THE SYSTEMATIC PORTFOLIO FRAMEWORK

A Governance Framework

for the Self-Directed Investor

By Gregory C. Eaves

April 2026

First Principles . Five-Scenario EV . Three-Prompt Sessions
Prediction Markets . Options Overlay . Layer 4 Due Diligence

Second Edition

This guide is published for educational purposes only. It does not constitute financial advice, investment advice, or a recommendation to buy or sell any security. All investment decisions involve risk, including the possible loss of principal. Past performance does not in any way influence or affect future results.

Consult a qualified financial professional before making any investment decision. Do not take money advice from strangers on the internet, especially not from this author.

Second Edition, April 2026.

Contents

Preface

How to Use This Book

Most investment books tell you what to buy. This one teaches you how to run a portfolio. There is a real difference. Knowing which exchange-traded fund (ETF) to own matters far less than knowing how to govern your decisions about it over time, when to add, when to exit, when to hold through pain, when to admit that a thesis is broken.

This guide presents a complete, rules-based governance framework for a self-directed investment portfolio. The framework has been refined through hundreds of sessions of real portfolio management across three tax wrappers. Every rule exists because its absence caused a problem, or because its presence prevented one.

The framework rests on six core commitments:

- Write everything down. Undocumented decisions do not exist.
- Separate policy from data. The Investment Policy Statement (IPS) governs what you will do. Session notes record what you did.
- Quantify convictions. Every tactical position carries a five-scenario expected value (EV) model with explicit probability weights.
- Pre-commit to exits. Define the conditions that invalidate a thesis before you open the position, not after it starts losing you money.
- Respect tax wrappers. For most US investors, there's a normal account, an IRA, and a Roth IRA, and all are taxed differently. The three accounts are one portfolio. Other jurisdictions (UK, Australia, Canada, etc.) also have different tax structures for different types of accounts, some with limits on contributions. Use this to your advantage. Where you hold an asset matters as much as what you hold.
- Use all available information. Prediction markets, energy inventory data, conflict trackers, etc. If it is relevant and readable, it belongs in your intelligence stack.

How this book is structured

Part I is new to this edition. I took nine canonical investment theories, I strip each to its first principles, and then I identify the surviving truths that the operational framework is built on. This puts clear limits on what you can actually be sure of, forcing you to think through and justify each buy or sell. Read it first.

Part II establishes the foundations: your obligation, your accounts, and asset location. Part III builds the portfolio architecture layer by layer. Part IV covers the governance constraints that prevent the most common self-directed investor mistakes. Part V covers the options overlay, both covered calls and tail hedges. Part VI covers intelligence: what to read, how to integrate it, and the two cascade protocols

(SS17 for geopolitical shifts, SS18 for macro shifts). Part VII covers session management and document architecture.

Each chapter closes with a workbook section, blank templates you fill in for your own portfolio. Work through them in order the first time. After that, return to individual chapters as reference.

A NOTE ON SPECIFICITY

This guide uses concrete examples throughout: specific EV thresholds, specific constraint limits, specific order naming conventions, and so on. The numbers are illustrative, not prescriptive.

Your obligation, your risk tolerance, and your tax situation will require you to calibrate the specific thresholds for your own portfolio.

The framework is the point. The numbers are placeholders until you replace them with yours. Your mileage may vary. Do not listen to strangers on the internet. That includes this author.

PART I

Theoretical Foundations

Before any framework can be built, the ground beneath it must be surveyed. Part I takes nine investment theories (Modern Portfolio Theory, Post-Modern Portfolio Theory, Behavioral Finance, Black–Litterman, Liability-Driven Investing, Goals-Based Investing, Factor Investing, the Efficient Market Hypothesis, and Core-Satellite), applies the Aristotelian first-principles method to each, and identifies what survives when every assumption is stripped away.

The chapters are short. The point is not encyclopedic coverage. The point is to isolate the irreducible truths that the operational framework is built on, and to show you exactly which conventional-wisdom assumptions the framework deliberately rejects.

Chapter 1

Why First Principles Matter

Investment theory is an unusually assumption-dense field. Every model rests on premises that are rarely stated plainly and even more rarely examined. A portfolio framework built on unexamined premises inherits all of their failure modes.

The Aristotelian method runs in five phases for each theory. Phase 1 surfaces every embedded assumption. Phase 2 strips the assumptions away and asks what propositions cannot be denied. Phase 3 reconstructs three distinct portfolio approaches from the surviving truths alone, as if no prior theory existed. Phase 4 maps conventional thinking against first principles, side by side. Phase 5 identifies the single sharpest insight that conventional analysis would never produce.

This is not an academic exercise. Every operational rule in Parts II through VII is traceable to a first principle from Part I. When the framework says “pre-commit to exits,” it is not invoking a slogan. It is applying Kahneman–Tversky loss aversion to the specific question of tactical position sizing. When the framework says “size the satellite to zero unless you can name the mechanism,” it is applying the Grossman–Stiglitz information paradox to a specific allocation decision.

> *The rule exists because the first principle demanded it. If you cannot state the principle, you cannot defend the rule.*

A note on humility

First-principles thinking does not produce certainty. It produces a tighter set of things one can honestly claim to know, and a clearer view of what must remain uncertain. The framework assumes uncertainty. Every rule is a hedge against a specific form of it.

Read Part I once in sequence. Then return to any individual chapter when you want to understand why a specific operational rule exists in Part IV or Part VI.

Chapter 2

Modern Portfolio Theory and Its Limits

Harry Markowitz published "Portfolio Selection" in the *Journal of Finance* in 1952 and changed how institutional investors thought about asset allocation for the next 70 years. The paper, and the larger body of work that followed, became known as Modern Portfolio Theory (MPT). It introduced the efficient frontier: the set of portfolios that produce the maximum expected return for a given level of risk. It won Markowitz the Nobel Prize in Economics in 1990.

MPT remains the dominant framework at most institutional allocators. It is also the most assumption-dense theory in finance.

The assumptions MPT makes

Markowitz assumes that investors are rational and risk-averse. That risk equals standard deviation of returns. That returns are normally distributed. That correlations are stable. That markets are frictionless. That the investment horizon is a single period. That expected returns can be estimated reliably. That all investors share the same information. None of these assumptions holds in practice.

The failure mode that matters most is the correlation assumption. MPT assumes stable correlations. In crises, correlations collapse toward 1.0 precisely when diversification is supposed to help. The 2008 financial crisis is the classic example. So is March 2020.

What survives when the assumptions are removed

A handful of propositions remain intact:

- Capital is finite. Every allocation has an opportunity cost.
- Different assets generate different cash flows with different timing and certainty.
- Holding multiple assets whose values do not move in perfect lockstep reduces combined-portfolio variance.
- Future returns are uncertain. No model eliminates this.
- Higher potential return requires higher uncertainty.
- A 50% loss requires a 100% gain to recover. Gains and losses of equal size are not equivalent.
- Time matters. The sequence of returns affects terminal wealth independently of average return.
- An investor's capacity to bear loss without changing behavior is finite and context-dependent.
- Mathematical optimization is only as valid as its inputs.

The Aristotelian move

> *Replace variance as your risk metric with maximum tolerable drawdown expressed*

> *in dollars, not percentages.*

Every other reform of MPT tinkers at the edges: better correlation estimates, better return estimates, downside variance instead of total variance. The first-principles move is different. It abandons the premise that risk is a statistical property of an asset. Risk is a relational property between an asset's behavior and an investor's specific capacity and goals.

Define the dollar amount you cannot afford to lose in any 12-month period without changing your behavior or your life plan. Build the entire portfolio around never breaching that number under any historical or plausibly simulated scenario. Asset allocation, diversification, position sizing, all fall out of that single constraint.

MPT cannot make this move because it would require acknowledging that the objective function it optimizes (mean-variance efficiency) is not the investor's actual objective (not going broke).

Chapter 3

Downside Risk and the Ruin Threshold

Brian Rom and Kathleen Ferguson published Post-Modern Portfolio Theory (PMPT) in 1991 as a direct critique of MPT's symmetric risk measure. PMPT replaces standard deviation with downside deviation, measuring risk only below a target return threshold (the Minimum Acceptable Return, or MAR). It uses the Sortino Ratio instead of the Sharpe Ratio.

PMPT is a meaningful improvement over MPT. It still inherits most of MPT's other assumptions.

What PMPT gets right

PMPT's core insight is correct: losses below an investor's required return are categorically different from gains above it. Return distributions are asymmetric. They exhibit fat tails, especially on the downside. A single volatility number cannot describe this.

What PMPT still gets wrong

It still uses a single statistical ratio (Sortino) to compress multi-dimensional risk into one number. It still assumes historical distributions predict future ones. It still optimizes toward a target return rather than around survival. And it ignores the path to a loss. A brief dip and a prolonged drawdown of the same final magnitude are treated identically, which is empirically wrong and psychologically wrong.

Graham and the margin of safety

Benjamin Graham's *The Intelligent Investor* (1949) reached the same conclusion about asymmetric loss four decades before PMPT, without the statistical machinery. Graham's margin of safety principle instructs the investor to size every position so that being substantially wrong still leaves capital intact. The construction is informal — a gap between price and estimated intrinsic value wide enough to absorb analytical error, bad luck, and regime change — but the intuition is the one PMPT later tried to formalize: capital loss is categorically different from opportunity cost, and a framework that treats them identically is not a risk framework.

Margin of safety survives the conditions that break statistical models, because it is not tuned to a historical return distribution. It is a buffer against ignorance. The ruin threshold that follows is the same buffer, stated in dollars rather than in valuation discount, and applied at the portfolio level rather than the position level.

The Aristotelian move

> *Define the investor's ruin threshold, not their target return, and make it the*

> *primary portfolio constraint.*

PMPT measures downside risk relative to a target. The first-principles move inverts the question. Instead of asking “what portfolio maximizes my Sortino Ratio relative to my target?”, ask “what is the level of loss from which I cannot recover, financially or behaviorally?”

That is the ruin threshold. Build a portfolio where the probability of reaching that threshold in any simulated scenario is effectively zero. The target return becomes secondary, a product of what remains after the ruin constraint is satisfied.

This move is the reason the operational framework (Part IV, Chapter 20) begins with a ruin analysis rather than a return target. Stage 3 of the drawdown protocol is the ruin threshold. Stages 1 and 2 are the pre-committed intermediate actions that keep you from ever reaching it.

Chapter 4

Behavioral Finance as Design Constraint

Kahneman and Tversky published "Prospect Theory: An Analysis of Decision Under Risk" in *Econometrica* in 1979. They showed that humans systematically deviate from expected utility theory: we are loss-averse, we overweight small probabilities, and we treat gains and losses asymmetrically. Kahneman won the Nobel in 2002. Tversky died in 1996, before the prize was awarded.

Richard Thaler extended this work to financial markets, documenting the disposition effect (selling winners too early, holding losers too long), mental accounting, and overconfidence. He won the Nobel in 2017.

Shefrin and Statman formalized Behavioral Portfolio Theory in 2000, modeling investors as holding layered "mental account" portfolios rather than mean-variance optimal ones.

What behavioral finance proved

Humans make decisions under uncertainty using heuristics, not optimization algorithms. Losses are experienced more acutely than equivalent gains. Investor behavior has direct consequences for investment outcomes independent of market behavior. The psychological sustainability of a strategy determines whether the strategy survives adverse conditions. The framing of a decision changes the decision made, even when the economic content is identical. Awareness of a bias does not reliably prevent its expression.

Where behavioral finance typically stops

Most behavioral finance interventions assume that if you show investors their biases, they will invest better. The evidence does not support this. Showing someone a cognitive bias chart does not reduce the bias. It makes them feel worse about the bias while continuing to exhibit it.

The Aristotelian move

> *Design the portfolio's decision architecture before designing the portfolio itself.*

Treat the investor's decision environment as the primary design problem and the portfolio as the output of that environment. Define which decisions the investor will make manually (none, if possible), which will be automated, and under what conditions human override is allowed. Then build the portfolio to function within that environment.

This move is the reason the operational framework pre-commits to exits, invalidation events, and time stops (Chapter 21). It is the reason the framework insists on written documentation for every decision

(Chapter 29). It is the reason the session structure (Chapter 28) separates intelligence from diagnosis from action: each phase protects a specific decision from a specific bias.

The framework does not ask you to be rational. It assumes you will not be, and structures the environment so that your predictable irrationality becomes less expensive.

Chapter 5

The Black–Litterman Track Record Requirement

Fischer Black and Robert Litterman developed the Black–Litterman model at Goldman Sachs in 1990 as a practical patch on MPT's sensitivity to expected return inputs. Black–Litterman starts from market-cap equilibrium and allows investors to express specific views, blending the two in a Bayesian framework.

The model produces more stable portfolios than raw MPT. It is still fragile in a specific way.

The fragility

Black–Litterman requires investors to quantify their confidence in each view. The Omega matrix (view uncertainty) is almost never calibrated rigorously in practice. Most investors assign confidence levels by feel. That feel is exactly the thing the model was supposed to discipline.

There is also the asymmetry problem. An investor who has never successfully predicted sector rotation will still assign a 70% confidence to their latest sector-rotation view, because conviction and accuracy are not the same thing, and humans systematically confuse them.

The Aristotelian move

> *Replace view confidence calibration with a track record requirement. No active deviation from the market portfolio is permitted without a documented history of accuracy in that specific type of view.*

Black–Litterman's core insight is correct: start from market equilibrium and deviate only where you have genuine information advantage. The first-principles move replaces the fabricated confidence number with a real one, your historical accuracy rate in this specific type of prediction.

If you have never tracked your sector-rotation calls, your confidence weight in a sector-rotation view is zero, regardless of how compelling the narrative feels. Building a personal prediction log becomes a prerequisite to active allocation. This might take years.

This move is the reason the operational framework (Chapter 17) requires an Adversarial Advantage Statement for every Layer 3 position before it is opened. It is the reason the framework separates Layer 1 (no views required, no track record required) from Layer 2 or Layer 3 (views required, and the view quality is audited every session against the EV model).

Chapter 6

Liability-Driven Investing for Individuals

Liability-Driven Investing (LDI) emerged from defined-benefit pension fund management in the 1990s and 2000s and became prominent after the Pension Protection Act of 2006 in the US. The core idea is that the fund's liability stream (future pension payments) is the benchmark, not a market index. The portfolio is constructed to match or hedge those liabilities.

LDI is more honest about its objective than most theories. It still introduces complexity that obscures the original insight.

What LDI proved

Every investor, individual or institutional, has future obligations that must be met from their asset pool. The relevant measure of investment success is not absolute return. It is the ability to meet specific obligations at specific times.

The mismatch between the timing of assets and the timing of obligations is the fundamental source of financial risk. Reducing the sensitivity of the asset-liability gap to interest rates, inflation, and longevity is a verifiable risk reduction.

Why LDI rarely reaches individual investors

Most individual investors have never heard of LDI. They manage their money against an index or a vague sense of "enough." This is a category error. An index is not an obligation. "Enough" is not a specification.

The Aristotelian move

> *Map every financial decision to a specific obligation. Refuse to take risks that are not compensated by a corresponding reduction in the probability of obligation failure.*

Build an explicit obligation schedule: when do you need money, how much, how certain are you of the amount? For every risk in the portfolio, ask a single question: does bearing this risk reduce the probability of meeting an obligation, or does it only raise expected return?

Risks in the second category are permissible only after all obligations in the first category are fully funded. Return generation is a secondary objective.

This is the reason Chapter 12 (The Obligation Framework) comes before everything else. The obligation is a dollar amount by a date. Required compound annual growth rate (CAGR) falls out of it. Risk tolerance falls out of it. The passive core sizing falls out of it. A portfolio without a named obligation is a portfolio without a governor.

Chapter 7

Goals-Based Investing and the Minimum Acceptable Outcome

Goals-Based Investing (GBI) was popularized in the 2000s and 2010s by practitioners including Ashvin Chhabra and by many wealth-management firms. It organizes an investor's portfolio around specific, named life goals (retirement, education, legacy, lifestyle) rather than around a single aggregate portfolio optimized for risk-adjusted return.

GBI is intuitively appealing and frequently oversold.

What GBI gets right

Investors genuinely hold multiple distinct financial objectives with different timelines, magnitudes, and emotional weights. A single aggregate portfolio metric (total return, Sharpe Ratio) cannot capture whether an investor is on track for all of their distinct objectives.

Goal clarity reduces decision error. Some goals are non-negotiable (survival, basic income, healthcare). Others are aspirational (wealth growth, legacy). These categories require different risk frameworks.

What most GBI implementations get wrong

They optimize toward ideal goal outcomes. They use Monte Carlo simulations with hidden input assumptions. They treat goals as static and stable over time. They assume goals can be funded independently, as if capital were not finite. They fail to provide a rigorous basis for goal prioritization when goals compete for the same dollar.

The Aristotelian move

> *Define the minimum acceptable outcome for every goal before defining the ideal outcome. Fund every minimum to near-certainty before allocating any capital to ideal outcomes.*

GBI correctly identifies that investors have multiple goals. It then optimizes toward ideal goal outcomes, which requires assuming more risk than necessary.

Distinguish the minimum acceptable outcome (which must be funded with near-certainty) from the ideal outcome (which is aspirational). Fund the floor of every goal before funding the ceiling of any goal.

This move requires asking an uncomfortable question that GBI frameworks almost never ask: what is the worst version of this goal I could still live with? The answer determines the risk-free portion of the portfolio. Everything above the floor is aspirational and can bear risk.

The floor-first principle runs through the operational framework. The passive core floor (55% of taxable net asset value, or NAV) is a floor, not a target. The Roth IRA (Individual Retirement Account) rule (Layer 1 only, never accessed before the obligation date) is a floor. The drawdown protocol's Stage 1 and Stage 2 actions are floor-defense, not return optimization.

Chapter 8

Factor Investing After Costs

Factor investing is rooted in the Fama–French Three-Factor Model (1992) and extended through the Five-Factor Model (2015) and momentum research (Jegadeesh and Titman, 1993). It identifies systematic, persistent return premiums associated with specific characteristics: market beta, size, value, profitability, investment conservatism, and momentum.

Factor investing is the best-documented source of structural return premium available to self-directed investors. It is also the most oversold.

The problem with how factors are sold

Factor products are marketed on gross historical premiums: "value has returned X%," "small cap has returned Y%." These numbers are before taxes, before fund costs, before trading friction, before the inevitable period of underperformance that tests behavioral durability. The decade from 2010 to 2020 was exactly that period for the value factor.

Academic factor research also suffers from the "factor zoo" problem. Hundreds of factors have been published. Many are data-mining artifacts rather than genuine return drivers. Distinguishing durable factors from spurious ones is an unsolved problem in the academic literature.

What survives

Securities with different structural characteristics have historically generated different long-run returns. Some return differences reflect durable structural features (a company being cheap relative to fundamentals). Return premiums that are well-documented and widely known face arbitrage pressure that may reduce or eliminate them over time. Implementation costs, taxes, and portfolio turnover have direct, computable effects on realized premium.

The Aristotelian move

> *Calculate the minimum factor premium required to break even after all implementation costs in your specific tax situation before investing in any factor product.*

Compute the break-even premium for your specific situation: marginal tax rate, account type, expense ratio, estimated turnover and its tax impact, realistic holding period. If the expected net premium does not survive this calculation with a meaningful margin of safety, the factor exposure is not worth taking.

In practice, this eliminates most actively managed factor products and many factor ETFs. It is also the calculation that factor-product manufacturers do not want you to do.

The operational framework's Layer 1 choices reflect this. Factor-tilted ETFs from Avantis and Dimensional (profitable value, size, low-investment) are preferred because they pair documented premium with low cost and low turnover. Momentum ETFs are generally avoided at the retail level because the turnover kills the premium after taxes.

Chapter 9

The Efficient Market Hypothesis as Diagnostic

Eugene Fama formalized the Efficient Market Hypothesis (EMH) in 1970. Markets come in three forms: weak (prices reflect all historical prices), semi-strong (prices reflect all public information), and strong (prices reflect all information, including private). The practical implication, in Fama's telling, is that consistent excess returns are not achievable without taking on additional risk. Fama won the Nobel in 2013.

EMH is one of the most debated propositions in finance. The debate is frequently useless because both sides are responding to caricatures.

The joint-hypothesis problem

EMH is difficult to test cleanly. Any test requires a joint hypothesis about the "correct" price model. If the test fails, either markets are inefficient or the price model is wrong, and you usually cannot tell which. This makes EMH harder to falsify than it looks.

What is actually true

Asset prices are set by the marginal transactions of buyers and sellers, each acting on their own information and objectives. Widely known and easily actionable information is rapidly incorporated into prices. Information that is costly to obtain, difficult to analyze, or restricted in use is incorporated more slowly and incompletely.

The cost of gathering, analyzing, and acting on information determines how efficiently that information is priced. Consistently outperforming the market requires either an information advantage, an analytical advantage, a behavioral advantage, or the acceptance of risks that the market has priced incorrectly.

Most investors, after costs, underperform the market over long periods. The degree of market efficiency varies across asset classes, conditions, and time periods.

Malkiel and the case for indexing

Burton Malkiel's *A Random Walk Down Wall Street* (1973) translated Fama's theoretical claim into a retail-facing argument that has held up across five decades of revised editions. Malkiel's move is empirical rather than theoretical. He surveys the evidence on technical analysis, on active fund management, and on most forms of stock-picking, and shows that after costs, taxes, and survivorship bias, the large majority of approaches fail to beat a broad index over long periods. The prescription that follows — a diversified, low-cost, index-heavy core held with discipline — remains the correct default for any investor who cannot answer the questions the diagnostic move below will demand.

Malkiel's contribution to this framework is the humility check. Most investors who believe they can beat the market cannot. The base rate is against them and they have not inspected it. If you cannot articulate, in specific and testable terms, why you are not one of the underperformers, you almost certainly are.

Bogle and the arithmetic of costs

John C. Bogle, founder of Vanguard and creator of the first retail index mutual fund in 1976, made the indexing case from a different angle than Malkiel in *The Little Book of Common Sense Investing* (2007). Bogle's argument is arithmetical rather than empirical. In aggregate, investors own the market, so their gross return must equal the market's return before costs. After costs, the collective result must fall short of the market by the size of those costs. This is a claim about identity, not about market efficiency, which is what makes it stronger than the random-walk case: you need not believe in efficient markets to believe in Bogle's arithmetic.

Bogle's contribution to this framework is the cost floor. Every active position, regardless of analytical quality, must clear a cost hurdle before it can justify displacing a passive alternative. Chapter 8's break-even premium calculation is this arithmetic applied at the position level, extended to include taxes, turnover, and realistic holding periods.

Graham and the investment test

Benjamin Graham's *The Intelligent Investor* (1949), written two decades before Fama formalized EMH, anticipated the diagnostic use of the efficiency question. Graham's central distinction — between investment (an operation that, on thorough analysis, promises safety of principal and an adequate return) and speculation (everything else) — maps directly onto the test this chapter requires before any active position is opened.

Graham does not deny that market prices can diverge from intrinsic value. He insists that acting on the divergence requires a specific, defensible analytical basis, and he treats the absence of such a basis as disqualifying. Read alongside Malkiel, the two form a complete retail-facing apparatus: Malkiel provides the humility check (the default is index), Graham provides the positive discipline (deviate only where the investment test is passed). The Aristotelian move below is their combined logic rewritten for a post-EMH world.

The Aristotelian move

> *Before any active investment decision, write down your specific, verifiable advantage over the marginal market participant for that investment. If you cannot, index.*

EMH is most often used as either a blanket argument for passive investing (ignoring that efficiency varies) or dismissed as obviously wrong (pointing at Warren Buffett). Both uses miss the point.

Treat EMH as a diagnostic tool. For every active position, force yourself to state the specific mechanism by which you expect to outperform the aggregate judgment of everyone else pricing this asset. Not a narrative. A mechanism. Who is on the other side of this trade? Why are they wrong? What information or analytical edge do you have that they lack? Why has the market not already priced this in?

If you cannot answer these questions with specificity, the position is speculation, not investment. This is the test every Layer 3 position must pass before it is opened.

Chapter 10

Core-Satellite as Null Hypothesis

Core-Satellite emerged in the 1990s and 2000s as a practical compromise between passive and active investing. A large passive core (typically 60% to 80% of the portfolio) sits alongside a smaller satellite portion (active, thematic, or alternative investments) intended to generate alpha or provide exposures the core lacks.

Core-Satellite is the default portfolio architecture at most wealth-management firms. It is also frequently the institutionalization of investment "fear of missing out" (FOMO).

Where the conventional 60/40 or 80/20 split comes from

Industry convention. That is the honest answer. It is not derived from any first-principles analysis of an investor's specific situation. It is the number that lets wealth managers justify their fees while not deviating so far from the index that they lose their clients to an ETF.

What is wrong with the satellite

The satellite exists, structurally, to satisfy the need to "do something." Most satellite positions are not based on genuine investment theses. They are based on narrative, recency bias, or professional anxiety about underperforming a benchmark.

Satellites also rarely provide the diversification they promise. If the satellite holds stocks, sector ETFs, or factor funds, its correlation to the core is often high, especially in crises, which is exactly when diversification is supposed to matter.

The Aristotelian move

> *Set your satellite allocation to zero unless you can name, in writing, the specific mechanism by which each satellite position will outperform. Track the prediction formally.*

Treat the satellite as a null hypothesis. The default allocation is zero. Every satellite position must actively earn its place by meeting a minimum standard of evidence.

That standard is a written thesis stating the return mechanism, the expected return above the core, the time horizon, and the exit criterion. The position is reviewed against this thesis, not against the market.

This move is the reason the operational framework sizes Layer 2 or Layer 3 by explicit thesis quality, not by a standing allocation. It is the reason the framework requires a time stop (Chapter 21). It is the reason the kill-switch (Chapter 28) forces a formal performance review of the tactical book against the

passive core every quarter. If the satellite cannot beat the core on a risk-adjusted basis over a meaningful period, the satellite disappears.

Chapter 11

The Seven Principles That Survived

Across all nine theories, a small set of recurring irreducible truths has emerged. These are the principles that cannot be denied once every assumption has been stripped away. They are the scaffolding of everything that follows in Parts II through VII.

PRINCIPLE 1: CAPITAL IS FINITE

Every allocation has an opportunity cost. Every dollar placed against one goal cannot also be placed against another. Portfolio design is always a resource-allocation problem across competing uses.

PRINCIPLE 2: SUCCESS IS MEETING OBLIGATIONS, NOT BEATING BENCHMARKS

The relevant measure of investment success is the ability to meet specific obligations and goals at specific times. Relative performance against a benchmark is a distraction, except where it feeds directly into obligation funding.

PRINCIPLE 3: CATASTROPHIC LOSS MATTERS MORE THAN AVERAGE RETURN

The probability and magnitude of ruin-level loss matters more than average expected return. Returns and losses of equal size are not equivalent. A 50% loss requires a 100% gain to recover. Sequence matters.

PRINCIPLE 4: BEHAVIOR IS THE PRIMARY DETERMINANT OF OUTCOMES

Investor behavior, not portfolio theory, is the primary determinant of actual investment outcomes. The portfolio that is maintained through adversity outperforms the optimal portfolio that is abandoned.

PRINCIPLE 5: MATHEMATICAL SOPHISTICATION DOES NOT REDUCE UNCERTAINTY

Complexity in portfolio construction does not eliminate the fundamental uncertainty of future returns. Complexity introduces fragility. A simple rule-based system that is actually followed outperforms a sophisticated optimizer that is abandoned.

PRINCIPLE 6: COSTS, TAXES, AND FRICTION ARE CERTAIN. ALPHA IS NOT

Implementation costs, taxes, and behavioral friction are certain drags on return. Alpha is uncertain. Minimize the certain drag before chasing the uncertain premium.

PRINCIPLE 7: THE CORRECT PORTFOLIO IS THE ONE THAT IS MAINTAINED

The correct portfolio is the one that is actually held through the full cycle, not the one that is mathematically optimal in a frictionless simulation. Design the portfolio around the investor who will hold it, not the investor

who theoretically should.

How these principles map to the operational framework

Principle 1 drives the sequence rules (Chapter 19) and the passive-core dominance rule. Principle 2 drives the Obligation Framework (Chapter 12) and the Liability-Driven structure of the three-account architecture (Chapter 13). Principle 3 drives the Drawdown Protocol (Chapter 20) and the four exit conditions (Chapter 21).

Principle 4 drives the entire pre-commitment apparatus: invalidation events, time stops, the three-prompt session structure. Principle 5 drives the framework's preference for simple rules over complex optimizers. Principle 6 drives the Layer 1 choice of low-cost factor ETFs and the LTCG Governance Rule (Chapter 18). Principle 7 drives the behavioral durability that every rule in Parts III through VII serves.

> *The framework is a system for turning these seven principles into specific orders on specific trading days.*

Parts II through VII are the translation layer. They take the principles that cannot be denied and turn them into operational rules that can actually be checked at a session open. If at any point you disagree with a rule, the best place to look for the disagreement is here, in Part I. Trace the rule back to the principle. Test the principle. Then decide.

PART II

Foundations

Obligation. Accounts. Asset location. Part II establishes the three foundation layers that every subsequent decision stands on. Without a named obligation, there is nothing to manage against. Without a clear account architecture, tax-inefficiency quietly eats returns. Without asset location discipline, the right assets end up in the wrong wrappers.

Chapter 12

The Obligation Framework

Most investors manage against an abstraction. “Beat the market.” “Grow my wealth.” “Retire comfortably.” These are not goals. They are wishes.

A governance framework requires a specific, measurable, time-bounded obligation.

12.1 Defining your obligation

Your primary obligation is a single statement in the form of: achieve a total portfolio value of $X by date Y. Everything in the framework flows from that statement. It determines your required compound annual growth rate (CAGR), your drawdown tolerance, your sizing of tactical positions, and your distribution policy. Net asset value (NAV) refers to the total dollar value of your portfolio across all accounts at current prices, and it is the denominator against which the obligation is tracked.

OBLIGATION FORMAT

Achieve a total portfolio value of $[TARGET] by [DATE].
Starting NAV: $[CURRENT].
Time horizon: [N] years.
Required CAGR: [X]%.
Distribution policy: [none / income required: $Y/year].

Required CAGR is computed as (Target / Starting NAV)^(1/Years) minus 1. That number is the most important figure in your framework.

If the required CAGR is below 7%, a passive index strategy will likely suffice. If it sits between 7% and 15%, a factor-tilted passive core with modest tactical overlay is appropriate. Above 15%, you are in territory where tactical positions and an options overlay become essential, and so does rigorous exit discipline, because the cost of holding a broken thesis grows faster than the potential upside.

12.2 Ruin analysis

The obligation framework requires a ruin analysis. A ruin event is a drawdown large enough that the required CAGR to reach your target becomes effectively unachievable given the remaining time.

The method is mechanical:

- Compute what a 20% drawdown from your high-water mark does to your required CAGR.
- Compute what a 25% drawdown does.
- Identify the drawdown level at which the required CAGR crosses your personal threshold of achievability (typically 30% to 35% CAGR). That is your ruin threshold.

- That threshold becomes your Stage 3 drawdown level. Build the entire drawdown protocol around it.

WORKED EXAMPLE

Starting NAV $500,000. Target $5,000,000 in 10 years. Required CAGR: 25.9%.

A 25% drawdown produces $375,000. Required CAGR from $375,000 to $5,000,000 in 10 years: 29.5%. Borderline achievable.

A 30% drawdown produces $350,000. Required CAGR: 30.4%. Effectively ruin.

So the Stage 3 Ruin threshold is set at 28% to 30% from your HWM. Everything above that level is protocol territory. Below it, the obligation formally needs to be reassessed.

12.3 The high-water mark

Your high-water mark (HWM) is the highest total portfolio NAV ever recorded at settled prices. Drawdown stages are measured from the HWM, not from a fixed starting value.

The HWM updates upward when the portfolio sets a new high, and it never moves downward. The stages move with it.

One rule matters more than any other: HWM updates only from post-session NAV values confirmed by a brokerage CSV (comma-separated values) export, not from a blip in the middle of the trading session. Forward-estimated or provisional NAVs do not propagate. This prevents drawdown stages from anchoring to aspirational figures.

Workbook: your obligation statement

Field	Your Value
Target Portfolio Value	
Target Date	
Current NAV (all accounts)	
Time Horizon (years)	
Required CAGR	
Distribution Policy	
Ruin Analysis: 25% drawdown NAV	
Ruin Analysis: required CAGR from there	
Ruin Threshold (Stage 3)	
Stage 1 (-15% from HWM)	
Stage 2 (-20% from HWM)	

Field	Your Value
Current High-Water Mark	

Chapter 13

The Three-Account Architecture

Most US-tax-paying self-directed investors treat their brokerage account, their traditional Individual Retirement Account (IRA), and their Roth IRA as three separate portfolios. This is a mistake. They are one portfolio expressed across three tax wrappers. This framework treats them as such.

For non-US investors, there are analogous account types in your home jurisdiction. There are RRSPs and TFSAs in Canada, the SIPP and Stocks and Shares ISA in the UK, and superannuation in Australia. Same thing, but different. Details vary across jurisdictions.

Fundamentals stay the same. Consult your local tax professional.

13.1 The three wrappers

Account	Tax Treatment	Withdrawals	Required Minimum Distributions (RMDs)	Best For
Taxable Brokerage	Dividends and realized gains taxable annually.	Any time. Capital gains rates apply.	None	Tax-efficient buy-and-hold: Layer 1 passive core, tactical ETFs.
Traditional IRA	Pre-tax contributions. Growth deferred.	Taxed as ordinary income at withdrawal.	Age 73+	Tax-inefficient assets: gold ETFs (28% collectibles rate sheltered), high-dividend stocks.
Roth IRA	After-tax contributions. Growth tax-free.	Tax-free after age 59.5.	None	Highest-conviction, highest-expected-return passive core.

13.2 The Roth IRA or TFSA is the most valuable account per dollar

This point deserves emphasis. Per dollar of assets, for US investors the Roth IRA (TFSA for Canadians; the after-tax contributions account for others) beats every other account structure because every dollar of growth is permanently tax-free. At a required CAGR above 20%, the tax-free compounding benefit over a 10-year horizon is substantial.

THE COMPOUNDING MATH

$7,000 contributed to a Roth IRA or TFSA, compounding at 20% CAGR for 10 years, becomes $43,358, all tax-free.

The same $7,000 in a Traditional IRA or RRSP becomes the same $43,358 before tax. At a 22% marginal rate, that is $33,819 after tax. $9,539 less.

The annual Roth or TFSA contribution deadline (April 15 for prior-year contributions in the US) is one of the most consequential deadlines in your calendar. Missing it is an unrecoverable permanent loss.

13.3 Constraint denominator policy

The floor and cap constraints in this framework (passive core floor, cluster cap, single-position cap) are calculated against taxable account NAV only. IRA account assets are governed by separate IRA-specific constraints.

This architecture preserves the precision of the taxable constraint calculations and allows the IRA accounts to be governed by simpler rules appropriate to their different roles.

Chapter 14

Asset Location

Asset location is the discipline of placing each holding in the tax wrapper where it generates the most after-tax value. It is one of the highest-value decisions in portfolio management and one of the most consistently neglected by self-directed investors.

14.1 The taxable account

The taxable account should hold tax-efficient assets: those that generate minimal annual tax drag and are best held for long periods at long-term capital gains (LTCG) rates. Factor-tilted passive ETFs are ideal. Low turnover, qualified dividends, LTCG treatment on gains, and the step-up in cost basis at death makes them exceptionally tax-efficient for multi-decade holds.

TAXABLE ACCOUNT: WHAT BELONGS, WHAT DOES NOT

ELIGIBLE: Low-turnover factor ETFs (passive core), sector tactical ETFs (governed by EV model), individual equities (governed by LTCG rule), options overlay (covered calls, protective puts).

PROHIBITED: High-dividend individual income stocks (dividends taxable annually), gold ETFs (collectibles taxed at 28%), any position held solely for yield without a thesis.

14.2 The Traditional IRA

The Traditional IRA (or RRSP) should hold assets that would otherwise generate significant annual tax drag in a taxable account. Gold ETFs are the canonical example, at least for the US tax system. Gold is taxed as a collectible at 28% in taxable accounts, but deferred entirely in an IRA. High-dividend income stocks are another: their dividends compound tax-deferred rather than being taxed at ordinary income rates each year.

The Traditional IRA is also the staging account for the Roth conversion ladder, the systematic conversion of Traditional IRA assets to Roth in low-income years.

The Roth conversion ladder

In January of each calendar year, assess your projected US federal taxable income for that year. Identify the top of your current marginal tax bracket. Convert Traditional IRA assets to Roth up to that amount. The converted amount is taxed as ordinary income in the conversion year, and thereafter grows and is withdrawn tax-free.

For investors with foreign-earned income excluded under the Foreign Earned Income Exclusion (FEIE), or those in temporary low-income years, the effective conversion rate can be very low, sometimes 0% if

the standard deduction covers the converted amount. This is a permanent, compounding benefit. Document the conversion amount, the bracket utilized, and the tax impact in your session log each year.

14.3 The Roth IRA

The Roth IRA (or TFSA) holds exclusively your highest-conviction, highest-expected-return passive core assets. In practice, this means a factor-tilted US equity ETF as the primary holding, with international equity as secondary. No tactical positions. No income positions. No options. No alternatives.

ROTH IRA RULE

Hold only Layer 1 passive core assets.

Deploy the annual contribution immediately on arrival: maximum whole shares of the primary equity ETF first, then international ETF, then residual to money market.

Tag every contribution with the tax year it applies to.

Never access before the obligation date.

14.4 Traditional IRA target end-state

Over time, the Traditional IRA should migrate toward a simplified two-sleeve structure:

Sleeve	Assets	Target Allocation	Notes
Gold / Alternatives	Gold ETF (e.g., IAU)	60%–70% of IRA NAV	Collectibles tax (28%) sheltered. Primary holding.
Passive Core Extension	Factor ETF (e.g., AVUS)	20%–35% of IRA NAV	Migrated from any S&P 500 overlap holdings.
Legacy Income (exit path)	Individual dividend stocks	≤20%, declining	No new additions. Exit at Q2/Q3 reviews.
Cash / Settlement	Money market	≤5%	Settlement only.

PART III

Portfolio Architecture

Layers. Passive core. Tactical EV. Individual equities. Part III builds the portfolio structure on top of the foundations from Part II. Every position belongs to exactly one layer. Layer assignment determines which rules apply.

Chapter 15

The Layer System

The layer system is the structural backbone of the framework. Every position in every account is assigned to exactly one layer. Layer assignment determines what governance rules apply: whether an EV model is required, what exit conditions exist, whether covered calls are permitted, and how the position interacts with the constraint dashboard.

15.1 The five layers

Layer	Role	EV Model	Exit Discipline	Options
Layer 1	Factor-tilted ETFs. Primary long-run return driver. Permanent.	None, ever.	None. Permanent holding.	Covered calls on primary ETF when floors met.
Layer 2	Defensive and income sector ETFs (utilities, staples, financials, etc.).	Required.	Hard Exit if EV deviation below -15%. No invalidation events. No time stop.	Permitted.
Layer 3	High-conviction thesis bets (energy, defense, healthcare, etc.).	Required.	Hard Exit. Thesis-invalidation events. Time stop at quarterly review.	Not while overweight.
Layer 4	Individual equities.	None (context only).	LTCG Rule. Annual due diligence protocol.	Position-specific.
Layer 5	Options overlay. Covered calls and long puts.	None.	Per options overlay policy.	N/A.

15.2 The Layer 1 permanence rule

Layer 1 is the only layer with an absolute permanence rule. No EV model. No break conditions. No thesis-invalidation events. No time stop. Layer 1 is never reduced to fund tactical opportunities. The rule has no exceptions.

The reasoning is direct. Factor-tilted passive equity has the best documented long-run return profile of any broadly accessible asset class. The primary risk of self-directed investing is behavioral: selling the permanent core during drawdowns to fund tactical bets. The permanence rule exists to make this impossible by policy, not by willpower.

THE PERMANENCE RULE

Layer 1 is never sold to fund any other layer.

It is never reduced in response to an EV model output, a cluster cap breach, a drawdown event, or any other governance trigger.

The passive core grows monotonically over time.

This is the most important single rule in the framework.

15.3 Layer assignment in practice

Assign every current and planned position to a layer before building the constraint dashboard. Unassigned positions are a governance gap. If you cannot assign a position to a layer, you should not own it under this framework.

Chapter 16

Building the Passive Core

The passive core is the engine of the portfolio. Everything else, tactical positions, options overlay, IRA sleeves, exists to supplement or protect it. The core should be factor-tilted, internationally diversified, and held across all three accounts.

16.1 Factor tilting

Decades of academic research support systematic tilts toward value, profitability, and size factors as sources of excess return over market-cap-weighted indices. The most practical implementations for self-directed investors are the Avantis family of ETFs (AVUS for US, AVDE for developed international, AVEM for emerging markets) and the Dimensional Fund Advisors (DFA) equivalents (both, as of April 2026). These tilt systematically toward profitable value companies with low asset growth, without requiring active security selection.

16.2 The core floor and dominance rules

PRIMARY ETF FLOOR

The primary passive core ETF must be at least 30% of taxable NAV before any tactical ETF purchase is permitted. This rule has no exceptions.

PASSIVE CORE DOMINANCE

The combined passive core (all Layer 1 ETFs) must be at least 55% of taxable NAV AND must exceed the combined tactical book in dollar value at all times. Violation of either condition activates the sequence rule: no new tactical buys until compliance is restored.

These thresholds are calibrated for a portfolio with a 20%-plus required CAGR. If your required CAGR is lower, you may prefer a higher passive core floor (60% to 70%). If it is higher, you may accept a lower floor temporarily, but the dominance rule (passive greater than tactical in dollar value) should always hold.

16.3 International allocation

A minimum 10% international allocation provides geographic diversification and currency exposure. The recommended split for most self-directed investors is 7% to 10% in a broad international developed-market ETF and 3% to 5% in a dedicated factor-tilted international ETF when available and affordable.

The case for international allocation is weaker than it was 30 years ago. Global equity markets co-move much more tightly than they once did. When US large caps sell off, European and Asian developed

markets usually sell off too, often with similar magnitude. Correlations between major developed-market indices have trended higher over the past two decades, driven by cross-listed multinationals, integrated supply chains, common global risk factors (rates, commodity prices, US dollar strength), and the simple fact that the largest companies on every exchange are exposed to the same end markets.

Some international exposure still earns its place, though, for three reasons. Currency is the first. Holding non-US equities gives you claim on non-USD cash flows. In a sustained US dollar-weakness regime (persistent US fiscal deterioration, for example, or a shift away from the US dollar as the dominant reserve currency, or a narrowing in the US growth premium), that exposure becomes a hedge rather than a drag. Valuation dispersion is the second. US equities have traded at substantial price-to-earnings premiums to most developed markets for much of the past decade (mid-2010s to mid-2020s). That gap may compress rather than widen further, in which case international outperforms mechanically from multiple convergence alone. And the tails are not fully correlated. A US-specific political, regulatory, or currency shock (a debt-ceiling crisis, a sudden tax regime change, a dollar liquidity event) may affect a domestic-only portfolio meaningfully more than a globally diversified one.

Size the allocation with those reasons in mind, not with a goal of replicating global market-cap weights. The 10% figure is a floor for the insurance value. Above 30% to 40%, the diversification benefit flattens, while the tax inefficiency of international dividend withholding starts to show up in taxable accounts.

16.4 DRIP policy

Dividend reinvestment (DRIP) should be enabled on all passive core positions in all three accounts. This is particularly important in IRA (or RRSP) accounts where small dividend amounts cannot be practically reinvested manually. Enable DRIP at account opening and confirm enrollment for each position.

Chapter 17

Tactical Positions and the Five-Scenario EV Model

Tactical positions exist to provide exposure to specific thesis-driven return opportunities that the passive core cannot capture: geopolitical disruption (the Strait of Hormuz closing in early 2026), sector repricing (the 2014 oil price collapse), regulatory change (Dodd-Frank, 2010; 2015 implementation), etc. Every tactical position requires a five-scenario Expected Value (EV) model, explicit exit conditions, and (for Layer 3) a thesis-invalidation statement and a time stop. Explain to yourself why you think you should do this, and when you will get out as things go bad.

17.1 The five-scenario EV model

Expected Value is the probability-weighted average of all possible outcomes. For a tactical position, you build five scenarios (Hard Bear, Soft Bear, Base, Soft Bull, Hard Bull), each with a price target and a probability weight. The weights sum to 100%.

THE EV FORMULA

EV = (w_HB x P_HB) + (w_SB x P_SB) + (w_Base x P_Base) + (w_SBull x P_SBull) + (w_HBull x P_HBull)

Deviation % = (EV - Live Price) / Live Price x 100

Positive deviation: the position is cheap vs. its expected outcome.

Negative deviation: the position is expensive vs. its expected outcome.

Building scenario targets

For each scenario, define a price target based on what the underlying asset would be worth under that scenario's conditions. Hard Bear might be a recession plus a thesis failure. Soft Bear might be only thesis delay. Base is your central case. Soft Bull is the thesis playing out as expected. Hard Bull is better than expected.

The weights are your probability assessments. Be honest. If you cannot articulate why you have assigned a specific weight, the weight is not credible and the EV model is not functioning as a governance tool.

17.2 Two parallel EV frameworks

Mature portfolios run two parallel EV frameworks on the same tactical positions: a geopolitical cascade (labeled SS17 throughout this framework) and a macro cascade (labeled SS18). Each uses its own weight vector across the five scenarios. Each produces its own EV and its own deviation percentage.

The two frameworks answer different questions. SS17 asks: how likely is the specific geopolitical thesis to resolve favorably? SS18 asks: what macroeconomic regime are we in, and how does that affect this

ticker? A position can be cheap on SS17 (geopolitical thesis intact) and expensive on SS18 (macro regime shifting against it), or vice versa. Holding both views makes the reasoning explicit.

Chapters 26 and 27 go deeper on the two frameworks and their cascade protocols.

17.3 EV break conditions

Condition	Trigger	Action Required
C1: Two-Session Rule	EV deviation below -10% for two consecutive sessions.	S1 clock starts. Heightened monitoring. Review scenario weights.
C2: S0 Hard Exit	EV deviation below -15% in any single session.	Immediate mandatory exit. No session delay. Non-negotiable.
C3: Thesis Invalidation (L3)	Pre-defined thesis-invalidation event occurs.	Mandatory exit review. Exit unless a new thesis is documented.
C4: Time Stop (L3)	EV negative at binding quarterly review date.	Mandatory exit at that review. No extensions.

17.4 The adversarial advantage statement (Layer 3)

Before opening any Layer 3 position, write a one-paragraph adversarial advantage statement. It answers a single question: why do you have an edge in this position that the market does not already price?

If you cannot answer this question convincingly, the position is speculation, not governed tactical investing. This is the Chapter 9 (EMH diagnostic) move applied at the position level.

17.5 Thesis-invalidation events

For every Layer 3 position, define before opening the position the specific observable events that would prove the thesis wrong. They must be concrete and verifiable. Examples:

- Energy/Geopolitical: strait traffic returns to above 75% of pre-disruption average for 5+ consecutive days.
- Defense: Congressional budget cut of more than 15% to baseline defense spending, or cancellation of a core program.
- Healthcare: Medicare negotiation expanded to more than 50 additional drugs, or a Centers for Medicare and Medicaid Services (CMS) ruling with more than $200B/10yr revenue impact.

Chapter 18

Individual Equities, the LTCG Rule, and Annual Due Diligence

Layer 4 individual equities are governed by a different set of rules than tactical ETFs. They carry no five-scenario EV model. Instead, Layer 4 is governed by the LTCG Rule, position-specific judgment, and (as of this edition) an Annual Due Diligence protocol triggered by each Form 10-K filing (the annual report that US public companies file with the Securities and Exchange Commission, or SEC).

18.1 The LTCG governance rule

LTCG RULE

Before selling any Layer 4 position, review the tax lot composition via your brokerage's unrealized gain/loss export.

Prioritize selling lots held more than 12 months (long-term capital gains, preferential capital gains rates) over lots held 12 months or less (short-term capital gains or STCG, ordinary income rates).

Where STCG treatment cannot be avoided, the tax cost must be explicitly acknowledged in the session log as a deliberate decision.

Exception: if unrealized losses become large enough that holding represents a greater expected harm than the STCG exit cost, sell regardless of holding period. That determination is a matter of session judgment, documented with explicit reasoning.

The LTCG rule applies only in the taxable account. In IRA accounts there are no capital gains taxes. All withdrawals are taxed as ordinary income (Traditional IRA or RRSP) or tax-free (Roth IRA or TFSA) regardless of holding period.

18.2 The Annual Due Diligence Protocol

The first edition of this framework governed Layer 4 positions through the LTCG Rule and session-level judgment alone. Experience revealed the asymmetry. Layer 3 positions carried explicit invalidation triggers. Layer 4 did not. A Layer 4 holding could drift for quarters on narrative strength without any formal structural review.

This edition closes that gap. Every Layer 4 holding (i.e., a listed US equity position) undergoes a structured, four-phase 10-K review within 30 calendar days of the counterparty's annual 10-K filing with the SEC (publicly available). This review is referred to throughout as due diligence (DD). The output feeds the SS18 Macro EV price-target recalibration for that ticker and produces an explicit invalidation-trigger list analogous to the Layer 3 triggers.

The four phases

Phase	Duration	Objective	Required Output
Phase 1: Business Model Clarity	~15 min	Revenue composition, customer concentration, unit economics. What has materially changed vs. prior year.	Structured summary with page citations. Delta vs. prior-year review flagged explicitly.
Phase 2: Competitive Position	~15 min	Stated competitive advantages, market position, research and development (R&D) and capex as % of revenue (3yr trend), barriers to entry.	Quantitative moat scorecard. Any material deterioration in R&D/capex or new named threat flagged.
Phase 3: Financial Quality	~15 min	Profitability, return on invested capital (ROIC), cash generation including free cash flow (FCF), capital structure, capital allocation. 3-year trends.	Scorecard. Year-over-year (YoY) degradation above 20% in ROIC, FCF conversion, or interest coverage flagged.
Phase 4: Risks and Valuation Context	~15 min	Top-5 material risks, industry cyclicality, strategic priorities, growth drivers, historical volatility.	Risk-to-valuation bridge. Deal-breaker risks named. Feeds SS18 price-target refresh.

Required outputs

Every annual DD review produces three artifacts, stored with the session documents:

- Refreshed SS18 five-scenario price targets for the ticker, with explicit rationale tying each scenario to specific bottom-up findings.
- A one-page thesis update summarizing Business Model / Competitive Position / Financial Quality / Risks, with 3-year trend deltas.
- An explicit invalidation-trigger list for the ticker, formatted identically to the Layer 3 triggers in Chapter 21. Minimum one C3-type fundamental-deterioration trigger plus one structural watch item.

Harvest gating

No Layer 4 harvest may be executed on stale due-diligence data. A price target that lacks a current review (within 12 months) is flagged as STALE and cannot be referenced for harvest decisions until refreshed.

If a Layer 4 position enters harvest candidacy while its DD is stale, the review window compresses to 14 days and harvest is deferred until the review completes.

ANNUAL DD DOES NOT APPLY TO

Layer 1 (permanent passive core, no thesis).

Layer 2 or Layer 3 (ETF structures, no 10-K to review, governed by SS17/SS18 and invalidation triggers).

Layer 5 options overlays (governed by Chapter 22).

The DD protocol is a Layer 4 rule only.

PART IV

Governance and Constraints

Constraints are the structural safeguards of the framework. They define the boundaries within which discretion operates. Part IV covers the constraint dashboard, the drawdown protocol, and the four exit conditions that are checked at every session open without exception.

Chapter 19

Hard Constraints and Sequence Rules

Constraints are the structural safeguards of the framework. Unlike guidelines, constraints are hard. Violating them is not a judgment call. It is a governance failure.

Constraints are checked at every session open before any analysis begins.

19.1 The constraint dashboard

The abbreviation CC in the dashboard stands for “covered call.” Covered calls are discussed in detail in Chapter 22.

Constraint	Formula	Threshold	Breach Action
Primary ETF Floor	Primary ETF / Taxable NAV	≥ 30%	No new tactical buys until compliant.
CC Eligibility Floor	Primary ETF / Taxable NAV	≥ 35%	No covered calls on primary ETF below this level.
International Floor	(Broad Intl + Factor Intl) / Taxable NAV	≥ 10%	Priority for next deployment.
Passive Core Floor	All L1 ETFs / Taxable NAV	≥ 55%	Sequence rule activated. No new tactical buys.
Passive > Tactical	L1 value > L2+L3+L4 value	Must hold	Sequence rule activated.
Cluster Cap	Sum of themed tactical / Taxable NAV	≤ 20%	No new buys in cluster. Monitor.
Single Position Cap	Any single tactical / Taxable NAV	≤ 15%	No additions. Monitor.
Cash Floor	Settled cash balance	≥ $5,000	Prioritize cash restoration.

19.2 The cluster cap

The cluster cap prevents concentration in thematically related tactical positions. If you hold three positions that all benefit from the same underlying thesis (energy disruption, defense spending), the combined value stays below 20% of taxable NAV.

A cluster cap breach does not require immediate selling. It prohibits new additions and triggers heightened monitoring with a binding compliance decision at the next quarterly review.

The intact-thesis rule matters here. If every position in the cluster still has a valid thesis, the breach is flagged and monitored. No forced selling. The cluster cap is about preventing excess concentration, not about forcing exits from winning positions.

19.3 Escalation tiers

Tier	Duration	Protocol
Tier 1: Advisory	1–2 sessions	Flag and monitor. Document in session note. Remediation path identified.
Tier 2: Elevated	3–4 sessions	Formal assessment required. Remediation timeline set. No new positions in affected category.
Tier 3: Emergency	5+ sessions	Formal determination: Hold or Exit. Funded remediation path required. Escalate if remediation does not execute on schedule.

Chapter 20

Drawdown Protocol and the High-Water Mark

The drawdown protocol is the framework's response to sustained portfolio losses. It defines, in advance, exactly what you will do at each stage of drawdown, before the emotional pressure to act (or to freeze) makes disciplined decision-making difficult.

20.1 The four stages

Stage	Trigger	Actions Required
Stage 0 Normal	Above Stage 1	Normal operations. All layers active. Full tactical book permitted subject to sequence rules.
Stage 1	-15% from HWM	Suspend all new tactical buys. Maximize options overlay income. No Layer 3 additions. Begin passive core consolidation.
Stage 2	-20% from HWM	Exit all Layer 3 positions. Reduce to Layer 1 + Layer 2 only. Maximize covered call income on remaining positions.
Stage 3 Ruin	-25% from HWM (or your calculated ruin threshold)	Full liquidation of tactical book. Portfolio collapses to Layer 1 passive core only. No recovery trading. IPS obligation formally reassessed.

20.2 Why pre-commitment matters

The value of the drawdown protocol is not the specific thresholds. The value is that the decision is made in advance, in a calm state, with full access to rational analysis. When Stage 2 is triggered by a real drawdown, you will not be in a calm state. The protocol makes the decision for you.

The Roth IRA or TFSA is never accessed under any drawdown stage. It is the most valuable per-dollar account and has the longest compounding runway. Protecting it is non-negotiable.

> *The protocol exists to make good decisions on bad days.*

20.3 HWM update protocol

The high-water mark updates upward when total portfolio NAV at settled prices exceeds the prior HWM. It never moves downward.

When the HWM updates, recalculate all three drawdown stage levels immediately. Record the update in the session note. Only CSV-confirmed post-session NAV values propagate to the HWM. Forward-estimated or provisional NAVs do not.

Chapter 21

Exit Discipline: Four Conditions

The four exit conditions are checked at every session open for every Layer 2 or Layer 3 position, without exception, before any analysis begins. This is the most important operational discipline in the framework. Skipping this check is a governance breach.

EXIT DISCIPLINE CHECKLIST

Run at every session open, for every Layer 2 or Layer 3 position:

[ ] C1: EV deviation below -10% for two consecutive sessions? (Two-Session Rule)

[ ] C2: EV deviation below -15% this session? (S0 Hard Exit, immediate)

[ ] C3: Thesis-invalidation event triggered? (Layer 3 only)

[ ] C4: EV negative at today's binding quarterly review date? (Layer 3 only)

21.1 The S0 Hard Exit

The S0 Hard Exit is the most consequential rule in the exit discipline framework. If any Layer 2 or Layer 3 position shows an EV deviation worse than -15% at the session open, it must be exited that session. No waiting for confirmation. No "it might recover." The -15% threshold was the pre-committed exit point when the position was opened. It is not negotiable.

The reason for pre-commitment is that a -15% EV deviation means the market is pricing the position 15% above your probability-weighted expected outcome. Either your scenario weights are wrong and need to be updated, or the market knows something you do not. Neither is a reason to hold.

21.2 The time stop

The time stop is the most underused risk management tool in tactical investing. A position with negative EV at a quarterly review date must be exited, regardless of the thesis narrative.

The time stop prevents the common failure of holding a losing position indefinitely on the grounds that "the thesis hasn't played out yet."

When opening a Layer 3 position, set the time stop date before the first trade. The time stop date is typically the next quarterly review (Jan. 1, April 1, July 1, Oct. 1). The position must show positive EV by that date or be exited.

21.3 Thesis-invalidation events

A thesis-invalidation event is a specific, observable occurrence that proves the thesis wrong. It differs from C1 and C2 (which are price-based) and C4 (which is time-based). C3 is thesis-based. The underlying premise has been falsified by real-world events, regardless of what the price is doing.

A position can be exited on C3 even if the price has moved favorably, if the thesis that justified the position no longer holds, i.e., if the Strait of Hormuz opens. A C3 event does not automatically force an exit. It forces a mandatory review. The outcome of the review must be documented: either a new thesis is written, or the position is exited.

PART V

The Options Overlay

Covered calls generate systematic income from the passive core and defensive tactical positions. Tail hedges protect against sharp left-tail events. Both require discipline. Both fail badly when that discipline lapses.

Chapter 22

Covered Call Strategy

The covered call overlay generates systematic income from the passive core and defensive tactical positions. Implemented correctly, it adds 1% to 3% annually to portfolio returns with minimal disruption to the underlying thesis. Implemented incorrectly, it caps upside at exactly the wrong moments and creates assignment risk at unfavorable prices.

22.1 The hard rules

COVERED CALL HARD RULES

(1) No covered call may be written at a delta above 0.25 at time of order entry.

(2) Minimum 5% out-of-the-money (OTM) floor at entry. Certain primary factor ETFs may drop to a 2% floor when their options chain is thin and the reduction is documented.

(3) VIX (the CBOE Volatility Index) divided by VIX3M (its 3-month counterpart) of 1.0 or greater (backwardation): widen OTM floor to minimum 7.5%.

(4) No covered calls on Layer 3 positions while overweight.

(5) No covered calls in IRA accounts.

(6) Primary ETF must be at least 35% of taxable NAV to write covered calls on it.

22.2 Strike selection

Strike selection must balance three competing objectives: maximum premium income, minimal interference with the underlying thesis, and liquidity sufficient to actually fill the order. The most common mistake is selecting strikes with insufficient open interest. The order sits unfilled, the premium opportunity is missed, and the cycle expires uncovered.

THE CC LIQUIDITY PRE-CHECK (MANDATORY)

Before entering any covered call order, verify open interest and bid/ask size at the target strike via Market Chameleon (marketchameleon.com) or Barchart Options (barchart.com). Both offer free tiers. (Both websites active as of April 2026.)

Pull the live chain from Nasdaq.com or Cboe.com. Confirm bid at the target strike is at least $0.15.

If open interest is zero, if the bid is at or below $0.05, or if ask size is insufficient to fill your contract count, do not enter the order.

Lower the strike, reduce contract count, or wait for better conditions. Record this check in your session log.

22.3 AON vs non-AON

All-or-None (AON) orders require the entire contract quantity to fill at once. For low-liquidity underlyings with sparse options chains, AON is the primary fill blocker. If you experience multiple consecutive AON failures on the same position, the root cause is almost certainly liquidity, not pricing. The fix is to remove the AON condition and reduce contract count, not to adjust the limit price.

22.4 VIX regime adjustment

The VIX/VIX3M ratio signals the current volatility regime. Below 1.0 (contango) means near-term fear is lower than medium-term fear, a normal environment. At or above 1.0 (backwardation) means near-term fear exceeds medium-term fear, an elevated-stress environment. In backwardation, widen the OTM floor to 7.5% minimum.

22.5 Income targets

A realistic income target for a well-run covered call program on a passive ETF position is 0.5% to 1.0% of position value per month when the VIX sits in the 15-to-25 range. At VIX below 15, premiums compress significantly. At VIX above 30, consider reducing contract count and widening strikes rather than chasing premium.

Chapter 23

Tail Hedges and Protective Puts

A tail hedge is a long put position on a broad market proxy, held as portfolio insurance against a sharp left-tail event. It is not expected to be profitable under normal conditions. It is insurance. The premium paid is the cost of protection against a Stage 2 or Stage 3 drawdown event.

23.1 The naked long put

The most common misconception about long put positions is that you must own the underlying to benefit. You do not. A long put gives the right to sell the underlying at the strike price. If the underlying falls sharply, the put gains value and can be sold to close for a profit, regardless of whether you own the underlying. The put is closed by selling the contract, not by exercising and delivering shares.

23.2 Sizing and strike selection

A tail hedge should be sized to partially offset a Stage 1 or Stage 2 drawdown event. A common approach is to purchase puts with a notional value equal to 10% to 20% of total portfolio NAV, struck at 15% to 20% below the current index level, with three to six months to expiration. The total cost of the hedge should not exceed 0.5% to 1.0% of portfolio NAV per year.

23.3 Close threshold

Long puts decay with time. The standard close discipline is to close the position if the mark reaches 80% of maximum theoretical value, or on the approach to expiration if the position is significantly underwater.

Do not hold to expiration as a default. The last 30 days of an OTM put’s life are mostly theta decay with limited remaining value.

PART VI

Intelligence and Information

The quality of your governance decisions is limited by the quality of your information. Part VI covers the sources, the prediction markets, and the two cascade protocols (SS17 for geopolitical shifts, SS18 for macro shifts) that keep your scenario weights current.

Chapter 24

Information Sources

Most self-directed investors underinvest in their information architecture. They rely on financial news headlines and brokerage research when institutional-quality data sources are freely available.

24.1 Verified free sources

The table below uses standard abbreviations: IMF (International Monetary Fund), EIA (US Energy Information Administration), FRED (Federal Reserve Economic Data), ACLED (Armed Conflict Location and Event Data Project), IEA (International Energy Agency), Cboe (Chicago Board Options Exchange), FOMC (Federal Open Market Committee), and IV (implied volatility).

Source	URL	What It Provides	Use in Framework
IMF Portwatch	portwatch.imf.org	Real-time maritime traffic at key chokepoints. 7-day moving average transit calls.	Thesis-invalidation monitoring for energy/geopolitical positions.
EIA Weekly Petroleum	eia.gov/petroleum/supply/weekly	US crude oil inventories, production, imports, refinery utilization. Wednesdays 10:30 ET.	Energy sector thesis monitoring. WTI price context.
FRED	fred.stlouisfed.org	800,000+ economic data series. Rates, CPI, PCE, GDP, yield curve.	Rate-sensitive position EV model inputs. FOMC context.
JPMorgan Guide to Markets	am.jpmorgan.com	Quarterly institutional-quality charts: valuations, earnings, sector performance.	Quarterly review context. Sector valuation benchmarking.
ACLED	acleddata.com	Armed conflict event data, geocoded and updated daily.	Geopolitical cluster thesis monitoring.
IEA Chokepoints Monitor	iea.org	Maritime chokepoint shipping flows drawing on Portwatch data.	Energy/shipping thesis secondary confirmation.
Market Chameleon	marketchameleon.com	Options chain OI, volume, IV by strike. Free tier.	MANDATORY CC pre-check before every options order.
Barchart Options	barchart.com	Options chain data, OI, volume, bid/ask depth. Free tier.	Secondary CC liquidity confirmation.
Nasdaq.com	nasdaq.com	Live options chain, bid/ask/size per strike.	Primary live chain for CC entry.
Cboe.com	cboe.com	Live options chain, OI, IV per strike.	Backup live chain for CC entry.
Polymarket	polymarket.com	Largest prediction market. Geopolitical contracts with volume.	Scenario weight revision inputs. SS17 trigger.
Kalshi	kalshi.com	Regulated US prediction exchange.	SS17 trigger. Cross-

Source	URL	What It Provides	Use in Framework
			reference with Polymarket.
The Economist	economist.com	Weekly macro and geopolitical analysis.	P1 market intelligence. Thesis monitoring.
Wikipedia (conflict pages)	en.wikipedia.org	Continuously updated conflict timelines with sourced chronology.	Rapid geopolitical orientation at session open.

24.2 Premium sources worth considering at scale

At portfolio NAV above USD $1,000,000 or so, certain premium data sources become cost-justified. Energy Aspects provides specialist oil-market research including chokepoint-specific scenario modeling used by trading desks. A Bloomberg Terminal provides real-time data across every asset class.

These are not necessary at smaller portfolio sizes. The free sources above cover the framework's needs adequately.

Chapter 25

Prediction Markets as Governance Inputs

Prediction markets are among the most underused tools in self-directed portfolio management. They aggregate the probability assessments of large numbers of informed participants, with real money at stake, into a single probability estimate that is often more accurate than expert consensus. As of April 2026, Kalshi and Polymarket are the most common in the US, with their odds often available through a simple search query.

25.1 What prediction markets are

A prediction market contract is a binary option on a real-world event. If you buy "Yes" on "Will the Strait of Hormuz return to normal traffic by June 1?" at 67 cents, you receive $1 if it resolves Yes and $0 if it resolves No. The price (67 cents) implies a 67% market-consensus probability. That probability is useful as a governance input even if you never trade the contract.

25.2 How they integrate into the framework

Prediction market probabilities feed directly into the scenario weights of the EV model. If a geopolitical event is the primary driver of a tactical position's thesis, the probability of that event resolving favorably is a first-order input to the scenario weights.

When that probability shifts materially, the scenario weights must be updated. The framework uses a 10%p shift rule: any prediction market probability relevant to a tactical position's thesis that shifts by 10%p or more from the prior session reading triggers a mandatory cascade review of all related positions. That is the SS17 protocol, covered in the next chapter.

25.3 Practical notes

- Polymarket and Kalshi may show different probabilities for similar contracts due to differing resolution criteria. Read both and note discrepancies.
- Check volume. A contract with $100,000 in volume is meaningfully more informative than one with $5,000.
- Resolution criteria matter enormously. Two seemingly similar contracts about the same event can resolve differently based on their specific language.
- Prediction markets can be thin in specific contracts. Cross-reference with physical data before making large weight changes.

Chapter 26

The SS17 Geopolitical Cascade

The SS17 cascade protocol is triggered when any prediction market probability relevant to a tactical position's thesis shifts by 10%p or more from the prior session reading. It mandates a systematic review of all tactical positions, not just the one most obviously affected.

26.1 Why a cascade review

Geopolitical and macro events rarely affect only one position. A shift in a key probability (Hormuz normalization, ceasefire extension, election outcome) affects not just the most obvious position but also related positions through shared scenario drivers.

The cascade review ensures that no position's scenario weights become stale because attention was focused on the most obvious one.

26.2 The protocol

SS17 PROTOCOL

(1) Identify the triggering shift: which contract, what probability, vs. what prior reading.
(2) Review ALL tactical positions.
(3) For each position, ask: does this shift change the probability of any of my five scenarios?
(4) Update weights only where warranted.
(5) Recalculate EV for every position with changed weights.
(6) Record written rationale for each weight change (and non-change) in the session note.
(7) Check all four exit conditions against updated EVs.

26.3 Written rationale requirement

Every SS17-triggered weight revision must be accompanied by written rationale explaining why the probability shift changes the scenario weighting for that specific position. This requirement exists to prevent mechanical weight-chasing, adjusting weights to fit the narrative of the moment rather than the underlying fundamentals of the thesis.

26.4 Prediction markets vs physical evidence

Prediction markets are often right. They are not always right. When prediction-market probability moves in one direction and physical evidence (Portwatch traffic data, EIA inventory data, ACLED conflict counts) moves in the opposite direction, the physical evidence wins, or at minimum earns a cautious weighting.

Document the conflict in the session note. State which signal you are weighting more heavily and why. A cascade that blindly follows a prediction market against contradictory physical evidence is a cascade that will eventually be wrong on exactly the moves that matter.

Chapter 27

The SS18 Macro Cascade

SS18 is the macro parallel to SS17. Where SS17 fires on geopolitical probability shifts, SS18 fires on macro indicator crossings. The two run in parallel and produce independent Macro EV readings for every position.

27.1 The five macro scenarios

The scenarios are distinguished by combinations of key macro indicators: the 10-year (10yr) US Treasury yield, the Federal Reserve's policy rate (Fed funds), US real gross domestic product (GDP) growth, and the Consumer Price Index (CPI) inflation rate. The Personal Consumption Expenditures index (PCE) is an alternative inflation measure. Later in the chapter, West Texas Intermediate crude oil futures (WTI) and the US Dollar Index (DXY) enter as trigger thresholds. Basis points (bp) measure one-hundredths of a %p: 20bp means 0.20%.

Scenario	Weight	10yr Yield	Fed Funds	GDP	CPI	Portfolio Impact
Hard Bear	~10%	>5.0% or <3.5%	Cuts >100bp or Hike >50bp	<0%	>4.5% or <1%	Risk-off. Defensives and gold bid. Growth hit.
Soft Bear	~28%	4.5–5.0%	Hike 25–50bp	0%–1.5%	3.8%–4.5%	Modal regime. Defensives outperform.
Base	~42%	4.0–4.5%	Hold or 1 cut	1.5%–2.5%	2.8%–3.8%	Moderate growth. Balanced.
Soft Bull	~15%	3.5–4.0%	2–3 cuts	2.5%–3.5%	2.0%–2.8%	Growth and factor-tilted core benefit.
Hard Bull	~5%	3.0–3.5%	4+ cuts	>3.5%	<2.0%	Growth surge. Defensives headwind.

The weights above are illustrative. They must be calibrated to the current macro environment at the start of every SS18 cycle. When a trigger fires, the weights shift per a pre-committed cascade direction.

27.2 Trigger thresholds

Trigger	Condition	Action
Rate Shift	10yr Treasury yield moves ≥ 20bp from baseline (prior session or prior FOMC reference)	Recompute Macro EV. Assess cascade direction.
Inflation Surprise	CPI or PCE release surprises ≥ 50bp vs consensus	Mandatory SS18 review. Cascade if surprise direction changes weights.
Fed Policy	Fed funds futures shift ≥ 25bp	Active SS18 watch. Cascade if confirmed.

Trigger	Condition	Action
Oil / Inflation Pre-Commit	WTI closes above pre-committed level (e.g., $90) on two consecutive US sessions	Apply pre-committed weight shift (e.g., SBear +2 / SBull -2).
Recession Signal	Prediction-market probability of 2026 recession crosses pre-committed threshold	Cascade toward bear scenarios.
Fed Chair / Regime Shift	New Fed chair confirmed or material policy-stance shift	Formal cascade. Review weights against new reaction function.

27.3 How SS17 and SS18 interact

A position carries two Expected Value readings: one from SS17 (geopolitical-weighted scenarios) and one from SS18 (macro-weighted scenarios). Both are computed. Both are reviewed. Neither is summed into a composite. They answer different questions and should be held as distinct signals.

When SS17 and SS18 agree on a position's direction, conviction rises. When they disagree, the disagreement itself is the signal. Investigate it. One framework is likely missing something the other is seeing.

The four exit conditions (Chapter 21) run against both EV readings. If a position breaches C2 (-15% deviation) on either SS17 Geo EV or SS18 Macro EV, the S0 Hard Exit fires. Exit discipline does not wait for both frameworks to agree.

PART VII

Session Management

The session is the unit of portfolio management. Every material decision happens within a session and is documented in a session note. Part VII covers the three-prompt session structure and the four living documents that anchor the framework.

Chapter 28

The Three-Prompt Session

A governance session is the unit of portfolio management. Every material decision, buying, selling, adjusting options, updating scenario weights, checking constraints, happens within a session and is documented in a session note.

Sessions are not daily. Most investors will run full governance sessions once or twice per week, with abbreviated check-in sessions on market-moving days.

28.1 Session structure

This second edition replaces the earlier four-prompt structure (P0/P1/P2/P3) with a tighter three-prompt flow. Pre-session setup is folded into Prompt 1.

Phase	Name	Contents
Prompt 1	Market Intelligence	Run price data. Upload CSVs. Geopolitical and macro developments. Prediction market readings. SS17 and SS18 trigger checks. Physical data refresh. Cascade proposal if any trigger fires.
Prompt 2	Portfolio Diagnosis	EV check for all tactical positions. Four exit conditions. Constraint dashboard update. NAV and HWM check. Active violation status. Options overlay status. IRA composition.
Prompt 3	Portfolio Actions	Proposed orders with governance citations. CC liquidity pre-check. Order sequencing. Pre-commitments for next session. Session state note draft. 30-day action calendar.

28.2 The session opening checklist

MANDATORY SESSION OPENING STEPS

Before any analysis:

(1) Run price data script. Paste live prices.

(2) Confirm all good-till-cancelled (GTC) order statuses.

(3) Check all four exit conditions for all tactical positions.

(4) Update constraint dashboard.

(5) Check HWM against current NAV.

(6) Pull fresh brokerage CSV exports.

(7) Check prediction market readings for SS17 trigger.

(8) Check macro indicators (10yr, VIX, WTI, DXY) for SS18 trigger.

28.3 The 30-day action calendar

Every session generates a 30-day action calendar, a prioritized list of actions required in the next 30 days, with their dates, required actions, and consequences of non-completion.

The consequence column is mandatory. It makes the cost of inaction explicit and converts the calendar from a wish list into a governance document.

28.4 Timing

Sessions should be timed to allow for overnight order execution; market orders go off when the next trading day begins. An optimal pattern is to run your session in the evening, after hours, before the next day's trading day, and confirm fills at the next session open. This rhythm separates analysis from execution and avoids the reactive decision-making that comes from watching prices move in real time. Do not watch business news in real time during market hours, and never, ever base a trade off of a headline seen in a chyron.

Investors in East Asian time zones should note the convention explicitly. "Tomorrow's open" means the next NYC market session, not tomorrow local time.

28.5 The kill-switch

Once per quarter, run the kill-switch check. If the tactical book has underperformed a named broad benchmark (total US market index) by more than 5% over the trailing 12 months, or if the Sortino ratio of the tactical book has fallen below 0.75, flag the kill-switch condition and enter a formal review.

Two consecutive quarterly reviews with the kill-switch flagged produces a binding decision: collapse the tactical book into Layer 1 and rebuild only if the review concludes the tactical thesis set has changed materially and the scenario weights are being systematically revised.

Chapter 29

Document Architecture

The framework requires four living documents, regenerated at each full governance session. The separation of policy from operational data is the most important structural principle. The Investment Policy Statement contains only policy. It never holds session-state data.

29.1 The four documents

Document	Version Convention	Contents	Update Frequency
Investment Policy Statement (IPS)	Major.Minor. Major = policy change. Minor = session update.	Policy only: constraints, layer rules, drawdown protocol, options overlay rules, obligation map, account roles, information sources.	Minor bump every session. Major version on policy changes only.
Systematic Portfolio Management (SPM)	Full version refresh each session.	Operational engine: violation narrative, three-prompt protocol, EV tables, order register, escalation status, 30-day calendar.	Every full session.
Portfolio State Workbook	v1.N, updated each session.	Spreadsheet: constraint dashboard, holdings, EV model, order register, IRA governance, calendar, opening protocol.	Every full session.
Session State Note (SSN)	Dated by session.	Point-in-time record: what happened, decisions made, open orders, violations, opening protocol for next session.	Every full session. Uploaded at next session open.

29.2 Version control

Version control is mandatory. Every document carries a version number. The IPS version tracks policy changes. A new major version means a constraint, rule, or structural element changed. Minor version bumps are session updates. The amendment log in the IPS records every version with a summary of what changed.

29.3 The upload protocol

UPLOAD AT EVERY SESSION

(1) IPS (current version).
(2) SPM (current version).
(3) Portfolio State Workbook.
(4) Session State Note (prior session).
(5) Holdings CSV for every account.

(6) Portfolio Summary CSV for every account.
(7) Unrealized Gain/Loss CSV for every account.
(8) Pending Activity CSV for every account.
(9) Balances CSV for every account.
(10) Price data output pasted at the start of Prompt 1.

29.4 Why documents, not memory

A governance framework that lives only in memory is not a framework. It is an intention. Documents are the mechanism by which a governance decision made on one day binds the investor on a different day, under different market conditions, under different emotional pressure.

> *Undocumented decisions do not exist.*

This is the single most important operational discipline in the framework. Every session produces documents. Every amendment bumps a version. Every SS17 cascade logs its written rationale. Every exception logs its reasoning. The documents are the framework.

The record-keeping discipline in this framework borrows from the regulated-professional world. Licensed brokers and compliance officers keep contemporaneous notes because those notes are subpoenable: they will be read out in court, on the record, and they are the primary defense when a client disputes a trade or a regulator asks why a position was opened. The self-directed investor has no such legal exposure, but the same discipline serves a different purpose. Documents written at the time a decision is made are the only defense against a future version of yourself who no longer remembers why the position was sized that way, why the thesis read that way, or whether the trigger that just fired was the one you actually committed to. Save the four files in at least two locations, a cloud service and a local drive, and treat the act of writing them as part of the decision itself, not a chore after it.

APPENDICES

Reference Material

Appendix A

Workbook Templates

Blank templates for direct use. Copy each into your own spreadsheet or session note and fill in the values for your portfolio.

A.1 Constraint dashboard template

Constraint	Limit $	Actual $	Actual %	Headroom $	Status
Primary ETF Floor (≥30% taxable)					
CC Eligibility Floor (≥35% taxable)					
Broad International (≥7% taxable)					
International Total (≥10% taxable)					
Passive Core Total (≥55% taxable)					
Passive > Tactical					
Cluster Cap 1 (≤20%)					
Cluster Cap 2 (≤20%)					
Single Position Cap (≤15%)					
Cash Floor (≥$5,000)					

A.2 EV model template

Ticker	HBear Tgt/Wt	SBear Tgt/Wt	Base Tgt/Wt	SBull Tgt/Wt	HBull Tgt/Wt	EV $	Live $	Dev %

A.3 Order register template

Order #	Security / Action	Type	Price	Est. Gross	Status / Notes

A.4 IRA holdings template

Account	Security	Qty	Price	Value	Sleeve	% IRA	Notes
Trad IRA					Gold/Alternatives		
Trad IRA					Passive Core Ext		
Trad IRA					Legacy Income		
Roth IRA					Passive Core Primary		
Roth IRA					Passive Core Intl		

A.5 Exit conditions per position template

Ticker	C1 Threshold	C2 Threshold	C3 Invalidation Trigger	C4 Time Stop
	EV < -10% x2	EV < -15%		
	EV < -10% x2	EV < -15%		
	EV < -10% x2	EV < -15%		
	EV < -10% x2	EV < -15%		

A.6 30-day action calendar template

Date / Window	Action Required	Consequence of Non-Completion

Appendix B

Quick Reference Cards

Single-page cards designed to sit in the session note. Print them. Laminate them. Use them as the first thing you check when a session opens or when a trigger fires in the middle of the week.

B.1 Session opening checklist

- Run the price data script. Paste live prices at the top of Prompt 1.
- Upload: IPS, SPM, Workbook, prior Session State Note, all CSV exports.
- Confirm every GTC order status in the brokerage interface.
- C1 check: any position with EV deviation below -10% for two consecutive sessions?
- C2 check: any position with EV deviation below -15% this session? (S0 Hard Exit)
- C3 check: any thesis-invalidation events triggered?
- C4 check: any Layer 3 position with EV below 0% at today's binding quarterly review date?
- Update the constraint dashboard. Flag any breaches.
- Check HWM. Update if NAV exceeds the prior HWM (CSV-confirmed only).
- Pull prediction market readings. Check SS17 trigger (≥ 10%p shift).
- Check macro indicators (10yr, VIX, WTI, DXY). Check SS18 trigger thresholds.
- Check physical data: Portwatch 7dMA, EIA Wednesday release, ACLED.

B.2 SS17 cascade protocol card

- Identify the triggering shift: which contract, what probability, vs. prior reading.
- Review ALL tactical positions, not only the most obviously affected.
- For each position, ask: does this shift change any scenario probability?
- Update weights only where warranted. Do not mechanically chase the headline.
- Recalculate EV for every position with changed weights.
- Write rationale for each weight change (and every non-change) in the session note.
- Re-run all four exit conditions against updated EVs.
- Record the SS17 trigger and outcome in the session documents.

B.3 SS18 macro cascade card

- Identify which macro threshold was crossed. Read the indicator. Confirm the level.
- For pre-committed triggers (e.g., WTI >$90 for two sessions), apply the pre-committed cascade direction.
- For non-pre-committed triggers, propose the cascade direction with written rationale.
- Recompute Macro EV for every position.

- Re-run all four exit conditions against updated Macro EVs.
- Document the cascade in the session note. Update SS18 weight vector in the SPM.

B.4 Covered call pre-entry checklist

- Verify primary ETF is at least 35% of taxable NAV.
- Check VIX/VIX3M ratio. If ≥1.0, widen OTM floor to 7.5%.
- Calculate 5% OTM floor: current price x 1.05 = minimum strike.
- Check delta at target strike. Must be ≤ 0.25.
- Run Market Chameleon or Barchart: OI and ask size at target strike.
- Pull live chain from Nasdaq.com or Cboe.com. Confirm bid ≥ $0.15.
- If OI is zero, bid is at or below $0.05, or ask size is insufficient: do not enter.
- Remove AON condition. Use GTC limit order.
- Record the liquidity check in the session log before order entry.

B.5 Drawdown protocol quick reference

Current NAV vs HWM	Stage	Immediate Actions
Above -15%	Stage 0 Normal	Normal operations. All layers active.
-15% to -20%	Stage 1	Suspend new tactical buys. Maximize CC income. No Layer 3 additions.
-20% to -25%	Stage 2	Exit all Layer 3. Reduce to Layer 1 + Layer 2. Maximize CC on remaining.
Below -25% (or ruin threshold)	Stage 3 Ruin	Exit entire tactical book. Layer 1 passive core only. Reassess obligation.

Appendix C

Glossary

Definitions in the order an investor encounters them when working through the framework.

Term	Definition
AON (All-or-None)	An order condition requiring the entire quantity to fill at once. Problematic for low-liquidity options chains. The primary cause of consecutive failed fills.
Adversarial Advantage Statement	A one-paragraph written justification required before opening any Layer 3 position. States why you have an edge in the position that the market does not already price.
Backwardation (VIX)	A condition where VIX/VIX3M ≥ 1.0, meaning near-term implied volatility exceeds medium-term. Signals elevated near-term market stress. Requires widening OTM floor on covered calls.
CAGR	Compound Annual Growth Rate. The rate at which a portfolio must grow each year to reach a target value from a starting value in a given number of years.
Cascade Review	A systematic review of all tactical positions triggered by an SS17 or SS18 event, ensuring no position's scenario weights become stale due to a shared underlying driver.
Cluster Cap	A hard constraint limiting the combined value of thematically related positions to ≤ 20% of taxable NAV.
C1 / C2 / C3 / C4	The four exit conditions. C1: two-session -10% rule. C2: single-session -15% S0 Hard Exit. C3: thesis invalidation (Layer 3). C4: time stop (Layer 3).
Contango (VIX)	A condition where VIX/VIX3M < 1.0, meaning near-term implied volatility is lower than medium-term. A normal, low-stress environment. Standard CC parameters apply.
Delta	A measure of an option's price sensitivity to the underlying asset's price. A delta of 0.25 means the option price moves $0.25 for each $1 move in the underlying. Hard limit: no CC above 0.25 delta.
DRIP	Dividend Reinvestment Program. Automatically reinvests dividend payments into additional shares. Enable on all Layer 1 positions in all accounts.
EV (Expected Value)	A probability-weighted average of all possible outcomes. EV = sum of (scenario weight x scenario price target) across all five scenarios.
EV Deviation	(EV - Live Price) / Live Price x 100%. Positive: position is cheap vs. expected outcome. Negative: expensive vs. expected outcome.
FEIE	Foreign Earned Income Exclusion. A US tax provision allowing qualifying expatriates to exclude foreign-earned income from US taxable income, potentially enabling low-cost Roth conversions.
GTC	Good Till Cancelled. An order type that remains active until filled or explicitly cancelled.
HWM (High-Water Mark)	The highest total portfolio NAV ever recorded at CSV-confirmed settled prices. Drawdown stages are measured from HWM. Updates upward only.
IPS	Investment Policy Statement. The policy-only governance document. Never contains session-state data.

Term	Definition
Kill-Switch	A quarterly check that flags when the tactical book has underperformed the benchmark by more than 5% over trailing 12 months, or when Sortino ratio falls below 0.75. Two consecutive flagged quarters trigger tactical book liquidation.
Layer 1 Permanence Rule	Layer 1 (passive core) is never sold to fund any other layer. No EV model, no exit conditions, no exceptions.
LTCG	Long-Term Capital Gains. Gains on assets held more than 12 months, taxed at preferential rates rather than ordinary income rates.
Naked Long Put	A long put position held without owning the underlying. Closed by selling the contract, not by exercising.
OTM (Out of the Money)	An option where the strike price is less favorable than the current market price from the buyer's perspective. For calls: strike > current price.
Passive Core	Layer 1 holdings: factor-tilted ETFs held permanently across all three accounts.
Prediction Market	A market where participants trade contracts on the outcomes of real-world events. Prices reflect consensus probability estimates with real money at stake.
Ruin Threshold	The drawdown level at which the required CAGR to reach your target becomes effectively unachievable. Defines Stage 3 of the drawdown protocol.
Roth Conversion Ladder	The systematic conversion of Traditional IRA assets to Roth in low-income years, to shift assets permanently into the tax-free wrapper.
S0 Hard Exit	Immediate mandatory exit triggered when EV deviation falls below -15% in any single session. Non-negotiable.
Sequence Rule	The rule preventing new tactical buys when the passive core floor (55%) or primary ETF floor (30%) is in breach.
SPM	Systematic Portfolio Management document. The operational engine: session procedures, EV tables, order register, calendar.
SS17	The geopolitical cascade protocol triggered by a ≥10%p prediction market probability shift relevant to any tactical position's thesis.
SS18	The macro cascade protocol triggered by macro indicator thresholds: 10yr yield shift, inflation surprise, Fed policy, oil/inflation pre-commit, recession signal, Fed chair/regime shift.
SSN	Session State Note. The point-in-time session record. Uploaded at the start of the following session.
STCG	Short-Term Capital Gains. Gains on assets held 12 months or less, taxed at ordinary income rates.
Tail Hedge	A long put position on a broad market proxy, held as portfolio insurance. Not expected to be profitable in normal conditions.
Thesis-Invalidation Event	A specific, pre-defined observable event that falsifies the thesis behind a Layer 3 position, triggering mandatory exit review.
Time Stop	A pre-committed deadline by which a Layer 3 position must show positive EV or be exited. Set before opening the position.

THE SYSTEMATIC PORTFOLIO FRAMEWORK

Second Edition

A Governance Framework for the Self-Directed Investor

By Gregory C. Eaves

April 2026

This guide is released for educational and entertainment purposes only. The framework described here represents a practical governance system for self-directed portfolio management. It is not financial advice. All investment decisions involve risk. You, too, can lose money, my friend, and you are not smarter than the market.

If you want to run this framework on your own portfolio, which I recommend you do, drag & drop the companion file ("The Systematic Portfolio Framework, Claude Operating Instructions") into a new AI chat conversation (ChatGPT, Claude, Gemini, DeepSeek, whatever; it doesn't matter), upload your brokerage CSV files, and then type "Run Prompt 1," let it do its work, read its output, answer its questions, then type "Run Prompt 2," let it do its work, read its output, answer its questions, and then finally write "Run Prompt 3," and read the results. See what it suggests.

For educational purposes only. Not financial advice.

If you have any questions, comments, or suggestions, or if you'd like to discuss things, please send me an email.

gceavesspf@gmail.com