30 percent used for full restoration.

Here, leverage is employed selectively after adverse deviations but within a defined ceiling. The portfolio borrows just enough to accelerate recovery without fully eliminating volatility drag. The leverage cap simultaneously reduces long-term risk and increases the long-term expected return. Capping leverage introduces a measurable trade-off among time, risk, and probability of recovery. The resulting trajectory is not linear but curved: small amounts of leverage slightly increase point-in-time risk while substantially improving long-term path efficiency. For instance, a 1 percent episodic leverage cap cannot restore full compounding efficiency, yet under Gaussian assumptions it introduces virtually no statistical long-term risk. Because the tangency portfolio cannot fall by 98 percent, the marginal probability of catastrophic loss from such minimal restoration is effectively zero.

Numerically, a maximum 10 percent leverage ratio applied ex post to a portfolio with an expected return of 10 percent and a standard deviation of 20 percent produces negligible long-term failure risk but only partial path restoration—yet 10 percent is materially more effective than 5 percent. Within this zone, Spiral Theory models a continuum of outcomes in which borrowing is used proportionally and prudently: enough to bend the compounding path upward but not enough to destabilize it.

Point B therefore represents the trade-off between the leverage ratio implied at Point C and the zero-leverage baseline of Point A. This relationship is nonlinear; one cannot simply interpolate between them. The line connecting A–B–C exhibits pronounced curvature because the relationship between leverage, volatility, and compounding is *multiplicative*, not additive. Five percent leverage is not "half as effective" as ten percent, nor is fifteen percent simply "one-and-a-half times" as effective. Expected return scales linearly with leverage, while variance scales with leverage squared. This convexity in risk combined with linearity in return produces a nonlinear path of geometric outcomes. Small increments of leverage can yield disproportionately large gains in compounding efficiency up to a threshold—beyond which additional leverage produces diminishing, and eventually negative, effects.

The Rational Zone (Points A–D)

Points A through D define the *rational zone* of the Spiral—the region in which leverage functions as a stabilizing force rather than a speculative one. Points along A–B–C can be viewed as representing probabilities of path restoration. At Point C—the minimum leverage required for full restoration—the probability of success approaches one. All points to the left of C involve some degree of failure risk, but with proportionally lower exposure.

This zone is critical because, across nearly all modeled conditions, *some* leverage—applied selectively and proportionally—restores compounding efficiency far more effectively than maintaining a static, unlevered position. The leap from no restoration to partial restoration is dramatically more powerful than the incremental gains from partial to full. The geometry of this region captures that asymmetry.

The line connecting A–B–C is nonlinear because the relationship between leverage, volatility, and compounding is multiplicative, not additive. Modest leverage produces disproportionately large improvements in expected wealth trajectories, while excessive leverage yields rapidly diminishing benefits and rising instability.

Between Points C and D lies the boundary of stability: leverage remains a rational tool but must be constrained by the physics of variance. Beyond D, the stabilizer becomes the source of instability.

A Simplified Expression: The 25 Percent Restoration Rule

In many ways, Spiral Theory collapses to a remarkably simple principle: *down is bad*.

When realized wealth falls meaningfully below its expected compounding path, restoration—not passivity—is the rational response. Empirically, when wealth declines by approximately 25 percent relative to its expected trajectory, borrowing 30 percent of portfolio value is sufficient to restore compounding efficiency under standard assumptions ($\mu = 9\%, \sigma = 15\%$).

When realized wealth exceeds expectations, the same policy operates in reverse: debt is repaid and surplus cash is harvested. This symmetry defines the discipline of the Spiral—exposure expands only when outcomes deteriorate and contracts when they improve. The result is a proportional response that keeps the portfolio centered around its expected growth path without forecasting or market timing.

In practice, this "25 percent restoration rule" translates the continuous identity of Spiral Theory into an actionable framework: **small, bounded leverage applied *after* adverse outcomes and symmetric harvesting applied *after* favorable ones**.

The process is dynamic but rule-based, operating entirely ex post and therefore compatible with real-world constraints. Fat-tail risk remains—large, non-Gaussian events can still breach modeled bounds—but these scenarios are addressed separately in the discussion of tail resilience and stress testing.

Metric	Value	Interpretation
Mean terminal deviation vs geometric path	–2.39%	Slight under-restoration due to time lag
Median terminal deviation	–1.02%	Near-perfect path efficiency at the median
P(	ΔW/W*	≤ 1%)
P(	ΔW/W*	≤ 3%)
Average terminal D/A	13.2%	Moderate, well-contained leverage at horizon
Paths ending with debt	34.7%	Two-thirds fully delevered by year 30
Mean terminal cash	742	Substantial liquidity accumulation
Average downside interventions	83.8	Episodic restoration (~3 per year)
Average upside interventions (harvests)	194.1	Frequent deleveraging/harvests (~6.5 per year)
Share of paths borrowing at least once	99.7%	System dynamically engages leverage
Share of paths harvesting	100%	Consistent upside discipline

Figure 10 - Rule Test: Conditional Restoration and Harvest (μ = 9%, σ = 15%, 30 years)

This rule test operationalizes Spiral Theory's central proposition: *bounded episodic leverage can restore geometric efficiency*. When portfolio wealth falls 25% below its expected compounding path, leverage up to 30% of assets is applied (cash-first, then debt). When wealth exceeds expectations, all debt is repaid and a fraction of gains (α = 10%) is harvested to cash. The result is a self-correcting dynamic that stabilizes compounding through time.

Over 10,000 thirty-year Monte Carlo paths (μ = 9%, σ = 15%), the strategy's median terminal wealth ended **within roughly one percent of the geometric benchmark**, with bounded drawdowns and modest long-run leverage (≈ 13%). The average portfolio accumulated large cash buffers, repaid debt in most scenarios, and maintained a stable balance sheet across cycles. These results illustrate that limited, rules-based borrowing combined with symmetric deleveraging can neutralize volatility drag in expectation—without breaching realistic institutional risk limits.

The finding is striking: **path efficiency can be achieved without full restoration or high leverage**. The stabilizer behaves like an endogenous "gearing keel," correcting deviations over time while remaining largely self-funded. The median outcome's one-percent proximity to the theoretical growth path demonstrates that the rule captures the essence of compounding restoration, translating Spiral Theory's mathematics into an implementable policy.

10. Real World Implications: Outcome-Based Planning

Modern Portfolio Theory and its descendants begin with the concept of *risk tolerance*. Investors are grouped into static models—60/40, 70/30, 80/20—based on their subjective willingness to accept volatility. Portfolio adjustments, framed as "rebalancing," restore these static weights rather than advance toward a measurable objective. Spiral Theory reframes this foundation.

In a time-based framework, the relevant variable is not *tolerance for risk* but the *quantity of assets relative to the expected return of the tangency portfolio.* The goal defines the risk; the portfolio merely solves for it. Each investor can be located precisely on the capital allocation line according to two parameters: (i) the quantity of money already accumulated, and (ii) the expected return of the tangency portfolio. As the quantity of money changes, the optimal position shifts dynamically and continuously—independent of market noise or personal risk preferences.

Any goal—whether defined as a retirement income, an endowment-style spending rate, or a future purchase—can therefore be expressed as a *required quantity of the tangency portfolio.* The most efficient path is to build that quantity first, then move *down* the capital allocation line by the reserve necessary to cross one's stability threshold. That threshold may correspond to Gaussian expectations or to the losses experienced in historical crises such as 2008 or the Great Depression.

Rebalancing thus follows a new rule: not to a *model portfolio*, but to a *target quantity* of the tangency portfolio. The reference point becomes absolute, not relative. This approach disrupts conventional asset-allocation logic, yet under Spiral Theory no internally consistent model can demonstrate it as inferior. It converts asset allocation from a behavioral heuristic into a solvable dynamic equation: each incremental change in wealth redefines the optimal mix of tangency exposure and reserve.

Empirical exploration should examine three dimensions:

1. **Comparative efficiency:** Measure how goal-based allocation performs relative to fixed-weight models in terms of time to restoration, drawdown frequency, and goal achievement.
2. **Reserve optimization:** Quantify the marginal benefit of additional reserves (cash or bonds) in extending path efficiency and mitigating sequence-of-returns risk.
3. **Implementation sequencing:** Test whether building the tangency exposure first, then the reserve, maximizes compounding efficiency and minimizes behavioral drag.

The implication is direct and profound. Outcome-based planning replaces the vocabulary of "risk tolerance" with the mathematics of solvable geometry. Each goal defines its own efficient point on the capital allocation line; asset allocation becomes a continuous process of maintaining sufficiency rather than restoring proportion. Risk, in this framing, is no longer the variance of returns—it is the probability of failing to meet an outcome.

11. Conversion & The Shiller Rule

The preceding discussion addressed transitions from legacy allocations toward outcome-based equilibrium under conditions of sufficiency or surplus. A related but distinct challenge arises when the investor begins in a deficit state—when the current quantity of money falls short of that required to meet the defined goal. Spiral Theory formalizes this condition explicitly. If the model implies that the investor must hold 125% exposure to the tangency portfolio to reach the target outcome, should the allocation jump immediately from 60/40 to 125% leveraged? Theoretically possible; practically, implausible.

The issue is not merely behavioral. It is structural. In deficit conditions, the theoretical optimum—borrowing to restore the compounding path—confronts real-world constraints on collateral, borrowing limits, and psychological risk tolerance. Yet ignoring the deficit has its own cost: compounding inefficiency and deferred restoration. The practical question becomes: *how can the balance sheet be structured to approximate restoration without violating feasibility or stability?*

Empirical evidence from over 1,000 household balance sheets suggests that meaningful progress can be achieved through indirect deficit restoration—most notably through the intelligent use of housing finance. For many families, home equity represents the single largest, least-productive asset on the balance sheet. Spiral Theory implies that *taking on a mortgage while holding the proceeds in cash or reserve assets* can improve long-term path efficiency by introducing optionality and smoothing consumption. Though counterintuitive, such "structural borrowing" increases flexibility and may reduce long-term failure probability relative to rigid debt aversion.

Comprehensive balance-sheet strategies therefore merit systematic exploration. Key research questions include:

1. **Deficit restoration pathways:** Under what conditions does partial leverage—via mortgage or line of credit—replicate the compounding path implied by full leverage?
2. **Collateral optimization:** How can illiquid home equity be converted into productive option value without inducing instability or consumption overreach?
3. **Dynamic sequencing:** What are the time-based rules for increasing or reducing deficit exposure as wealth converges toward sufficiency?

The analytical lineage of this problem traces to Robert Merton's *Intertemporal Capital Asset Pricing Model* (ICAPM), in which the optimal allocation shifts continuously with wealth, labor income, and time horizon. Spiral Theory shares this intuition but extends it from a frictionless, continuous framework to a discrete, path-dependent one. Where Merton models smooth adjustments in an idealized economy, Spiral Theory models episodic adjustments within a world of collateral limits, tax frictions, and behavioral thresholds.

The two frameworks converge on the same insight: *wealth, not risk tolerance, is the state variable that determines allocation.* But they diverge in implementation. Merton's share is continuous and endogenous to an ideal market; the Spiral share is discontinuous and constrained by financing conditions. The former describes how an investor *should* behave in an unconstrained setting; the latter describes how an investor *can* behave within the realities of balance-sheet management.

The hypothesis is that a properly structured balance sheet—combining mortgage leverage, reserve cash, and restoration rules—can approximate the compounding trajectory implied by Merton's continuous solution while maintaining real-world feasibility. This synthesis of theory and implementation reframes personal finance as *household-level capital structure optimization.*

11. Additional Topics for Further Exploration

This section articulates practical implications that follow from Spiral Theory's time-based, constraint-aware framing. For each topic, I state the issue, a testable hypothesis, and the direction for empirical exploration. The aim is not exhaustive proof but a research agenda that translates the Spiral from geometry to implementation.

Composition of the tangency portfolio

In theory, there exists a single tangency portfolio—defined precisely by the intersection of the efficient frontier and the capital allocation line. In this two-asset world, all investors hold a combination of the risk-free asset and the tangency portfolio, scaled by preference or leverage. In practice, however, this concept is almost universally misunderstood. It is estimated that the vast majority of licensed professionals, and even many academics, conflate the tangency portfolio with familiar market benchmarks. Professors and practitioners alike often substitute the S&P 500 because it is convenient, accessible, and feels "representative."

That substitution is more consequential than it appears. It reintroduces idiosyncratic and home-bias risk into what should be a diversified global baseline, then misattributes the resulting variance to whatever overlay or rule is being tested. The working hypothesis here is that a true tangency proxy must be globally diversified and effectively free of idiosyncratic risk. A broad global equity index—on the order of several thousand companies—is the appropriate baseline, even if its expected return (≈ 9%) and volatility (≈ 15%) are somewhat lower than a concentrated U.S. index.

The research agenda should remain open rather than prescriptive. There is not one empirical tangency portfolio but a *universe* of potential candidates. Each defines a distinct surface against which the restoration rule can be tested. Key areas for inquiry include:

1. **Substitution effects:** Quantify how replacing a global basket with domestic large-cap equity distorts the C–D corridor (time to restoration, breach probabilities, and harvest cadence).
2. **Sleeve-level restoration:** Evaluate whether applying the rule at the asset-class or sleeve level—particularly when incorporating low-correlated assets such as gold, commodities, or digital stores of value—accelerates recovery or improves near-term stability.

3. **The role of fixed income:** Explore how including bonds within the tangency portfolio alters the compounding path. If, in theory, expected return falls from 9% to 8% but volatility drops from 15% to 10%, do superior outcomes emerge when viewed over time rather than by point-in-time Sharpe ratios? This inquiry tests whether lower arithmetic efficiency can, through compounding, produce higher geometric efficiency—and whether the restoration rule amplifies or neutralizes that advantage.

The objective is not to assert a single correct composition, but to understand how variations in the underlying definition of the tangency portfolio influence the behavior and efficacy of restoration over time.

Currency risk and funding currency

The world is multi-currency, and restoration ultimately occurs in the currency of consumption. Borrowing costs, term structures, and basis spreads differ across currencies, while post–gold-standard history is limited for many regimes. These differences shape both the cost and efficacy of restoration.

The working hypothesis is **portability**: Spiral Theory remains valid globally, provided that the risk-free asset and the restoration funding source are coherently defined relative to the investor's consumption currency. The theory itself is currency-agnostic; what matters is internal consistency between the unit of measurement, the cost of leverage, and the compounding path. A U.S. investor could, in principle, implement the rule using yen, euros, or a currency basket as the funding base—potentially capturing carry benefits or diversifying rate exposure—but would then assume additional exchange-rate and policy risks that must be modeled explicitly.

Empirical exploration should focus on three fronts:

1. **Funding asymmetry:** Quantify how cross-currency borrowing differentials (e.g., USD-funded vs. JPY-funded strategies) alter restoration speed, path variance, and breach probabilities.
2. **Policy interaction:** Model conditional rate responses—such as central-bank easing or tightening after shocks—and test how ex-post restoration implicitly leans into or against monetary stabilization.
3. **Crisis behavior:** Examine historical episodes in which domestic funding choices created unintended currency stances relative to the global tangency portfolio, magnifying or mitigating drawdowns.

The objective is to determine whether the framework's time-based restoration logic is *portable* across funding and consumption currencies, and to identify the conditions under which currency choice becomes a source of either stabilization or hidden leverage.

Fat tails and regime risk

Practitioners are rarely true Gaussian thinkers. They implicitly account for depressions, wars, and structural breaks—periods when correlations collapse and volatility becomes discontinuous. The concern is straightforward: in genuine tail states, leverage accelerates failure rather than fosters stability. If restoration depends on access to credit, liquidity, or functioning markets, extreme events can interrupt the mechanism precisely when it is most needed.

The working hypothesis is that Spiral Theory can accommodate fat tails through *rule-based adaptation*, not abandonment. By modestly adjusting parameters or cadence, the restoration process can remain viable under non-Gaussian conditions. The objective is not to eliminate tail risk—no model can—but to ensure that restoration behaves proportionally rather than pathologically when regimes shift.

Empirical work should examine several avenues:

1. **Proactive cash sweeps:** After sustained up-moves, systematic harvesting to cash or credit reserves (a Shiller-style anti-extrapolation rule) may reduce exposure before regime shifts.
2. **Low-dose restoration:** Following large shocks, applying small, incremental leverage (e.g., 10% after a 30% decline) may preserve convexity benefits while containing tail exposure.
3. **Structural liquidity:** Test the stabilizing value of cash and pre-committed credit facilities as buffers that sustain restoration capacity during stress periods.
4. **Institutional smoothing:** For spending institutions—endowments, foundations, pension plans—evaluate the substitution of borrowing for distributions during restoration windows to smooth path efficiency without forced asset sales.

The goal is to identify whether and how Spiral Theory's proportional restoration can remain self-stabilizing in environments characterized by fat tails, liquidity shocks, and regime transitions—maintaining its time-based coherence even when normality fails.

Options overlays

The baseline Spiral model omits explicit option value, yet its rule-based restoration dynamics implicitly create option-like exposures. When wealth rises far above the compounding path, the rule prescribes harvesting—a behavior economically equivalent to systematically selling calls. When wealth falls well below the path, the rule prescribes restoration through buying—functionally akin to selling puts. In both states, the investor is providing liquidity to the market and harvesting risk premia that are typically embedded in option pricing.

The hypothesis is that explicit, *state-contingent options overlays* can magnify this effect. By formalizing the restoration process in derivative terms, one can systematically monetize volatility. For instance, a policy that sells calls after a +25% advance or sells puts after a −25% decline synthetically extends the restoration rule into the derivatives domain. The premium collected in both directions can either (i) enhance expected return directly through alpha accumulation, or (ii) fund the purchase of deep out-of-the-money protective puts, converting harvested volatility into asymmetric tail protection.

Empirical exploration should consider the following directions:

1. **Premium efficiency:** Quantify the alpha potential of rule-based call and put selling relative to standard volatility-selling benchmarks (e.g., short-VIX or covered-call indices).
2. **Defensive reinvestment:** Test frameworks that recycle collected premium into systematic tail hedges—evaluating how many deep out-of-the-money puts can be funded and how this alters drawdown depth and recovery speed.
3. **Interaction with restoration:** Assess how derivative overlays interact with the baseline leverage rule—whether the synthetic exposure from option writing complements, substitutes, or amplifies physical restoration through borrowing.

The broader research agenda is to map the *universe of overlays* consistent with Spiral Theory's geometry: those that monetize volatility in surplus states, reinforce restoration in deficit states, and recycle option premium into capital preservation. This intersection between compounding, convexity, and credit may reveal that the Spiral is not only a geometric structure of risk and return, but also a framework for systematic volatility management.

General-equilibrium effects and adoption

A natural question arises for any systematic framework: what if everyone does this? One view predicts amplified volatility—an economy where every balance sheet responds to the same triggers, magnifying both rallies and drawdowns. Another predicts stabilization—as speculative, pro-cyclical leverage is replaced by rule-based, negatively correlated financing that responds inversely to price movements.

The working hypothesis is **stabilization at the margin**. Because restoration is activated by adverse deviations, aggregate adoption would tend to increase buying pressure in declining markets and reduce leverage during expansions. At the portfolio level, this converts idiosyncratic volatility into a systematic stabilizer; at the macro level, it ties the supply of financing to states of distress rather than exuberance. In principle, this dynamic could compress the amplitude of cycles and dampen volatility, though beyond some threshold, feedback effects may reappear.

Empirical exploration should proceed in three dimensions:

1. **Agent-based modeling:** Simulate heterogeneous investors with varying leverage caps, reaction speeds, and liquidity constraints to identify the adoption rate at which restoration ceases to be price-taking and begins to influence returns.
2. **Market-impact analysis:** Measure how rule-based restoration interacts with existing liquidity providers—does it compete with, replace, or complement market-making and central-bank stabilization functions?
3. **Cycle amplitude testing:** Examine whether partial adoption dampens volatility (through countercyclical re-leveraging) or amplifies it (through synchronized rebalancing), and whether thresholds of stabilization or instability can be empirically identified.

The broader question is whether Spiral Theory, at scale, becomes a self-stabilizing market mechanism or a new source of correlated behavior. The hypothesis remains that *at the margin*—within plausible levels of adoption—rule-based restoration reduces pro-cyclical flows, tightens the linkage between financing and adverse states, and improves the resilience of both portfolios and the system they inhabit.

Structural debt and housing finance

Traditional portfolio theory treats debt as an external financing choice, separable from portfolio optimization. In reality, long-dated, low-cost mortgage debt is the dominant liability on most household balance sheets and a defining feature of real-world path efficiency. Its cost is typically below the expected return of the tangency portfolio and, at times, even below the yield on cash. This persistent spread between borrowing cost and expected return represents a structural advantage that standard models ignore.

The working hypothesis is that mortgage debt, when integrated within Spiral Theory's framework, improves long-horizon path efficiency by introducing *option value*—the ability to restore the compounding path during adverse states. Because a home cannot be rebalanced, liquidity must exist elsewhere. Holding a cash reserve reframes cash not as a drag on return but as a *stored option*, capable of being exercised precisely when adversity creates the widest efficiency gap. Ironically, the more one worries about fat tails, the greater the value of maintaining mortgage leverage—because it preserves optionality and avoids premature balance-sheet rigidity.

Spiral Theory formalizes this relationship in three ways:

1. **Quantification:** It measures the option value of mortgage leverage as a function of volatility, funding spread, and compounding horizon.
2. **Rules for deployment:** It prescribes how and when structural cash reserves can be released—using adversity as the trigger—to restore the portfolio to its target path.
3. **Sequence of priorities:** It highlights the path-efficient order of operations: first build exposure to the tangency portfolio, then accumulate a buffer, and only after the buffer is established, consider debt reduction.

Empirical exploration should compare alternative configurations—"own home outright + 70/30 portfolio" versus "mortgage + structural cash reserve + restoration rule at the portfolio level"—to quantify differences in compounding, drawdown resilience, and recovery time. The objective is to determine whether integrating structural debt explicitly into the restoration framework leads to superior geometric efficiency and a more rational, rule-based sequencing of household financial decisions.

Debt–cash spreads and platform frictions

Funding costs are heterogeneous, even among investors with similar assets and risk profiles. Margin rates, collateral haircuts, and cash yields can differ by hundreds of basis points depending on the lending platform, relationship status, or account structure. Yet many investors implicitly assume that the quoted platform rate represents "the market," overlooking the degree to which balance-sheet architecture and custody choice shape their effective cost of capital.

The working hypothesis is that *spread optimization*—minimizing the differential between borrowing and cash yields—is a first-order driver of path efficiency. In Spiral Theory, this spread directly defines the slope of the feasible restoration corridor. Achievable improvements in funding cost can shift points

C and D—the regions of efficient leverage and restoration speed—more than most asset allocation changes within a conventional portfolio.

Empirical exploration should pursue three directions:

1. **Mapping the frontier:** Quantify how the efficient frontier shifts as a function of the debt–cash spread, holding expected asset return and volatility constant.
2. **Segment realism:** Document realistic spread bands by investor type—retail, advisory platform, institutional, and ultra–high-net-worth—to identify where access frictions are most binding.
3. **Restoration dynamics:** Simulate how narrowing the spread shortens time to restoration and expands the rational corridor—the range within which financing acts as a stabilizer rather than an amplifier of volatility.

The objective is to demonstrate that small, persistent differences in platform spreads compound into large differences in long-term outcomes. In a time-based framework, the cost of liquidity is not a secondary detail but a primary determinant of efficiency. Reducing the debt–cash spread expands the feasible domain of Spiral Theory, bringing more investors within reach of stable, rule-based restoration.

Tax-deferred accounts and implementability

Retirement and tax-deferred accounts often cannot be margined, prompting the objection that Spiral Theory is "useless" in those settings. In practice, this constraint reframes rather than invalidates the framework. Spiral Theory operates at the *household balance-sheet level*, not the account level. Structural debt outside a tax-advantaged wrapper can substitute for margin inside it, preserving the same functional exposure.

Millions of households follow the heuristic of *maximizing retirement contributions and paying down the mortgage*. Spiral Theory quantifies the hidden cost of that approach: the **perpetual volatility drag** that arises when all leverage is extinguished and all compounding is locked inside untouchable wrappers. The model challenges conventional tax assumptions by demonstrating that what matters is not tax deferral in isolation but *the interaction between taxation, compounding, and path efficiency through time*.

For example, if the core portfolio were a single, low-cost global equity index fund with no rebalancing, realized taxes would differ sharply from those in a typical multi-asset strategy. If, instead, the base portfolio were never sold and only long-term gains *above a defined target* were harvested, the effective tax rate on compounding would change again. Spiral Theory identifies these as quantifiable dimensions of volatility drag within both taxable and tax-deferred environments and reframes the underlying tax assumptions themselves.

Empirical exploration should focus on:

1. **Quantifying drag:** Measure the opportunity cost of full tax deferral relative to partial liquidity with restoration, under realistic spreads and statutory tax rates.

2. **Tax-structure modeling:** Simulate long-term outcomes across jurisdictions and account types, isolating how realization rules, capital-gain treatment, and deductibility of interest alter path efficiency.
3. **Integrated household policy:** Evaluate optimal combinations of retirement saving, mortgage structure, and restoration borrowing as a unified household policy rather than isolated tax or investment decisions.

Spiral Theory's broader implication is that taxation cannot be separated from time. The framework measures not only after-tax returns but the *temporal efficiency* with which those returns compound. By integrating financing, compounding, and taxation into a single system, it opens a new field of inquiry: how tax policy, asset structure, and household balance-sheet design jointly determine the real geometry of wealth over time.

13. The Initial Acceptance of Spiral Theory

Professional Feedback

As part of this research, the author presented the Spiral Theory framework to approximately 1,500 licensed financial professionals over a twelve-month period. The largest session hosted about 900 participants, while ten smaller workshops ranged from 20 to 50 attendees. Two additional webinars were conducted with unrecorded participant counts. The material was approved for continuing-education credit, reinforcing its professional legitimacy, though not all participants received formal credit.

The audience consisted primarily of elite practitioners—financial advisors among the top producers in their respective firms. More than 80% had previously delivered at least one lending solution to a client, yet fewer than 5%, and likely less than 1%, had a deliberate, forward-looking strategy for balance-sheet design. Within the sample, only one professional reported leading with cost of debt as a primary competitive advantage. Sessions were held in San Francisco, New York, Columbus, Houston, Tampa, Orlando, Kansas City, New Orleans, and Philadelphia. Larger audiences were lecture-style, while smaller ones were interactive, including a four-hour classroom session with thirty senior advisors focused on applied modeling and discussion.

The professional response was strikingly consistent. Once the stabilizer mechanism was understood—that the restoration rule is an identity with a time lag, not a forecast—the proof was rarely contested. The logic that borrowing can act as a negatively correlated asset, stabilizing compounding rather than amplifying risk, was viewed as both novel and persuasive. The notion that "you can't rebalance against your house" resonated immediately, as did the idea of holding mortgage proceeds in cash as an option reserve for restoring path efficiency.

However, the sessions revealed deep structural and cognitive barriers to implementation. Nearly all participants agreed that existing tools are inadequate. Presently, no major financial-planning software easily models portfolio decisions that incorporate debt as an endogenous variable. Debt decisions are treated as exogenous to outcomes—assumed static or irrelevant to risk-return optimization. Monte Carlo simulations, in which risk is assumed to increase with time, remain the industry standard. It is widely accepted that outcomes diverge over time; the idea that a disciplined, rule-based approach could *narrow* the distribution of outcomes over time was viewed as ideal in theory but foreign in practice—an "auto-pilot" toward convergence that most professionals had never considered possible.

Professionals expressed strong interest in tools that would allow them to "see" these dynamics—to visualize the interplay between leverage, compounding, and path efficiency. Yet the sessions also revealed an extraordinary disconnect between professional intuition and financial theory. Fewer than 5% of participants could clearly describe the capital allocation line, the tangency portfolio, or the logic of a two-asset world. Fewer still recognized the names Markowitz, Sharpe, Tobin, or Merton, despite their foundational role in the mathematics underlying every model used in practice.

This divergence between theory and profession is profound. Practitioners live in a world of path-dependent outcomes, behavioral constraints, and client goals. Academics, as the next section explores, remain anchored to point-in-time models. Spiral Theory sits squarely in the middle: it unifies these two worlds by showing how the time-based dynamics that govern real portfolios can be derived directly from the static geometry of Modern Portfolio Theory.

The professional feedback thus serves as both validation and challenge—affirming the intuitive appeal of the model while underscoring the need for education, visualization, and tools that connect the formal language of theory to the practical realities of financial planning.

The Gulf Between Professionals and Academics

Perhaps the most revealing finding of this research is not the reaction to Spiral Theory itself, but what that reaction exposed about the state of professional understanding. In the course of presenting to more than 1,500 licensed advisors, it became clear that the foundational principles of portfolio theory are not widely taught, and in many cases, not even recognized.

The evidence is striking. Neither the FINRA Series 7 nor the NASAA Series 65 or 66 examinations appear to require substantive mastery of Modern Portfolio Theory. "Portfolio analysis" is mentioned only briefly under suitability and diversification; there is no clear requirement to understand or apply the Capital Allocation Line, the tangency portfolio, or even the mathematical meaning of standard deviation. In practice, few advisors track the standard deviation of client portfolios, and fewer still use it as a measure of risk. It is not simply that advisors measure risk differently—they often do not know what "standard deviation" formally represents.

To define the terms that are missing from professional discourse: the Capital Allocation Line (CAL) represents all possible combinations of the risk-free asset and the tangency portfolio, tracing the linear relationship between expected return and risk at a point in time. The tangency portfolio is the unique, optimally diversified mix of risky assets that maximizes the Sharpe ratio—the slope of the CAL. It is the only portfolio on the efficient frontier that, when combined with the risk-free asset, produces a straight line of optimal choices for any investor, independent of preferences. The standard deviation measures the dispersion of returns around their mean; in this geometry, it is the axis of risk.

Yet in professional practice, these ideas are largely absent. Advisors commonly rely on model portfolios, Monte Carlo projections, or firm-supplied “risk scores” derived from questionnaires rather than from the mathematical structure of portfolio theory. They speak of diversification and volatility but without reference to the efficient frontier or the CAL. Many are unaware that, in the canonical model, *every* investor holds the same risky portfolio—the tangency portfolio—and differs only in their proportion of the risk-free asset.

The implications are profound. Advisors operate in a world of path-dependent outcomes, constrained by client behavior and real-world frictions. Academics, by contrast, remain anchored to point-in-time equilibria, analyzing instantaneous risk–return trade-offs under static assumptions. The result is a system bifurcated by perspective: one side sees volatility as an obstacle to behavior; the other, as a parameter in an equation.

This gulf is not merely academic—it is structural and existential. Professionals are entrusted with the practical stewardship of wealth, yet most have never been trained in the frameworks that define optimal behavior under uncertainty. Academics, meanwhile, refine models that assume away the very frictions—taxes, borrowing constraints, path dependence—that dominate reality. Spiral Theory sits precisely in the gap between them. It extends the geometry of Modern Portfolio Theory into time, where professionals live, and introduces the constraints and asymmetries that academics often ignore.

The findings of the past twelve months sound an alarm for both communities. For academics, the challenge is to return to the field—to model the world as it is, with credit, compounding, and constraint. For professionals, the challenge is to reclaim the foundations of the discipline—to understand not only what they do, but *why* the underlying theory demands it. The bridge between point-in-time optimization and over-time compounding is not optional; it is the discipline itself.

The Academic Response

The academic response was disciplined and precise. Scholars wield the point-in-time geometry of modern finance with great fluency, and their initial critiques were both clarifying and constructive. Three early conjectures proved incorrect as stated; each, once corrected, materially strengthened the framework.

First, the idea of a “shadow” capital allocation line created by low, structural leverage was rejected. The capital allocation line at a point in time does not shift upward merely because the investor carries leverage through time. The correction is straightforward: leverage does not generate a new line; compounding generates a lag. In expectation, volatility drag causes realized wealth to fall below the arithmetic forecast, so the *observed* long-horizon locus of risk–return is a time-lagged shadow of the instantaneous CAL, not a distinct frontier. Framed this way, the presence of “two lines” is no longer an ontological claim but a timing claim: one is the exact, point-in-time relation; the other is its geometric counterpart once compounding and variance have worked. The mathematics is standard and the interpretation uncontroversial.

Second, my initial hypothesis that "Tobin's line bends"—that 20% leverage is categorically different from 80% leverage even under Gaussian assumptions—was rightly rejected at a point in time. Under the canonical assumptions, the capital allocation line remains linear and exact. The implication is decisive for Spiral Theory: any departure from linearity must arise *over time*, not *at a time*. Spiral Theory is therefore an overlay on the temporal dimension: a rule that governs how the portfolio moves across successive points, not a modification of the instantaneous geometry. This is where superposition or Bayesian language has been both helpful and contentious. Conceptually, the portfolio occupies two descriptions: a point-in-time state with a well-defined mean and variance, and an over-time state in which compounding transforms those same parameters into a distribution of wealth paths. The controversy centers on the object of "risk." If risk is defined strictly as period standard deviation, the point-in-time model is complete. If, however, risk is the probability of failing to meet an outcome (a path property), then time enters as an essential variable, and rules that condition on past deviations acquire analytical standing. The present work adopts the latter definition for planning problems and retains the former for instantaneous pricing.

Third, the spiral image itself misled some readers. Considerable energy was spent searching for a closed-form "spiral formula." That search was misplaced. The figure is a map, not a single equation. It connects a finite set of solvable points—A through G—each defined by explicit constraints (e.g., do nothing; partial restoration; least-leverage restoration over a horizon; fastest restoration without value destruction; thresholds worse than doing nothing; 50/50 failure; ruin). Joined in parameter space, these points resemble a spiral, but the contribution is the characterization of the points, not the existence of a closed curve. Future work may produce a functional envelope; present rigor rests on the solvability of the points themselves.

The largest gap between academic priors and the present thesis concerned expectation through time. Many readers initially interpreted "restoration" as a single-period claim. That was never the hypothesis. The claim is that borrowing, applied episodically and proportionally, can offset volatility drag *in expectation, over time*. Once this is stated explicitly—holding the point-in-time CAPM/Tobin structure fixed; recognizing that the CAL "lags itself" under geometric compounding; treating the spiral as a set of constraint-defined points; and defining risk for planning as outcome failure rather than period variance—the foundation for academic receptivity improves markedly. The overlay does not contradict the instantaneous model; it completes the planning problem that the instantaneous model does not address.

In sum, the academic critique sharpened four commitments. First, preserve the point-in-time linearity of the capital allocation line without exception. Second, locate all nonlinearity in the temporal aggregation of returns, where compounding and variance interact. Third, replace the search for a closed-form spiral with the enumeration of solvable constraint points. Fourth, be explicit that "risk," in the planning sense, is a path property tied to goal attainment. With these clarifications, Spiral Theory is recast not as a rival to the standard geometry but as its time-based extension: a disciplined rule for moving through states in a way that narrows the distribution of outcomes while respecting the exactness of the underlying point-in-time model.

12. The Bridge: Integration with Canonical Theory

To situate the stabilizer rule within established theory, this section connects its mechanism to canonical frameworks in portfolio optimization and corporate finance. The aim is not to replace prior results but to show how they fit together once time and financing are made explicit.

From Allocation to Financing: The Missing Bridge to Modigliani–Miller

Modern Portfolio Theory is an asset-selection model: choose a risky sleeve and scale it along a straight capital-allocation line, with financing treated as exogenous. Modigliani–Miller is a financing model: in frictionless markets, leverage does not affect value, but in practice firms treat liquidity as an option and manage balance sheets episodically. Portfolio theory absorbed the first and largely ignored the second, leaving a gap between investment logic and the CFO's balance-sheet framework.

Spiral Theory closes that gap without disturbing either foundation. It keeps MPT's point-in-time results intact—no new frontier, no new Sharpe—and imports the corporate-finance insight that matters through time: financing flexibility is an option. A rules-based policy—borrow after adverse realizations, repay after favorable ones—neutralizes the compounding penalty that MPT leaves unpriced. Debt is not used to chase a steeper CAL; it keeps realized compounding aligned with the slope MPT already implies.

This reframes household and endowment choices excluded from traditional MPT. A fixed-rate, non-callable mortgage with liquid reserves is not "cash at a negative spread"; it is an option to adjust the quantity of money when it matters. Owning outright extinguishes that option; term financing plus cash preserves it—the same logic that governs corporate credit lines, operating cash, and maturity policy.

Spiral Theory therefore treats financing as a third control variable—alongside asset mix and spending—showing how episodic, rules-based debt can restore compounding efficiency under the same assumptions that support both Modigliani–Miller separation and mean–variance analysis. Where financing is costly, callable, or constrained, the option's value is bounded and the uplift attenuates; where funding is reliable and cheap relative to expected returns, the corporate-finance logic carries through cleanly.

Together the models interlock: **MPT defines the sleeve, Modigliani–Miller explains leverage neutrality, and Spiral Theory supplies the through-time financing policy that preserves expected compounding in realized outcomes.**

Connection to other theories

Kelly and log-utility. Kelly maximizes expected log-wealth by choosing a constant exposure and continuously rebalancing. Its objective is optimal growth given a utility criterion. Spiral Theory's objective is different: it offsets the arithmetic–geometric gap for a given risky sleeve so that the long-run geometric return equals the sleeve's arithmetic expectation. The policy acts *ex post*—after deviations—not through continuous, predictive control. Under the canonical assumptions, the identity concerns eliminating a compounding penalty rather than maximizing expected utility. The two coincide in special cases but are not equivalent.

Volatility harvesting and rebalancing. The so-called rebalancing premium arises from diversification and convexity across imperfectly correlated assets. It generally presumes some mean reversion or cross-asset dispersion. Spiral Theory requires neither. It operates on a single risky sleeve and the risk-free asset, where the borrowing option alone offsets volatility drag in expectation under i.i.d. Gaussian returns. Any additional rebalancing premium from multiple assets is orthogonal to the variance-offset channel examined here.

Volatility targeting and risk parity. Volatility targeting scales exposure *ex ante* based on realized or forecast volatility to stabilize prospective risk. Spiral Theory instead conditions on realized return shortfalls relative to expectation and adjusts *ex post* through financing. The two can coexist, but their mechanisms and state variables differ: one manages variance, the other manages compounding.

CPPI and floor-based methods. Constant-proportion portfolio insurance (CPPI) preserves a floor by reducing risky exposure after losses—selling into drawdowns. Spiral Theory does the opposite on the financing margin. After losses it adds financed exposure to offset the compounding penalty, subject to constraints. CPPI's objective is capital preservation; Spiral Theory's is restoration of long-run compounding efficiency.

Market timing and forecasting. Spiral Theory makes no forecasts. Actions are triggered mechanically by realized deviations from the expected path. Borrowing follows the drag; any uplift is an identity of compounding, not a prediction premium.

Relation to the Merton Share. In Merton's intertemporal framework with stable investment opportunities and standard risk aversion, the investor holds a fixed fraction of wealth in the risky portfolio. That fraction—the Merton share—rises with expected excess return and falls with volatility and aversion. Spiral Theory leaves this decision untouched. Conditional on whatever risky share the investor or planner selects, the episodic financing rule acts only on the *financing margin*, after the fact, to offset the compounding penalty that otherwise causes realized growth to fall below the model's expectation.

The rule does not alter portfolio composition or depend on forecasts. It borrows proportionally after adverse realizations and repays after favorable ones so that exposure remains near the chosen share while geometric efficiency is restored. In practice, borrowing costs and leverage limits bound this overlay; a utility-consistent implementation sets those bounds so average exposure remains near the target.

Spiral Theory therefore complements rather than contradicts Merton's structure. It preserves the optimal risky share as determined by preference and expectation, while adding a time-based mechanism to reconcile the arithmetic and geometric dimensions of compounding. In this sense, it extends the logic of continuous-time optimization into the discrete, frictional world where investors actually live, providing a rule for maintaining equilibrium through time rather than only at a point in time.

Merton chooses how much risk to hold; Spiral Theory shows how to finance that choice through time so that realized compounding converges, in expectation, to the model's return.

13. The Open Door: From Divergence to Design

Realized geometric returns systematically fall short of arithmetic expectations because variance compounds against itself. The stabilizer rule—borrowing when behind, deleveraging and building cash when ahead—closes this gap in expectation. Under canonical assumptions, this is not conjecture but identity: the long-run geometric return converges to the arithmetic mean, $A_r = E_r$. Large-scale Monte Carlo simulation confirms the result; volatility drag disappears, and realized compounding restores the capital allocation line's expected slope.

More precisely, while point-in-time risk remains unchanged, risk through time fades. In the buy-and-hold world, dispersion grows linearly with horizon: $Var[\log{(W_t/W_t^*)}] \propto t$. Under the stabilizer rule, the cross-sectional dispersion of $W/W^*$is stationary. Each intervention resets the risky notional to the expected path $W^*(t)$, transforming compounding into a sequence of independent, single-period deviations with fixed variance. The process "forgets" prior shocks. Accuracy becomes a property of cadence (Δ), not horizon. In empirical terms, mean absolute and root-mean-square percentage errors remain flat over time; the tracking band does not widen. This reconciles the theorist's belief that wealth converges to the average with the practitioner's observation that wealth diverges—both are correct, but under the stabilizer, convergence dominates.

The results do not end the discussion—they begin it. If a simple rule can close the gap between arithmetic and geometric growth, then something fundamental has been found. The shortfall long treated as an unavoidable fact of randomness can, at least in theory, be offset by using credit as a stabilizer rather than a wager. In one sense, that is the paper's central result: borrowing when behind and repaying when ahead has measurable value. In another, it points to something larger. What we call "risk" may partly reflect how financing is managed through time. Volatility drag is not just a cost of variance—it is a policy variable.

The implications are profound. If risk through time can be tamed by design, then portfolio management becomes a problem of financing cadence rather than fixed allocation. This implies a new taxonomy of investors. For return-target investors—such as endowments or sovereign funds—the stabilizer acts as a compounding keel, maintaining exposure and shortening the time to target. For goal-based investors—such as retirees, insurers, and defined-benefit plans—the same mechanism stabilizes confidence: instead of cutting risk after losses, they can use controlled debt to preserve trajectory. In both cases, debt becomes a policy instrument rather than a speculative variable, and the quantity of money emerges as an endogenous dimension of optimization alongside mean and variance.

This reconceptualization places individuals and institutions squarely within the Modigliani–Miller frame: value depends on assets and liabilities together, managed through time. Financing flexibility offsets volatility drag just as corporate leverage arbitrage offsets capital-structure noise. In this sense, Modern Portfolio Theory is not wrong—it is incomplete. Its point-in-time truths remain intact, but its temporal application requires a stabilizer to prevent divergence between arithmetic and geometric growth.

The theoretical and empirical findings together establish a new class of balance-sheet policies—rules that stabilize compounding across horizons. They unify asset allocation, spending, and financing into a single design problem: goal → sleeve → buffer → preferences. Once the financing cadence is chosen, risk ceases to accumulate geometrically and becomes stationary, bounded by the intervention interval.

The next step is not prescription but measurement. How much of this effect survives once borrowing costs, taxes, or institutional limits are introduced? Every theory that meets the real world must pass through the same narrowing funnel—what is possible in theory becomes what is permissible in practice. If even part of the stabilizing effect survives, then personal, institutional, and corporate finance can be governed by a common law of compounding: the disciplined use of credit transforms variance from an adversary into an asset.

Thus the paper concludes where finance begins anew. Volatility drag, long seen as randomness's tax on compounding, is revealed instead as a controllable variable—a matter of cadence and policy. On the other side of this door lies a universe of design: how to structure balance sheets, endowments, and retirements so that risk fades through time while growth remains. The mathematics is finished; the exploration has just begun.

Appendix A: Doctrine & Theory

Definitions. By household I mean the single steward of a portfolio. By volatility drag I mean the shortfall of compounded from expected arithmetic return. By episodic financing I mean the practice of borrowing after adverse realizations and repaying after favorable ones, within limits.

Assumptions. Prices follow a stationary process with finite second moments; borrowing at the risk-free rate is available within bounds; taxes and frictions may be set to zero or specified.

Proposition 1 (separation through time). Given the assumptions without frictions, episodic financing restores the long-run geometric growth of any chosen risky sleeve to its arithmetic expectation. At a date, financing is indifferent; through time, it preserves the promise of the one-period model.

Sketch. The compounding shortfall equals a function of variance; the financing rule contributes the same in expectation by acting only after realized deviations. Hence restoration.

Proposition 2 (primacy under frictions). With spreads, caps, taxes, collateral, and non-tradable risks, household capital structure becomes a first principle: the option value of liquidity is high, constraints bind early, and outcomes depend on financing policy.

Corollary. The portfolio choice (what to hold) and the financing policy (how to act through time) are distinct arts. The former is given by mean–variance; the latter, by admissible bounds that keep ruin improbable. Where bounds are respected, the rule stabilizes compounding; where they are not, certainty approaches but welfare declines.

Replication Code: Volatility Drag

```
# Spiral Theory — (Test 0: Baseline Volatility Drag)

import numpy as np, pandas as pd

seed   = 42
mu     = 0.09
sigma  = 0.15
dt     = 1/12
years  = 30
n_paths= 10_000
S0     = 100.0

rng = np.random.default_rng(seed)
n_steps = int(years / dt)
Z = rng.normal(size=(n_paths, n_steps))
drift = (mu - 0.5*sigma**2)*dt
vol   = sigma*np.sqrt(dt)

S = np.empty((n_paths, n_steps+1))
S[:,0] = S0
for t in range(1, n_steps+1):
    S[:,t] = S[:,t-1]*np.exp(drift + vol*Z[:,t-1])

# Expectations
tgrid = np.arange(n_steps+1)*dt
arith_expect = S0*np.exp(mu*tgrid)
geom_expect  = S0*np.exp((mu - 0.5*sigma**2)*tgrid)

terminal = S[:,-1]
results = dict(
    mean_terminal_price = np.mean(terminal),
    median_terminal_price = np.median(terminal),
    mean_vs_arith_pct = (np.mean(terminal)/arith_expect[-1] - 1)*100,
    median_vs_geom_pct = (np.median(terminal)/geom_expect[-1] - 1)*100
)

df = pd.DataFrame([results])
print("\nAppendix D.2 — Test 0: Baseline Volatility Drag\n")
print(df.to_string(index=False))
```

Replication Code: Two-Sided Keel

```
# Spiral Theory — Test 1 (Unbounded, Costless, Cash-First, Proportional
Restoration)
# Monte Carlo simulation relative to the geometric expected path

import numpy as np
import pandas as pd

# -----------------------
# Parameters
# -----------------------
seed    = 42
mu      = 0.09                 # annual arithmetic mean return
sigma   = 0.15                 # annual volatility
dt      = 1/12                 # monthly step
years   = 30
n_paths = 10_000
S0      = 100.0
friction = 0.0                 # costless as requested (keep at 0.0)

# -----------------------
# GBM simulation (lognormal prices; normal log returns)
# -----------------------
rng = np.random.default_rng(seed)
n_steps = int(years / dt)

Z     = rng.normal(size=(n_paths, n_steps))
drift = (mu - 0.5 * sigma**2) * dt
vol   = sigma * np.sqrt(dt)

S = np.empty((n_paths, n_steps + 1))
S[:, 0] = S0
for t in range(1, n_steps + 1):
    S[:, t] = S[:, t-1] * np.exp(drift + vol * Z[:, t-1])

# Geometric expected (compounding) benchmark: W*_t = S0 * exp((mu -
0.5*sigma^2) * t)
tgrid  = np.arange(n_steps + 1) * dt
Wstar  = S0 * np.exp((mu - 0.5 * sigma**2) * tgrid)

# ---------------------------------------
# Policy simulator: proportional restoration, cash-first, unbounded, costless
# ---------------------------------------
def run_policy(policy_type, delta=None, interval_months=None):
```

```
    """
    Proportional restoration (the 'keel'):
      phi_t = W*_t / W_t
      target risky notional E_target = phi_t * W_t  (baseline sleeve = 100% of
equity)
      current risky notional E_curr = q_t * S_t
      trade notional ΔE = E_target - E_curr; Δq = ΔE / S_t

    Financing hierarchy:
      - If ΔE > 0 (buy): use cash up to ΔE; borrow the residual (unbounded).
Equity unchanged.
      - If ΔE < 0 (sell): sell risky; repay debt first up to proceeds;
leftover to cash. Equity unchanged.

    Intervention schedule:
      - 'time': intervene every k months (1, 3, or 12).
      - 'trigger': intervene when |(W_t - W*_t) / W*_t| >= delta, i.e., wealth
deviation from geometric path.
    """
    # State variables
    risky_units   = np.full(n_paths, 1.0)   # start fully invested: q_0 = 1
unit
    cash          = np.zeros(n_paths)
    debt          = np.zeros(n_paths)
    wealth        = np.full(n_paths, S0)    # q*S + cash - debt
    interventions = np.zeros(n_paths, dtype=int)

    for t in range(1, n_steps + 1):
        # Mark holdings to market with realized return
        risky_units *= (S[:, t] / S[:, t-1])
        risky_value = risky_units * S[:, t]
        wealth      = risky_value + cash - debt

        # Decide whether to intervene
        if policy_type == 'trigger':
            dev = (wealth - Wstar[t]) / Wstar[t]
            do  = np.abs(dev) >= delta
        elif policy_type == 'time':
            do  = np.full(n_paths, (t % interval_months == 0))
        else:
            raise ValueError("policy_type must be 'trigger' or 'time'")

        if not np.any(do):
            continue
```

```
        idx = np.where(do)[0]
        St  = S[idx, t]
        Wt  = wealth[idx]
        qt  = risky_units[idx]

        # Proportional restoration multiplier and target risky notional
        phi      = Wstar[t] / Wt                      # gross exposure
multiple vs equity
        Etarget  = phi * Wt                           # target risky notional
(baseline sleeve = 100%)
        Ecurr    = qt * St
        dE       = Etarget - Ecurr                    # trade notional
        buy      = dE > 0
        sell     = dE < 0

        # --- Buy: cash-first, then debt (unbounded) ---
        if np.any(buy):
            i = idx[buy]
            need = dE[buy]
            # units to purchase
            risky_units[i] += need / S[i, t]
            # finance
            use_cash = np.minimum(cash[i], need)
            cash[i] -= use_cash
            debt[i] += (need - use_cash)
            # optional friction (inactive at 0.0)
            if friction != 0.0:
                cost = friction * need
                use_for_cost = np.minimum(cash[i], cost)
                cash[i] -= use_for_cost
                debt[i] += (cost - use_for_cost)

        # --- Sell: repay debt first, excess to cash ---
        if np.any(sell):
            i = idx[sell]
            proceeds = -dE[sell]
            risky_units[i] += dE[sell] / S[i, t]      # dE<0 reduces units
            # optional friction (inactive at 0.0)
            net = proceeds
            if friction != 0.0:
                net -= friction * proceeds
            repay = np.minimum(debt[i], net)
            debt[i] -= repay
            cash[i] += (net - repay)
```



```
        # Post-trade wealth for intervened paths (equity unchanged by
financing)
        risky_value = risky_units[idx] * St
        wealth[idx] = risky_value + cash[idx] - debt[idx]
        interventions[idx] += 1

    # Terminal deviation relative to geometric benchmark
    terminal_dev = (wealth - Wstar[-1]) / Wstar[-1]
    return dict(
        mean_terminal_dev_pct   = float(np.mean(terminal_dev) * 100),
        median_terminal_dev_pct = float(np.median(terminal_dev) * 100),
        pct_within_pm1          = float(np.mean(np.abs(terminal_dev) <= 0.01)
* 100),
        pct_within_pm3          = float(np.mean(np.abs(terminal_dev) <= 0.03)
* 100),
        avg_interventions       = float(np.mean(interventions)),
        median_interventions    = float(np.median(interventions))
    )

# -------------------------
# Run the requested experiments
# -------------------------
rows = []

# Time-based: monthly, quarterly, annual
for months, label in [(1, "monthly"), (3, "quarterly"), (12, "annual")]:
    r = run_policy('time', interval_months=months)
    r['policy'] = f"Time-based ({label})"
    rows.append(r)

# Trigger-based: ±5%, ±10%, ±15% (wealth deviation vs geometric path)
for delta in [0.05, 0.10, 0.15]:
    r = run_policy('trigger', delta=delta)
    r['policy'] = f"Trigger-based (±{int(delta*100)}%)"
    rows.append(r)

df = pd.DataFrame(rows)[[
    'policy',
    'mean_terminal_dev_pct',
    'median_terminal_dev_pct',
    'pct_within_pm1',
    'pct_within_pm3',
    'avg_interventions',
    'median_interventions'
]]
```



```
print("\nAppendix D.2 — Unbounded, Costless, Cash-First, Proportional
Restoration\n")
print(df.to_string(index=False))

# Save CSV for convenience
df.to_csv("Appendix_D2_unbounded_costless_cashfirst_proportional.csv",
index=False)
print("\nSaved: Appendix_D2_unbounded_costless_cashfirst_proportional.csv")
```

Replication Code for a Simple Rule

```
# Spiral Theory — Rule Test with Cash Harvest:
# "If W <= 0.75·W*: set D/A = 30% (cash-first); if W > W*: repay any debt,
then harvest α of (W - W*) to cash"
# GBM, monthly steps, costless trading.

import numpy as np
import pandas as pd

# ------------------------
# Parameters
# ------------------------
seed        = 42
mu          = 0.09          # annual arithmetic mean return
sigma       = 0.15          # annual volatility
dt          = 1/12          # monthly step
years       = 30
n_paths     = 10_000
S0          = 100.0
friction    = 0.0           # keep at 0.0 (costless)
dev_trig    = -0.25         # downside trigger: wealth ≤ (1 - 25%) of expected
path
DA_target   = 0.30          # post-trade Debt/Assets target when downside
trigger fires
harvest_a   = 0.10          # upside harvest rate α: sell α * (W - W*) to cash
when W > W*

# ------------------------
# GBM simulation (lognormal prices; normal log returns)
# ------------------------
rng = np.random.default_rng(seed)
n_steps = int(years / dt)
```

```
Z     = rng.normal(size=(n_paths, n_steps))
drift = (mu - 0.5 * sigma**2) * dt
vol   = sigma * np.sqrt(dt)

S = np.empty((n_paths, n_steps + 1))
S[:, 0] = S0
for t in range(1, n_steps + 1):
    S[:, t] = S[:, t-1] * np.exp(drift + vol * Z[:, t-1])

# Geometric expected (compounding) benchmark
tgrid = np.arange(n_steps + 1) * dt
Wstar = S0 * np.exp((mu - 0.5 * sigma**2) * tgrid)

# ------------------------
# Helpers
# ------------------------
def borrow_to_target_DA(assets0, debt0, DA_star):
    # Solve (Debt0 + B) / (Assets0 + B) = DA_star  => B* = (DA*Assets0 - 
Debt0)/(1 - DA)
    num = DA_star * assets0 - debt0
    den = (1.0 - DA_star)
    B_star = np.where(den > 0.0, num / den, 0.0)
    return np.maximum(0.0, B_star)

# ------------------------
# Policy simulator
# ------------------------
def run_rule_2530_with_harvest(dev_trig=-0.25, DA_target=0.30, 
harvest_a=0.10):
    risky_units   = np.full(n_paths, 1.0)   # start fully invested: q0 = 1
    cash          = np.zeros(n_paths)
    debt          = np.zeros(n_paths)

    n_int_down    = np.zeros(n_paths, dtype=int)  # downside interventions
    n_int_up      = np.zeros(n_paths, dtype=int)  # upside interventions 
(repay and/or harvest)
    n_borrow      = np.zeros(n_paths, dtype=int)  # times we actually borrowed
    n_repay       = np.zeros(n_paths, dtype=int)  # times we repaid debt
    n_harvest     = np.zeros(n_paths, dtype=int)  # times we harvested to cash

    for t in range(1, n_steps + 1):
        # Mark-to-market
        risky_units *= (S[:, t] / S[:, t-1])
        risky_val = risky_units * S[:, t]
```

```
wealth    = risky_val + cash - debt

# Deviation vs geometric expected path
dev = (wealth - Wstar[t]) / Wstar[t]

# ---- Downside: restore to D/A target when dev ≤ trigger ----
do_down = dev <= dev_trig
if np.any(do_down):
    idx = np.where(do_down)[0]
    St  = S[idx, t]

    # deploy positive cash first
    deploy = np.maximum(cash[idx], 0.0)
    if np.any(deploy > 0.0):
        risky_units[idx] += deploy / St
        cash[idx]        -= deploy

    # recompute assets and debt after cash deployment
    risky_val_i = risky_units[idx] * St
    assets0     = risky_val_i + cash[idx]
    debt0       = np.maximum(debt[idx], 0.0)

    # borrow to reach target D/A
    B = borrow_to_target_DA(assets0, debt0, DA_target)

    need_borrow = B > 0.0
    if np.any(need_borrow):
        i          = idx[need_borrow]
        St_i       = S[i, t]
        debt[i]   += B[need_borrow]
        risky_units[i] += B[need_borrow] / St_i
        n_borrow[i]    += 1

        if friction != 0.0:
            cost = friction * B[need_borrow]
            use_for_cost = np.minimum(cash[i], cost)
            cash[i] -= use_for_cost
            debt[i] += (cost - use_for_cost)

    n_int_down[idx] += 1

# ---- Upside: repay debt first; then harvest α·(W - W*) into cash ---
-
# Recompute after any downside trades
risky_val = risky_units * S[:, t]
```

```
        wealth     = risky_val + cash - debt
        dev        = (wealth - Wstar[t]) / Wstar[t]

        do_up = dev > 0.0
        if np.any(do_up):
            idx = np.where(do_up)[0]
            St  = S[idx, t]

            # 1) repay any debt by selling risky
            has_debt = debt[idx] > 0.0
            if np.any(has_debt):
                i         = idx[has_debt]
                repayAmt = debt[i].copy()
                units_to_sell = repayAmt / St[has_debt]
                risky_units[i] -= units_to_sell
                debt[i] = 0.0
                n_repay[i] += 1

            # 2) harvest α * (W - W*) to cash (sell risky; proceeds to cash)
            # recompute wealth for idx after potential repayment
            risky_val_idx = risky_units[idx] * St
            wealth_idx    = risky_val_idx + cash[idx] - debt[idx]
            excess        = np.maximum(wealth_idx - Wstar[t], 0.0)
            harvest_amt   = harvest_a * excess

            do_h = harvest_amt > 0.0
            if np.any(do_h):
                i      = idx[do_h]
                units  = harvest_amt[do_h] / St[do_h]
                risky_units[i] -= units
                cash[i]        += harvest_amt[do_h]
                n_harvest[i]   += 1

                if friction != 0.0:
                    cost = friction * harvest_amt[do_h]
                    use_for_cost = np.minimum(cash[i], cost)
                    cash[i] -= use_for_cost
                    # if cost > cash, sell tiny extra next step; ignored here

            n_int_up[idx] += 1

    # Terminal metrics
    risky_terminal  = risky_units * S[:, -1]
    wealth_terminal = risky_terminal + cash - debt
    terminal_dev    = (wealth_terminal - Wstar[-1]) / Wstar[-1]
```



```
    DA_terminal = np.divide(np.maximum(debt,0.0),
                            np.maximum(risky_terminal + cash, 1e-12))

    return dict(
        mean_terminal_dev_pct   = float(np.mean(terminal_dev) * 100),
        median_terminal_dev_pct = float(np.median(terminal_dev) * 100),
        pct_within_pm1          = float(np.mean(np.abs(terminal_dev) <= 0.01)
* 100),
        pct_within_pm3          = float(np.mean(np.abs(terminal_dev) <= 0.03)
* 100),
        avg_DA_terminal         = float(np.mean(DA_terminal) * 100),
        share_with_debt_T       = float(np.mean(debt > 0) * 100),
        mean_cash_T             = float(np.mean(cash)),
        avg_downside_interv     = float(np.mean(n_int_down)),
        avg_upside_interv       = float(np.mean(n_int_up)),
        share_paths_borrowed    = float(np.mean(n_borrow > 0) * 100),
        avg_borrow_events       = float(np.mean(n_borrow)),
        share_paths_repaid      = float(np.mean(n_repay > 0) * 100),
        avg_repay_events        = float(np.mean(n_repay)),
        share_paths_harvest     = float(np.mean(n_harvest > 0) * 100),
        avg_harvest_events      = float(np.mean(n_harvest)),
        params_dev_trig         = dev_trig,
        params_DA_target        = DA_target,
        params_harvest_alpha    = harvest_a
    )

# -----------------------
# Run the rule test
# -----------------------
res = run_rule_2530_with_harvest(dev_trig=dev_trig, DA_target=DA_target,
harvest_a=harvest_a)
df = pd.DataFrame([res])
print("\nRule Test — If W ≤ 0.75·W*: set D/A = 30%; If W > W*: repay debt,
then harvest α·(W - W*) to cash\n")
print(df.to_string(index=False))
```

Acknowledgements

I am indebted to Michael Gibbs, Charles J. Cuny, Mahendra Gupta, Laurence Siegel, Jeffrey Jaffe, Chris Merker, Steve Vanourny, and Francis McMahon for mentorship, sharp criticism, tough love, and guidance. Their comments improved this work; any remaining errors are mine.

Much of the framing owes to conversations with my children, whose insistence that two seemingly contradictory perspectives can coexist proved central here: debt can be risky at a point in time and, under stated assumptions, improve outcomes in expectation over time. The Schrödinger's-cat metaphor I use in teaching captures that "two-states-at-once" intuition; it is not developed formally in this note.

I am also grateful for the intellectual lineage that made this project possible; had Milton Friedman not mentored Gary Becker, and Becker not mentored Michael Gibbs, this work would likely have taken a different form.

Caution & disclaimer

Nothing here is investment advice. Borrowing introduces real risks and constraints; outcomes depend on costs, collateral, taxes, and tail behavior. Test before use, and fit the policy to your balance sheet, not the other way around.

