

MUTANTDeFi RESEARCH | GRADUATE TEXT IN QUANTITATIVE FINANCE

Full-Distribution Financial Science

Full-distribution foundations for portfolio activity, trading, and future conviction

The argument of the volume



Mean–variance states a closed two-moment world. The rest of the volume replaces it with a hurdle-net reading and a mean–omega frontier.

Author and Publication Information

Author William Lloyd Gleim, MS

Affiliation MutantDeFi Research Lab, ConfigSecretAI, Inc.

Correspondence 4225 Executive Square, Suite 600, La Jolla, CA 92037, United States

Contact mutantdefi.com · X: @mutantdefi, @billgleim

Imprint MutantDeFi Graduate Texts in Quantitative Finance

Structure Volume I, Measurement · Volume II, Efficiency

Edition Second edition, August 2026

PUBLICATION SCOPE

This is an independent academic and professional treatment of full-distribution measurement, portfolio activity, and time-bound market state. Scope lives on this leaf, not as a stamp on every page.

Contents

Author and Publication Information	3
Abstract	8
Preface	9
I Measurement	11
Introduction: An Overview of the Arc	13
1 The Classical Composition: Markowitz, CAPM, and Value-at-Risk	16
1.1 Mean–variance selection	17
1.2 The Capital Asset Pricing Model	17
1.3 Value-at-Risk	17
1.4 The composition, and why it is a second-moment orthodoxy	18
1.5 Failure modes	18
1.5.1 Fat tails and the Mandelbrot critique	19
1.5.2 Quantile blindness and incoherence of VaR	19
1.5.3 What the orthodoxy cannot see	19
1.6 Time, intentionally absent	20
1.7 Problems	20
2 Third and Fourth Moments: The Case for Surf Kinematics	22
2.1 Skewness and kurtosis as first-class objects	22
2.2 The tie that variance cannot break	23
2.3 Where the higher moments live: price and volume	24

2.4	Design requirements for a higher-moment state reading	24
2.5	Problems	25
2.6	When shape becomes load-bearing	26
2.7	Optional extension: horizon-bound future-state valuation	27
2.8	Market classes with higher-moment dominance	27
2.9	Why cryptocurrency markets belong to the class	28
2.10	Measurement discipline	29
2.11	Problems	29
3	Surf Kinematics: Heuristic Mathematics for Trading and Portfolios	31
3.1	The kinematic metaphor	32
3.2	The kinematic triad	32
3.3	Alignment gating and trade activity	33
3.4	Regime classification	35
3.5	The risk overlay: volatility compression	35
3.6	Portfolio-level heuristics	36
3.7	What the heuristics achieve, and what they lack	36
3.8	Problems	37
4	The Omega Ratio: A Full-Distribution Characterization That Subsumes the Kinematics	38
4.1	Definition and elementary properties	39
4.2	The hurdle net	40
4.3	Canonical coordinates	41
4.4	The subsumption theorem	43
4.5	Problems	45
	Formal Omega Function Theory	46
4.6	Nguyen's setup and the reciprocal convention	46
4.7	Properties without continuity assumptions	46
4.8	Standard dispersion, mixtures, and finite samples	47
4.9	Interpretation and limits	47

4.10 Problems	47
5 Pricing and Optimizing Horizon-Bound Future States	49
5.1 One objective, many mechanisms	49
5.1.1 Rule-based gating	50
5.1.2 Supervised learning	50
5.1.3 Reinforcement learning	51
5.2 Portfolio optimization: omega risk budgeting	51
5.3 Trading activity as information	52
5.4 Evaluation discipline	53
5.5 Problems	53
II Efficiency and the Frontier	55
To the Reader of Volume II	57
6 Gleim–Markowitz Efficiency and the GM! Frontier	59
6.1 Setup	60
6.2 The ranking identity	60
6.3 Numbered definitions	61
6.4 Menu and selection	64
6.5 Dual picture: the (G, L) Omega frontier	65
6.6 Relation to Theorem 6.1 and Table 6.1	66
6.7 What this chapter defines	66
6.8 What Volume I already supplied	67
6.9 Problems	67
7 Replacement Proposals: Omega Successors for the Classical Infrastructure	69
7.1 The replacement principle	70
7.2 Replacing Sharpe and Sortino: performance measurement	71
7.3 Replacing VaR: risk reporting and capital	73
7.4 Replacing CAPM: pricing and systematic exposure	74
7.5 Replacing the load-bearing mean–variance infrastructure	76

7.6	What the proposals do not claim	78
7.7	The GM futures canon	79
7.8	Interpretive benchmark: a full-distribution instrument and the VIX	81
7.9	Problems	82
A	Limited Empirical Exhibits	85
A.1	A.1 Exhibit 1: Loss-magnitude control (controlled A/B validation, 2025) . . .	85
A.2	A.2 Exhibit 2: Gate quality from kinematic alignment (program validation, 2026) . . .	85
A.3	A.3 Exhibit 3: The minimal canonical rule and out-of-sample discipline (pre-registered battery, 2026)	86
B	Peer-Review Colloquy: A Devil’s Advocate Reading	87
B.1	On the polarity (GM-1)	87
B.2	On hurdle dependence	87
B.3	On the mirror theorem (GM-3)	88
B.4	On the subsumption theorem (GM-4)	88
B.5	On the omega security line (GM-5)	88
B.6	On subadditivity and the combined hurdle	89
B.7	On estimation under fat tails	89
B.8	On the empirical exhibits	89
B.9	On compression’s upside cost	89
B.10	On comparing instruments of different epistemic kinds	90
B.11	On the reciprocal Omega convention	90
B.12	On private-sample constants	90
B.13	On the named frontier	90
C	Statistical Significance Roadmap	92
D	Data Availability and Contact	95
D.1	Availability boundary	95
	Endnote on the Name	97
	Bibliography	98

Abstract

Classical portfolio theory—Markowitz selection, the CAPM, and Value-at-Risk—rests on two premises this volume relaxes together: that the first two moments of the return distribution suffice, and that time may be left out. We develop a full-distribution, time-inherent alternative organized around the hurdle net and the Omega ratio. A surf-kinematic system of bounded, volume-aware state readings is shown to be a heuristic estimator of a canonical hurdle-net triad—mass, direction, and increment—through an exact mirror symmetry and a subsumption theorem. The resulting objective is installed across gating, learning, and portfolio construction, and is used to propose like-for-like successors to Sharpe, Sortino, VaR, and related infrastructure.


The book is in two volumes. Volume I develops the measurement: hurdle net, Omega ratio, and a kinematic triad installed across gates, learners, and books. Volume II numbers the optimality objects Volume I draws. Tail-aware risk is the loss mass. Mean-omega dominance replaces mean-variance dominance. Omega-efficiency (Gleim–Markowitz efficiency) replaces Markowitz efficiency. The GM! Frontier is the image of the omega-efficient set in the tail-aware plane. Gleim–Markowitz optimality is the max- Ω selection on that set, characterized interiorly by the omega security line. The framework is called GM futures conviction theory: a research treatment of distribution-wide, time-bound information movement of which future conviction is a central object. It is not an index, product, feed, strategy, or software interface.

Keywords: portfolio selection; higher moments; Omega ratio; hurdle nets; distributional shape; time-inherent measurement; future conviction; Gleim–Markowitz efficiency; GM! Frontier.

Empirical exhibits are limited to Appendix A; data access is described in Appendix D.

Preface

Two-moment risk models fail in practice. That fact is motivation for replacement of the second-moment orthodoxy; it is not a foundation of the proofs that follow.



The Gaussian premise against fat-tailed reality.

Two-moment risk models fail in practice. That fact is motivation for replacement of the second-moment orthodoxy developed in Chapter 1; it is not a foundation of the proofs that follow.

This book is in two volumes. Volume I states a polarity and builds a measurement language. It shows that the first two moments of a return distribution do not suffice, that a hurdle net carries the rest of the shape, and that a single omega objective can be installed across gating, learning, and budgeting. It draws a mean-omega envelope (Figure 5.1, Table 6.1) and proves

an interior first-order condition (Theorem 6.1). Formal objects keep their established labels (Definitions 4.1–4.3, Propositions 4.1–4.2, Theorems 4.1–4.2 and 6.1, Principles 1.1 through 6.6, Tables 6.1–6.3, canon GM-1–GM-5).

Volume II numbers the objects Volume I used without defining: tail-aware risk, mean–omega dominance, omega-efficiency, the GM! Frontier, and Gleim–Markowitz optimality. It proves the ranking identity that places Ω on the same plane whose risk axis is the loss mass, then states the like-for-like replacements those objects complete.

The open mathematical edges remain open: estimator consistency of the kinematic triad, a subgradient security line under atoms and constraints, and sampling error of loss-mass estimators in the tail-index band the theory cares about. Hiding those edges would be a worse textbook.

Volume I

Measurement

MUTANTDEFI

IN THIS VOLUME

The two-moment system, then a full-distribution measurement: higher moments, a kinematic triad, the hurdle net, and one omega objective.

MUTANTDeFi

An Overview of the Arc

IN THIS CHAPTER

The volume moves from the two-moment orthodoxy through a kinematic triad to the hurdle net, omega, and the GM! Frontier.



The arc of the volume: two-moment orthodoxy to the GM! Frontier.

Quantitative finance spent its first half-century inside two moments. The mean carried reward; the variance carried risk; and an extraordinary edifice—efficient frontiers, equilibrium pricing, regulatory capital—was erected on the premise that the first two moments of the return distribution suffice. This book is about what happens when they do not, and about a constructive path from that failure to a working full-distribution practice.

The arc runs as follows. Chapter 1 presents the classical composition—Markowitz mean-variance selection, the Capital Asset Pricing Model (CAPM), and Value-at-Risk (VaR)—with the respect it deserves and the failure modes it earned: fat tails, undefined variances,

quantile blindness, and incoherence. Chapter 2 introduces the third and fourth moments as first-class objects and motivates surf kinematics: a family of bounded, volume-aware market-state readings built precisely because skewness and kurtosis carry decision-relevant information that variance discards. An interlude then locates the markets and mechanisms in which full-distribution methods are load-bearing. Chapter 3 develops the surf kinematic mathematics as a working system—the speed–velocity–acceleration triad, alignment gating, regime classification, and the non-parametric risk overlay we call volatility compression—applied to both trade activity and portfolio construction. Chapter 4 introduces the omega ratio of Keating and Shadwick as the full-distribution characterization that subsumes the kinematic heuristics: each member of the kinematic triad is shown to be an estimator of a canonical hurdle-net quantity with exact, provable symmetry properties. A second interlude records Nguyen’s measure-theoretic treatment of the Omega function. Chapter 5 completes the measurement program by carrying the omega objective across every mechanism and domain in which the kinematic heuristics ever operated—rule-based gating, supervised learning, reinforcement learning, portfolio risk budgeting, and the measurement of trading activity itself—so that one objective, stated once, governs the whole.

Volume II then numbers the efficiency, frontier, and optimality objects that Chapter 5 and Table 6.1 already use, and proves the algebraic bridge that makes omega a ranking primitive on the tail-aware plane. Chapter 7 carries the program to its institutional conclusion: like-for-like replacement proposals for the CAPM, VaR, the Sharpe and Sortino ratios, and the load-bearing mean–variance infrastructure that presumes them.

Two departures, given equal footing throughout, define the post-Sharpe class this book develops. The first is shape: the full return distribution replaces its first two moments. The second is time: canonical single-period instruments intentionally collapse event order, whereas the readings developed here are indexed, differenced, and evaluated through windows of a distribution in motion.

Throughout this volume the framework is called GM futures conviction theory (or, for short, GM futures theory). Its core claim remains full-distribution and time-inherent: distributional shape and the order of change matter. The framework does not require a forecast horizon, a conditional pricing model, or an exchange-traded futures contract.

One useful application asks a forward question: not merely “how is the asset priced

today?” but “how should a state at a declared future horizon be characterized and valued from information available today?” In that application, a conditional distribution can be evaluated through the same hurdle-net and Omega apparatus. The dedicated treatment later in the book is an extension of the theory, not its definition.

The name denotes a research framework, not a product, index, feed, strategy, or implementation. Exchange-traded futures are one important domain, not the framework’s boundary.

The pedagogical claim of the book is simple to state. The heuristics of Chapter 3 estimate the object of Chapter 4. Naming that object prescribes the canonical replacements and unifies three parallel engineering traditions into a single optimization discipline. The reader who finishes Chapter 5 should be able to look at a bounded activity reading, a signed momentum channel, or an asymmetric learning reward and recognize, in each, the same distributional object wearing different clothes. The reader who finishes Chapter 6 should be able to state which books are undominated on the tail-aware plane and which book, if any, is Gleim–Markowitz optimal.


A note on evidence. This is a book of principles, and it deliberately excludes implementation identifiers of every kind. Where an empirical result is necessary to make a case, it is summarized in Appendix A in limited, program-level form; complete datasets, protocols, and validation records are available from MutantDeFi as described in Appendix D.

1

The Classical Composition: Markowitz, CAPM, and Value-at-Risk

IN THIS CHAPTER

Mean, variance, and a quantile form a closed system. This chapter states that system so the rest of the volume can refuse it.



The classical mean-variance envelope that Chapter 1 then refuses.

1.1 Mean–variance selection

Modern portfolio theory begins with Markowitz’s 1952 insight that portfolio selection is a joint problem in reward and risk, formalized as mean and variance [1]. Let $R \in \mathbb{R}^n$ be the vector of asset returns with mean $\mu = \mathbb{E}[R]$ and covariance $\Sigma = \text{Cov}(R)$, and let w be portfolio weights. The canonical program is

$$\min_w w^\top \Sigma w \text{ subject to } w^\top \mu = \mu^*, w^\top 1 = 1 \quad (1.1)$$

Tracing the solution over target returns μ^* yields the efficient frontier: the upper boundary of attainable (σ, μ) pairs, a hyperbola in mean–standard-deviation space. Two premises do all the work. First, that investors care about the distribution of wealth only through its first two moments (or, equivalently, that returns are elliptically distributed, of which the Gaussian is the canonical case). Second, that Σ exists, is stable, and is estimable.

1.2 The Capital Asset Pricing Model

If every investor solves (1.1) against the same inputs and can lend and borrow at a riskless rate r_f [2], equilibrium follows. Sharpe, Lintner, and Mossin showed that all investors then hold the same tangency portfolio—the market M —and that expected excess returns are priced by a single systematic exposure [3, 4, 5]:

$$\mathbb{E}[R_i] - r_f = \beta_i(\mathbb{E}[R_M] - r_f), \beta_i = \text{Cov}(R_i, R_M)/\text{Var}(R_M) \quad (1.2)$$

The Capital Asset Pricing Model (CAPM) is mean–variance logic promoted to a theory of prices: the only risk that earns a premium is covariance with the market, because all idiosyncratic variance diversifies away. Everything above the second moment is, by construction, unpriced.

1.3 Value-at-Risk

Where Markowitz and the CAPM allocate, Value-at-Risk (VaR) polices. Popularized by the RiskMetrics program and adopted into regulatory capital frameworks [12, 13], VaR summarizes the risk of a loss variable L by a single quantile. At confidence level $\alpha \in (0, 1)$,

$$VaR_\alpha(L) = \inf\{\ell \in \mathbb{R} : \Pr(L \leq \ell) \geq \alpha\}, \quad (1.3)$$

the smallest loss threshold that will not be exceeded with probability α . Its appeal lies in its communicability (one number, one horizon, one probability), and its composition with the mean–variance world is seamless: under Gaussian returns $VaR_\alpha = \mu_L + z_\alpha \sigma_L$, so that VaR is once again a two-moment statistic.




Figure 1.1: The Gaussian premise against fat-tailed reality. Panel (a), linear scale: a Gaussian and a Student- t ($\nu = 3$) density look like siblings. Panel (b), log-density scale: the tails diverge by orders of magnitude. The shaded region is tail mass the Gaussian model omits—precisely the mass that decides solvency.

1.4 The composition, and why it is a second-moment orthodoxy

The three components interlock. Markowitz defines efficiency in (μ, σ) ; the CAPM prices assets by their contribution to portfolio σ ; VaR reports risk as a scaled σ under the same distributional premise. Call this the second-moment orthodoxy: a closed, mutually reinforcing system in which the variance (or its square root) is simultaneously the choice variable, the priced quantity, and the reported exposure.

1.5 Failure modes

1.5.1 Fat tails and the Mandelbrot critique

Mandelbrot demonstrated in 1963 that speculative prices exhibit heavy tails, clustering, and scale-invariance inconsistent with the Gaussian premise [8], findings extended by Fama [9] and codified as stylized facts by Cont [11]. Under stable laws with tail index $\alpha < 2$,

$$\Pr(|X| > x) \sim x^{-\alpha}, 1 < \alpha < 2, \quad (1.4)$$

the variance is infinite: program (1.1) minimizes a quantity that does not exist, and β in (1.2) divides by it. Even in the tamer finite-variance regime of Student-t returns, the fourth moment diverges for degrees of freedom $\nu \leq 4$, so sample variances converge slowly and tail probabilities are underestimated by orders of magnitude (Figure 1.1). Cryptocurrency markets often exhibit heavy tails, jumps, and clustered volatility, but their frequency and scale depend on asset, venue, sampling interval, and period. No universal extreme-event cadence is asserted here.

1.5.2 Quantile blindness and incoherence of VaR

VaR's defects are structural, not merely parametric. First, quantile blindness: (1.3) says nothing about the magnitude of loss beyond the threshold, which under fat tails is exactly where the risk lives. Second, incoherence: Artzner, Delbaen, Eber, and Heath showed that VaR fails subadditivity—the risk of a combined book can exceed the sum of its parts—so VaR can penalize diversification [14]. Expected Shortfall (Conditional VaR),

$$ES_\alpha(L) = \mathbb{E}[L | L \geq VaR_\alpha(L)], \quad (1.5)$$

repairs coherence and reads the tail's magnitude [15], and is the correct classical stepping stone—but it remains a one-sided summary of one tail at one level, not a characterization of the distribution.

1.5.3 What the orthodoxy cannot see

The common root of these failures is informational: the second-moment orthodoxy compresses the return distribution to (μ, σ) and therefore cannot distinguish distributions that differ only in asymmetry or tail weight. Chapter 2 makes that statement exact—exhibiting return

distributions the orthodoxy scores as identical that no rational trader would regard as interchangeable—and begins the constructive response.

1.6 Time, intentionally absent

There is a second omission, orthogonal to the moments, and it is deliberate. The 1952 program is explicitly single-period: allocate at the start, evaluate at the end [1, 6]. The CAPM prices one period of equilibrium; the Sharpe ratio compresses an entire return history into two numbers from which the ordering of events has been erased; VaR is a snapshot at a fixed horizon. Nothing in the orthodoxy observes the distribution move. This absence is a modeling choice, made knowingly for tractability, and it is honored here as such—dynamic extensions exist [7], but they restore dynamics to the investor’s problem while retaining the moment premise in the risk object. The instruments of this book take the opposite pole: time is inherent, not appended.

PRINCIPLE 1.1 (THE TIME POLARITY)

Sharpe-grade instruments are time-absent by intent: they summarize a return distribution as if it were delivered all at once. Post-Sharpe instruments in the GM futures frame are time-inherent by construction: their primitive quantities are readings, increments, and windows of a distribution in motion, and removing the time index destroys them. Portfolio activity theory is the study of the second pole.

Sharpe-grade instruments are time-absent by intent. Post-Sharpe readings are time-inherent; remove the time index and they vanish.

1.7 Problems

1.1 Derive the closed-form solution of (1.1) using two Lagrange multipliers, and show that the frontier is a hyperbola in (σ, μ^*) .

1.2 Show that under Gaussian returns VaR_α and ES_α are both affine in σ , and conclude

that in the Gaussian world ES adds communicable caution but no new information ranking.

1.3 Construct a two-position portfolio of short, deep out-of-the-money digital options for which $Var_{0.95}$ is superadditive. Interpret the failure in the language of diversification.

1.4 For a Student-t return with $\nu = 3$, compute the ratio of true $\Pr(X < -5\sigma)$ to its Gaussian estimate, and comment on Figure 1.1(b).

1.5 Show that the Sharpe ratio of a return series is invariant under any permutation of the series, and that the velocity and acceleration of Chapter 3 are not. State precisely what information the permutation destroys, and relate it to Principle 1.1.

MUTANTDeFI

2

Third and Fourth Moments: The Case for Surf Kinematics

IN THIS CHAPTER

A two-moment tie is not an economic tie. Skew and kurtosis force a state reading that can see the tails.

2.1 Skewness and kurtosis as first-class objects

For a return X with mean μ and standard deviation σ , define the standardized third and fourth moments

$$\gamma_1 = \mathbb{E}[(X - \mu)^3]/\sigma^3 \text{ (skewness)}, \gamma_2 = \mathbb{E}[(X - \mu)^4]/\sigma^4 - 3 \text{ (excess kurtosis)} \quad (2.1)$$

Skewness measures asymmetry between the gain and loss sides; excess kurtosis measures tail weight relative to the Gaussian. Neither is exotic to economic theory: expected utility with a positive third derivative implies a preference for positive skewness, and the general result of Scott and Horvath is that rational moment preferences alternate in sign—odd moments desired, even moments penalized [16]. Asset-pricing evidence agrees: systematic co-skewness commands a premium beyond beta [17, 18]. The orthodoxy of Chapter 1 does not deny these preferences; it is simply unable to express them.

2.2 The tie that variance cannot break

EXAMPLE 2.1 (A TWO-MOMENT TIE)

Figure 2.1 exhibits two return densities constructed as mirror-image skew-normal laws with identical mean ($\mu = 1.00$) and identical variance ($\sigma^2 = 2.42$). Every statistic of Chapter 1 scores them equal: same Markowitz coordinates, same Gaussian VaR at every level. Yet one carries its risk as a fat left tail ($\gamma_1 = -0.85$) and the other as a fat right tail ($\gamma_1 = +0.84$). A trader long the first and a trader long the second hold economically different positions; only the third moment can say so.

Every statistic of Chapter 1 scores these densities equal. Only the third moment can say they are different positions.




Figure 2.1: The two-moment tie of Example 2.1. Panel (a): mirror-skewed densities with identical mean and variance. Panel (b): variance ties; the third moment separates. Every instrument of Chapter 1 is indifferent between these positions.

2.3 Where the higher moments live: price and volume

Higher moments are not merely a property of the marginal return distribution; they are generated by market microstructure, and microstructure leaves a second observable trace: volume. A one-percent move on three times average volume and the same move on half average volume are different events—the first is conviction, the second is noise—but any statistic of price alone scores them identically. The empirical signatures of fat tails (clustered volatility, jumps, liquidity cascades) are volume-coupled phenomena [11]. A market-state reading that intends to respect the third and fourth moments must therefore read price and volume jointly.

The incumbent volume-aware object in institutional practice is the volume-weighted average price (VWAP), the standard benchmark for execution quality [10], and the comparison is instructive in both directions. VWAP demonstrates that volume-weighting is already accepted infrastructure; but VWAP is a first-moment, execution-side object—it answers “at what volume-weighted price did trading occur”—and it is time-collapsing in exactly the Sharpe-grade sense, compressing a session into one number. The kinematic use of volume is state-side and distribution-wide: volume enters as a conviction weight on higher-moment readings, instant by instant rather than as a session average. VWAP weights the first moment by volume; the kinematics of Chapter 3 weight the activity and asymmetry channels by volume. The two are complements, not competitors—one grades executions, the other reads state—and the contrast locates precisely what is new here.

2.4 Design requirements for a higher-moment state reading

The considerations above yield four requirements, and they are the specification that surf kinematics answers in Chapter 3.

1. Non-parametric normalization. Under heavy tails, moment-based normalizations (z-scores, standard deviations of standard deviations) inherit the estimation pathologies of Chapter 1. Rank and percentile statistics remain well-behaved for any tail index, so the state reading should be built from rolling ranks, not moments.
2. An undirected activity channel. Kurtosis is direction-blind: it says “the tails are active,”

not which one. The reading needs a bounded, undirected magnitude channel that rises whenever tail-generating activity rises.

3. A signed asymmetry channel. Skewness is directional. The reading needs a signed channel that is large only when directed movement is supported by activity—separating conviction from noise-driven oscillation.
4. A change channel. Higher-moment regimes turn quickly (volatility clustering, jump onset). The reading needs the rate of change of the signed channel, normalized so that “fast for these conditions” is comparable across regimes.

PRINCIPLE 2.1 (MOMENT-AWARE STATE READING)

A market-state representation adequate to third- and fourth-moment structure comprises a bounded undirected activity reading, a signed activity-weighted direction reading, and a normalized change reading, each constructed from joint price–volume information using rank-based (non-parametric) normalization.

Note the division of labor: the first three requirements answer Chapter 1’s shape omission, while the fourth answers its time omission—a change reading is destroyed by any permutation of the sample, so a state reading satisfying all four sits, by construction, on the time-inherent pole of the GM futures polarity (Principle 1.1).

Chapter 3 constructs exactly this object and puts it to work; Chapter 4 will reveal that the same three-channel shape is not an accident of engineering taste but the shadow of a canonical distributional decomposition.

2.5 Problems

2.1 Prove that for any random variable with finite third moment, a symmetric density implies $\gamma_1 = 0$, and give a counterexample to the converse.

2.2 Using Taylor expansion of expected utility to fourth order, derive the Scott–Horvath sign pattern: $U' > 0$, $U'' < 0$, $U''' > 0$, $U'''' < 0$ imply preference for mean and skewness and aversion to variance and kurtosis.

2.3 Construct (analytically or numerically) two densities with identical first three moments that differ in kurtosis, and explain which channel of Principle 2.1 distinguishes them.

2.4 Explain why a rolling percentile rank of a volatility reading is invariant to any strictly increasing transformation of that reading, and why this matters under unknown tail index.

Where Full-Distribution Methods Matter

2.6 When shape becomes load-bearing

Mean and variance are sufficient only under restrictive combinations of distribution, preference, and decision horizon. When those conditions fail, higher moments are not decorative refinements. They identify economically distinct ways in which the same mean and variance can be produced.

- Heavy or power-law tails. Extreme observations carry more probability mass than a Gaussian model permits; variance may converge slowly or fail to exist.
- Skewness and jump asymmetry. Downward gaps, upward squeezes, and capped or uncapped payoffs create different gain/loss geometry even when variance ties.
- Volatility clustering. Risk arrives in regimes. A single unconditional variance averages quiet and turbulent states that should not be treated as equivalent.
- Discontinuity and liquidity gaps. Prices jump across levels where a continuous diffusion assumes intermediate trading was possible.
- Path dependence. Liquidation, margin, stop, barrier, and rebalancing rules make event order part of the payoff rather than a nuisance.

PRINCIPLE 2.2 (FULL-DISTRIBUTION RELEVANCE)

Full-distribution methods become especially valuable when economically relevant probability mass is concentrated in asymmetric tails, when the distribution changes through time, or when the trading mechanism itself transforms the return distribution.

2.7 Optional extension: horizon-bound future-state valuation

At decision time t , declare a horizon h before estimating the object. For a horizon-state variable $Y_{t,h}$ and hurdle $\tau_{t,h}$, define

$$G_{t,h} = \mathbb{E}[(Y_{t,h} - \tau_{t,h})_+ | I_t], L_{t,h} = \mathbb{E}[(\tau_{t,h} - Y_{t,h})_+ | I_t], \Omega_{t,h} = G_{t,h}/L_{t,h}.$$

Within this application these are conditional future-state quantities. A current quote, spread, volatility reading, or order-book state belongs in I_t ; it is evidence about the future distribution, not a substitute for it. Changing h changes the economic question and generally changes the distribution, hurdle, dependence, and ranking.

Remark 2.1 (Scope). Horizon declaration is required when making a horizon-bound forecast or valuation. It is not a requirement of GM futures theory itself, which also covers realized, historical, contemporaneous, and rolling-window distributions.

2.8 Market classes with higher-moment dominance

No market label guarantees non-Gaussianity, and no market remains permanently in one regime. The classification is conditional: inspect the empirical distribution, market mechanism, and decision horizon before selecting the statistic.

Market class	Distributional property	Why two moments can mislead
Options and short-volatility books	convexity, truncation, rare large losses	small steady gains can mask a catastrophic left tail
Commodities and electricity	spikes, seasonality, storage constraints	physical bottlenecks create jumps and regime-specific tails
Emerging and frontier markets	liquidity gaps, controls, event risk	observed variance can omit non-trading and reopening jumps
Distressed credit	default atoms, recovery asymmetry	a continuous symmetric model cannot represent the payoff support
Catastrophe and reinsurance risk	rare, clustered, extreme losses	the loss tail, not ordinary-period variance, determines solvency
Leveraged and path-dependent products	compounding, barriers, liquidation	permuting returns changes wealth while moments remain unchanged
Illiquid alternatives	stale marks and appraisal smoothing	reported low variance may be a measurement artifact

2.9 Why cryptocurrency markets belong to the class

Cryptocurrency markets are an important example, but not the definition of the class. Their relevance follows from several structural properties.

- Continuous trading: information and deleveraging propagate through weekends and regional liquidity cycles.
- Fragmented venues and instruments: spot, perpetual swaps, dated futures, options, and lending transmit stress through basis, funding, and collateral.

- Endogenous leverage: liquidation engines turn price moves into forced order flow, producing feedback and asymmetric cascades.
- Variable market depth: identical notional orders can have different impact by asset, venue, and regime.
- Rapid structural change: listings, protocol events, custody failures, regulation, and participant composition can move the distribution itself.
- Observable activity channels: volume, open interest, funding, and cross-venue dislocations carry state information omitted by price-only summaries.

No universal cryptocurrency constant. Tail index, skewness, kurtosis, jump intensity, and dependence vary by asset, venue, sampling interval, and period. This text therefore does not promote a fixed threshold, cadence, or calibration derived from a private sample. Cryptocurrency is a vivid laboratory, not an exclusive domain.

2.10 Measurement discipline

A full-distribution workflow declares the return definition, sampling interval, horizon, and economic hurdle; inspects support, atoms, missing observations, and mark construction; estimates tail and asymmetry diagnostics with uncertainty; compares regime-conditional distributions; retains hurdle-net gain and loss masses; and stresses dependence, costs, slippage, and liquidity. A marginal distribution alone is not a trading system.

2.11 Problems

1. Give two distributions with equal mean and variance but different economic risk for a shortfall-sensitive investor. Identify the feature that separates them.
1. Explain why volatility clustering can make an unconditional fourth moment difficult to interpret even when it exists.
1. Map a liquidation cascade from initial shock to changed return distribution and identify the endogenous feedback step.
1. Name three non-cryptocurrency markets where full-distribution characterization is load-

bearing and justify each choice.

1. Design a rolling comparison of tail behavior in high- and low-liquidity regimes without assuming a fixed universal threshold.
1. State two reasons Omega can be estimated badly despite making no Gaussian assumption.
1. Explain why $\Omega_{t,h1}$ and $\Omega_{t,h2}$ need not rank two assets in the same order, even when both use properly scaled hurdles.


MUTANTDeFi

3

Surf Kinematics: Heuristic Mathematics for Trading and Portfolios

IN THIS CHAPTER

Speed, velocity, and acceleration are a heuristic triad—bounded readings of a distribution in motion.



Volatility compression amputates the left tail.

3.1 The kinematic metaphor

The kinematic metaphor framing is earned and not decorative. Raw market action is an ocean; a family of compression operators (bounded loss, trailing protection, session resets) creates a rideable surface; and the state reading tells the trader where to paddle in. The mathematical content of the metaphor is a decomposition of market state into speed (how much supported activity), velocity (which way, with how much support), and acceleration (how the direction is changing)—the three channels demanded by Principle 2.1.

3.2 The kinematic triad

Let r_t denote returns, u_t volume, and let $\phi_t = u_t / \text{MA}_w(u)_t$ be the relative-volume factor, where MA_w is a length- w moving average. Construct a small family of volume-coupled volatility readings—for instance volume-weighted return dispersion, volume-scaled average true range, and realized volatility scaled by ϕ_t —and let $X\%_t$ denote the rolling percentile rank of reading X over a trailing window.

DEFINITION 3.1 (SURF SPEED)

The speed is the clipped mean of the ranked readings,

$$S_t = \text{clip}((1/m) \sum_X X\%_t, 0, 1) \in [0, 1]: \quad (3.1)$$

a bounded, undirected, non-parametric reading of volume-supported activity, equal to zero at minimum observed activity and one when all readings are simultaneously extreme.

DEFINITION 3.2 (SURF VELOCITY)

The velocity weights signed returns by speed and smooths,

$$V_t = \text{MA}_{w_v}(S \cdot r)_t \quad (3.2)$$

V_t is large and positive only when price is rising and the rise is supported by elevated activity; noise-driven oscillation (large $|r_t|$, small S_t) produces negligible velocity.

DEFINITION 3.3 (SURF ACCELERATION)

With $\Delta V_t = V_t - V_{t-1}$,

$$A_t = \text{MA}_{w_a}(\Delta V)_t / (\text{MA}_{w_n}(|\Delta V|)_t + \varepsilon), \quad (3.3)$$

a dimensionless reading of directional change normalized by its own typical magnitude, so that “accelerating” means accelerating relative to current conditions.

3.3 Alignment gating and trade activity

DEFINITION 3.4 (ALIGNMENT GATES)

Long and short entry alignment are

$$\text{align}_t^+ = (V_t > 0) \wedge (A_t > 0), \text{align}_t^- = (V_t < 0) \wedge (A_t < 0): \quad (3.4)$$

the market must be moving in the entry direction with increasing conviction.

In program-level validation, replacing naïve momentum cascades (moving-average velocity and its difference, without activity weighting) with the kinematic gates reduced false entries—entries followed immediately by a protective stop—by roughly a third (Appendix A). Alignment is typically embedded in a weighted quality score: a linear combination of gate indicators (structural extrema, volatility band membership, alignment, trend-weakness,




Figure 3.1: The kinematic decomposition on a synthetic path with quiet, trending, and choppy regimes. From top: price with long-alignment windows shaded; speed S_t (bounded, undirected); velocity V_t (signed, activity-weighted); acceleration A_t (normalized change). Alignment requires agreement of the signed channels.

mispricing, volume confirmation) with entry above a threshold, and competitive selection when several instruments qualify at once.

3.4 Regime classification

Fat-tailed markets alternate between directional and mean-reverting (choppy) conditions, and the same signal deserves different treatment in each.

DEFINITION 3.5 (CHOP RATIO)

With a fast dispersion reading D_t^{fast} (e.g., a short-window average true range) and its slow average,

$$\chi_t = D_t^{fast} / \text{MA}_{w_c}(D_t^{fast})_t \quad (3.5)$$

$\chi_t > 1$ indicates volatility above its local norm—regime-transitional, choppy conditions; $\chi_t < 1$ indicates contraction—typically directional trend.

Classification uses thresholds with hysteresis (distinct entry and exit levels) to prevent flapping, and the regime conditions the risk overlay: stops tighten and entry thresholds rise in heavy chop, relax in clean trend.

3.5 The risk overlay: volatility compression

The kinematic reading decides where; the overlay decides how much can be lost. Its logic is non-parametric surgery on the payoff itself.

DEFINITION 3.6 (VOLATILITY COMPRESSION)

Volatility compression is the composition of (i) a hard loss bound $L < 0$, so realized return is $R_{bounded} = \max(R, L)$; (ii) trailing protection, which converts an excursion above a profit threshold into a ratcheting floor; (iii) profit capping at a session target; and (iv) calendar resets, which discretize risk into sessions and prevent multi-day compounding of error.

PROPOSITION 3.1 (DISTRIBUTIONAL EFFECT OF COMPRESSION)

Compression maps the return distribution to one that is left-truncated at L (with the truncated mass an atom at the bound), right-preserved up to the cap, and therefore right-skewed with bounded loss—raising γ_1 and bounding the loss tail regardless of the underlying tail index, including regimes where variance or expected shortfall of the raw return is infinite.

This is the decisive contrast with Chapter 1's VaR: a quantile reports the tail; compression amputates it. The theoretical signature is distributional rather than hit-rate based: a process can tolerate many small losing observations if its loss tail is bounded and its gain mass remains sufficient. Whether an implementation achieves that transformation is an empirical question confined to Appendix A.

3.6 Portfolio-level heuristics

The same grammar lifts to the portfolio. Quality scores induce a cross-sectional ranking for capital allocation; per-position compression bounds each marginal loss; and a portfolio slope protection rule—exit all positions when aggregate open profit retreats by a buffer from its running peak—bounds the portfolio drawdown by construction, playing the role that no covariance estimate can play when second moments are unstable or undefined.

3.7 What the heuristics achieve, and what they lack

The kinematic system satisfies every requirement of Principle 2.1 and is empirically load-bearing. But it is a heuristic system: its channels are built from windows, ranks, and smoothing constants; its gates are tuned; and nothing in its construction names the distributional object being estimated. Three consequences follow. The channels are coordinate-dependent (change the readings and the triad changes); the gates are not portable across mechanisms (a rule threshold, a learning target, and a policy reward must each be re-engineered); and there is no exact symmetry to appeal to when asking whether long and short are treated evenhandedly. Chapter 4 removes all three limitations at once, by identifying the object the triad was estimating all along.

3.8 Problems

3.1 Show that S_t of Definition 3.1 is invariant to any strictly increasing transformation applied to each underlying reading, and explain the significance under unknown tail index (cf. Problem 2.4).

3.2 Prove Proposition 3.1: compute the skewness of $\max(R, L)$ for R symmetric about 0 and $L < 0$, and show it is positive.

3.3 A trader proposes replacing the hysteresis of the chop classifier with a single threshold. Construct a dispersion path on which the single-threshold classifier switches regimes every bar while the hysteresis classifier switches once.


3.4 Portfolio slope protection bounds drawdown without covariances. Exhibit a two-asset return sequence on which mean-variance risk parity increases allocation to the asset about to crash while slope protection exits both.

4

The Omega Ratio: A Full-Distribution Characterization That Subsumes the Kinematics

IN THIS CHAPTER

The hurdle net measures; omega ranks. Together they give an exact mirror of a configuration and its long–short opposite.



The hurdle net reads gain and loss mass on the whole distribution.

4.1 Definition and elementary properties

DEFINITION 4.1 (OMEGA RATIO [19])

For a return R with distribution function F and hurdle τ ,

As a function of the hurdle, $\Omega(\cdot)$ is determined by—and determines—the entire distribution.

$$\Omega(\tau) = \int_{\tau}^{\infty} (1 - F(r))dr / \int_{-\infty}^{\tau} F(r)dr = \mathbb{E}[\max(R - \tau, 0)] / \mathbb{E}[\max(\tau - R, 0)] \quad (4.1)$$

PROPOSITION 4.1 (ELEMENTARY PROPERTIES)

- (i) Ω requires no distributional assumption and is determined by, and determines, the entire distribution: as a function of τ , $\Omega(\cdot)$ characterizes F . (ii) Ω is strictly decreasing in τ wherever defined. (iii) $\Omega(\mathbb{E}[R]) = 1$: the omega ratio equals one exactly at the mean. (iv) All moments—including γ_1 and γ_2 of Chapter 2—influence $\Omega(\tau)$ through the tail integrals; in particular the two densities of Example 2.1, tied in mean and variance, are strictly separated by $\Omega(\tau)$ at any hurdle below the mean.

Proof of (iii). $\max(x, 0) - \max(-x, 0) = x$ pointwise, hence $\mathbb{E}[\max(R - \tau, 0)] - \mathbb{E}[\max(\tau - R, 0)] = \mathbb{E}[R] - \tau$. At $\tau = \mathbb{E}[R]$ the numerator and denominator of Definition 4.1 are equal.

Property (iv) is the promised resolution of Chapter 2's tie, and the ranking accords with the Scott–Horvath preferences: at a common sub-mean hurdle, the right-skewed member of Example 2.1 carries the higher omega. The omega ratio is thus a performance measure with the moment preferences built in [20, 21].

4.2 The hurdle net**DEFINITION 4.2 (HURDLE-NET DECOMPOSITION)**

Write

$$G(\tau) = \mathbb{E}[\max(R - \tau, 0)], L(\tau) = \mathbb{E}[\max(\tau - R, 0)], \Omega = G/L \quad (4.2)$$

The pair (G, L) carries strictly more information than the quotient: Ω fixes only the ray G/L , while the total crossing mass $G + L$ measures how much of the distribution is in play at the hurdle. Two exact identities organize everything that follows.

PROPOSITION 4.2 (NET-MEAN IDENTITY)


$G(\tau) - L(\tau) = \mathbb{E}[R] - \tau$ for every τ . The net of the hurdle net is the excess mean; the gross, $G + L$, is pure distributional dispersion about the hurdle.

PRINCIPLE 4.1 (TWO PRIMITIVES, INDEPENDENTLY LOAD-BEARING)

The hurdle net and the omega ratio are independent new-generation primitives, not one concept and its detail. The hurdle net (G, L) is the measurement primitive: it carries mass, net, and time-resolved movement, and it powers reporting, budgeting, and attribution without a ratio ever being formed—the loss mass alone carries the risk report of Chapter 7. The omega ratio $\Omega = G/L$ is the ranking primitive: a scale-free functional fit for comparison and optimization. Each stands without the other; the canonical triad of the next section draws on both.

4.3 Canonical coordinates**DEFINITION 4.3 (CANONICAL TRIAD)**

Define the opportunity mass, the direction, and the increment:



The canonical triad: opportunity mass, direction, increment.

$$M = \tanh(\log(1 + G + L)) \in [0, 1), D = (G - L)/(G + L) \in (-1, 1), \Delta\Omega_t = \Omega_t - \Omega_{t-1} \quad (4.3)$$

The increment $\Delta\Omega_t$ is the triad's time-bearing member: a reading of the hurdle net in motion, it places the triad on the time-inherent pole of the GM futures polarity (Principle 1.1).

THEOREM 4.1 (EXACT MIRROR SYMMETRY)

For every gain $k > 0$ the map $D_k(\Omega) = \tanh(k \log \Omega)$ satisfies $D_k(1/\Omega) = -D_k(\Omega)$: a configuration and its long-short mirror take values of equal magnitude and opposite sign, with neutrality $D_k(1) = 0$ exactly at $G = L$ (equivalently, by Proposition 4.2, at $\tau = \mathbb{E}[R]$). Moreover the hurdle-net direction is itself such a coordinate:

$$D = (G - L)/(G + L) = (\Omega - 1)/(\Omega + 1) = \tanh(1/2 \log \Omega), \quad (4.4)$$

so D and $\tanh(\log \Omega) = (\Omega^2 - 1)/(\Omega^2 + 1)$ are strictly increasing transformations of one another: same sign, same zero, same ranking, same mirror.

Proof. Oddness: $\log(1/\Omega) = -\log \Omega$ and $\tanh(-x) = -\tanh(x)$. The identity: divide $(G - L)/(G + L)$ through by L and substitute $\Omega = G/L$; then $(\Omega - 1)/(\Omega + 1) = (e^{x/2} - e^{-x/2})/(e^{x/2} + e^{-x/2})$ with $x = \log \Omega$. The identities are algebraic consequences of the hurdle-net definitions; numerical checks diagnose implementation error but are not evidence for the theorem.

Figure 4.1 is Theorem 4.1 made visible, and it repays slow reading. Take a configuration with $\Omega = 3$ and its exact long–short mirror with $\Omega = 1/3$. On the raw ratio scale these read 3.00 and 0.33—numbers so lopsided that any learner, ranker, or reader will weight the bullish case nearly ten times the bearish one. On the canonical curve the same two configurations land at exactly equal heights on opposite sides of zero: the dotted chord connects them, and it passes through the origin. Fold the panel along the vertical axis and each curve lands on its own reflection with only the sign flipped; neutrality sits at the origin because $D = 0$ exactly when $G = L$, which by Proposition 4.2 is exactly the hurdle at the mean. That is the visceral content of the theorem: long–short evenhandedness is not asserted, it is drawn. Panel (b) shows the mass coordinate; Figure 4.2 then applies the machinery to Example 2.1, where the two-moment tie is broken decisively ($\Omega = 6.47$ against $\Omega = 4.44$ at a common sub-mean hurdle).




Figure 4.1: Canonical hurdle-net coordinates. Panel (a): the direction maps $\tanh(\log \Omega)$ and $\tanh(1/2 \log \Omega) = (G - L)/(G + L)$ as odd functions of $\log \Omega$: Ω and $1/\Omega$ mirror exactly (Theorem 4.1). Panel (b): the opportunity mass M : zero when nothing crosses the hurdle, saturating as crossing mass grows.

4.4 The subsumption theorem

We can now say precisely what Chapter 3 built.




Figure 4.2: The omega ratio resolves the two-moment tie of Example 2.1: identical mean and variance, decisively different Ω at a common hurdle, with the ranking matching the Scott–Horvath moment preferences of Chapter 2.

THEOREM 4.2 (KINEMATIC SUBSUMPTION)

The kinematic triad is a heuristic estimator of the canonical triad, member by member:

Kinematic channel	Canonical quantity	Shared properties
speed S_t	mass M	bounded $[0, 1)$, undirected, monotone in activity
velocity V_t	direction D	signed, bounded, zero at neutrality, mirror-symmetric
acceleration A_t	increment $\Delta\Omega$	change of the signed channel, regime-normalized

Speed and mass are both bounded undirected activity readings, zero at zero activity and saturating under extremes; velocity and direction are both signed channels vanishing exactly at neutrality (zero supported drift; $G = L$) with symmetric treatment of the long and short

sides (heuristically for V_t , exactly for D by Theorem 4.1); acceleration and $\Delta\Omega$ are the respective change readings. The kinematic triad estimates these quantities from price–volume kinematics; the canonical triad computes them from the distribution itself.

Corollary 4.1 (What is gained by the canonical form). Replacing the heuristic channels by their canonical counterparts confers: coordinate-independence (the canonical triad depends on F and τ only, not on the choice of volatility readings, windows, or smoothings); exact long–short symmetry (Theorem 4.1, available only approximately to V_t); and auditability (every canonical value is a closed-form transformation of (G, L) and can be re-derived and parity-checked).

The heuristics are not discarded—they remain fast, robust estimators, and their empirical record stands. The claim of this chapter is taxonomic: the kinematics are omega objects wearing kinematic clothes; with the object named, it can be optimized directly, across every mechanism at once. That program is Chapter 5.

4.5 Problems

4.1 Prove Proposition 4.1(ii): $\Omega(\tau)$ is strictly decreasing wherever $0 < F(\tau) < 1$.

4.2 Prove Proposition 4.2 from $\max(x, 0) - \max(-x, 0) = x$, and deduce that $D = (\mathbb{E}[R] - \tau)/(G + L)$: the canonical direction is the excess mean normalized by crossing mass.

4.3 Show that M of Definition 4.3 is invariant under $G \leftrightarrow L$ and strictly increasing in $G + L$, and compute its value when the entire unit of distributional mass sits exactly at the hurdle.

4.4 Two configurations satisfy $\Omega = 3$ and $\Omega = 1/3$. Compute $\tanh(\log \Omega)$ and D for each, verify the mirror, and explain why the raw ratio values 3.0 and 0.33 would bias a symmetric learner while the canonical values do not.

4.5 Exhibit two hurdle nets with the same D but different M , and describe the economic difference between them.

Formal Omega Function Theory

4.6 Nguyen's setup and the reciprocal convention

Vu Ngoc Nguyen's *Omega Function: A Theoretical Introduction* [33] supplies a measure-theoretic treatment that separates what Omega requires from what continuous-density derivations merely assume.

Let F be a cumulative distribution function with finite mean μ and non-trivial domain (A, B) , the interior of $x : 0 < F(x) < 1$. Nguyen defines

$$I_1(x) = \int_A^x F(t)dt = \mathbb{E}[(x - X)_+], \quad I_2(x) = \int_x^B [1 - F(t)]dt = \mathbb{E}[(X - x)_+].$$

His orientation is $\Omega_N(x) = I_1(x)/I_2(x)$, a loss-over-gain ratio that increases with x . This textbook uses the common finance orientation $\Omega(x) = I_2(x)/I_1(x)$, a gain-over-loss ratio that decreases with x .

Convention warning. $\Omega = 1/\Omega_N$. A monotonicity statement, endpoint limit, or derivative imported from one orientation into the other must be inverted. The underlying gain and loss integrals are identical.

4.7 Properties without continuity assumptions

Nguyen proves, for a finite-mean CDF with non-empty non-trivial domain:

1. I_1 and I_2 are Lipschitz continuous.
2. Ω_N is continuous and strictly increasing on (A, B) ; equivalently, this text's Ω is continuous and strictly decreasing.
3. Ω_N runs from 0 to ∞ across the domain; the reciprocal runs from ∞ to 0.
4. $I_2(x) - I_1(x) = \mu - x$, hence either ratio equals one at $x = \mu$.

5. One-sided derivatives exist even when F has atoms; derivative jumps record point mass.

Nguyen also derives a reconstruction formula for F from Ω_N , its one-sided derivative, and μ , and proves that equal full Omega functions imply equal CDFs under the stated assumptions. A single Omega value does not characterize F ; the entire curve does.

4.8 Standard dispersion, mixtures, and finite samples

Nguyen defines the standard dispersion through the slope of the increasing orientation:

$$\omega = 1/\Omega'_N(\mu+) = I_1(\mu) = I_2(\mu) = (1/2)E|X - \mu|.$$

Thus ω is one half of mean absolute deviation about the mean and requires only a finite first absolute moment. For a mixture $F = \sum_i a_i F_i$, the mean is $\sum_i a_i \mu_i$, while the dispersion includes both weighted component dispersions and a correction for component means away from the mixture mean.

For non-constant observations $x_1, \dots, x_n$, empirical standard deviation σ , and $\omega = (2n)^{-1} \sum_i |x_i - x|$, Nguyen proves

$$2 \leq \sigma/\omega \leq \sqrt{(2n)}.$$

The upper bound grows with sample size: neither statistic immunizes the analyst against sparse extremes.

4.9 Interpretation and limits

The formal results justify retaining gain and loss masses before forming a ratio, declaring the hurdle, reporting the curve for due diligence, and attaching uncertainty to empirical tail integrals. They do not encode serial dependence, costs, liquidity, or path-dependent execution unless the return variable is defined to include them. Nguyen's empirical examples illustrate properties; they do not prove universal portfolio superiority.

4.10 Problems

1. Convert the endpoint and monotonicity properties of Ω_N into this text's reciprocal convention.
1. Use $I_2 - I_1 = \mu - x$ to prove that either convention equals one at μ .

1. For a two-point distribution, compute I_1 , I_2 , both Omega orientations, and ω .
1. Explain why equality of one Omega value does not imply equal distributions, while equality of full curves does under Nguyen's assumptions.
1. Verify $2 \leq \sigma/\omega$ for a symmetric two-point sample and identify the equality case.

MUTANTDeFi

5

Pricing and Optimizing Horizon-Bound Future States

IN THIS CHAPTER

One published triad feeds gates, learners, and rewards. A change in the measurement moves every mechanism together.

5.1 One objective, many mechanisms

Chapter 4 established the hurdle net and its canonical triad of mass, direction, and increment; this chapter installs the omega objective, computed from that triad, as the single objective across every mechanism in which the Chapter 3 heuristics operate. The unifying principle is that a threshold rule, a supervised learner, and a reinforcement-learning policy are three search procedures over the same decision surface, and the canonical triad gives them one shared coordinate system.

PRINCIPLE 5.1 (ONE PRIMITIVE, THREE MECHANISMS)

Publish the canonical triad $(M, D, \Delta\Omega)$ once, and let every decision mechanism consume it: rule-based gates threshold it; supervised learners regress on and toward it; reinforcement learners observe it as state and are paid in it as reward. A change in the underlying hurdle-net measurement then propagates identically to all three, making their behaviors directly comparable and jointly auditable.

5.1.1 Rule-based gating

The kinematic alignment gate of Definition 3.4 translates verbatim: enter long when $D \geq +\beta$ (with adequate mass, $M \geq \gamma$), short when $D \leq -\beta$, flat otherwise. By Theorem 4.1 the long and short rules are exactly mirror-symmetric—a property the heuristic gate could only approximate. A sobering empirical note follows: in pre-registered comparisons, this minimal rule has proven remarkably hard to beat out of sample by any richer overlay (Appendix A), a finding that recommends thresholding the canonical direction directly and treating every additional layer as a hypothesis to be tested, not an improvement to be assumed.

5.1.2 Supervised learning

A regression learner targeting the raw ratio Ω inherits an unbounded, right-skewed label on which symmetric loss functions misbehave. The canonical direction repairs this by construction: the label $\tanh(\log\Omega)$ (equivalently D , by Theorem 4.1) is bounded in $(-1, 1)$, treats a bullish configuration and its bearish mirror equidistantly from neutrality, and is log-linear near $\Omega = 1$ —small hurdle-net imbalances enter the label linearly in log-odds fashion while extremes saturate, matching their diminishing decision relevance. Kernel methods suited to non-Gaussian structure [28] then learn the map from state to canonical direction without distributional assumptions. It is a point of historical interest that this label was adopted empirically—as the transformation that made the learning problem well-behaved—before the hurdle-net formalism identified it as the exact canonical coordinate.

The figure is a schematic, not a numbered definition. Chapter 6 supplies the numbered dominance relation, the efficient set, and the named frontier that the schematic already uses.




Figure 5.1: Mean-omega dominance on a tail-aware risk axis (schematic). Panel (a): because variance underprices the left tail, the mean-variance selection A carries risk its own coordinate cannot see; the omega-efficient A^* attains higher expected return at equal tail-aware risk. The mean-omega frontier is the true envelope on that axis; the mean-variance choice, re-scored, sits inside it. Panel (b): the mechanism is the fat left tail of A against the compressed tail of A^* —the property the omega ratio prices and rewards.

5.1.3 Reinforcement learning

Policy learning closes the loop from prediction to action [29]. The omega objective enters twice. As reward asymmetry: paying $+|r|$ for returns above the hurdle and $-\lambda|r|$, $\lambda > 1$, for returns below it makes expected reward maximization an omega maximization—the loss aversion is the denominator of Definition 4.1 made operational. As alignment shaping: scaling realized reward up when the action agrees with the predicted canonical direction, and scaling penalty up when it opposes it, with $\Delta\Omega$ —the canonical increment—pricing the change in hurdle-net evidence over the holding interval, exactly as acceleration priced the change in velocity. Holding-time penalties and prompt-exit bonuses reproduce, declaratively, the loss-duration bounds that the compression overlay of Chapter 3 enforces imperatively: reward design and risk overlay are two implementations of one thesis.

5.2 Portfolio optimization: omega risk budgeting

At the portfolio level the omega objective replaces the variance objective of Chapter 1 wholesale: contributions to portfolio Ω replace contributions to portfolio variance, and the mean-variance frontier gives way to a mean-omega frontier. The comparison must be made

on a risk axis that prices the full distribution; on such an axis the omega-efficient frontier is the true envelope and the mean–variance choice is revealed as dominated, carrying left-tail risk that variance never measured (Figure 5.1). The practical portfolio grammar of Chapter 3—cross-sectional ranking, per-position bounds, portfolio slope protection—survives intact, with ranking now performed in the canonical direction D and budget expressed in omega contribution, and with compression continuing to guarantee bounded loss in regimes where moments fail entirely.

Volume II names the envelope, the dominance relation that generates it, and the selection rule on it. What this section asserts is the replacement of the objective: variance is not the quantity being equalized, and (σ, μ) is not the plane on which efficiency is read.

5.3 Trading activity as information

The final domain is the trading activity itself. A strategy’s timing—when and how often it acts—is a channel that may or may not carry information, and the canonical mass gives the natural reference against which to measure it. These are time-resolved statistics by construction—computed window by window over a distribution in motion—and they are therefore the chapter’s most distinctly activity-theoretic instruments (Principle 1.1).

DEFINITION 5.1 (ACTIVITY–OPPORTUNITY DIAGNOSTICS)

Fix a minimal canonical reference rule (threshold on D with mass floor on M). For strategy s over evaluation windows $w_1, \dots, w_K$:

$$\rho_s(w) = n_s(w)/b(w), \text{AO}_s = \text{Spearman}(n_s(w_k)_k, M(w_k)_k), \quad (5.1)$$

where $n_s(w)$ is the strategy’s trade count on window w , $b(w)$ the reference rule’s band-crossing count, and $M(w)$ the mean opportunity mass.

PRINCIPLE 5.2 (FREQUENCY AS INFORMATION)

$\rho_s \approx 1$ says the strategy times like the minimal reference and its timing channel adds nothing; $\rho_s \gg 1$ flags manufactured entries the hurdle net does not motivate (timing noise and fee drag); $\rho_s \ll 1$ flags selectivity, whose information content the second statistic then adjudicates: a significantly positive AO_s is evidence that the strategy trades more when there is more to trade, while $AO_s \leq 0$ marks a timing channel decoupled from—or opposed to—the opportunity structure. Jointly, (ρ_s, AO_s) separate how much a strategy trades from whether that quantity is informed.

5.4 Evaluation discipline

Cross-mechanism optimization demands cross-mechanism evaluation, and the discipline is worth stating as doctrine. Decouple signal from execution: a feed or state reading is judged by out-of-sample rank correlation with forward returns (information coefficients), calibration, and coverage; a trader is judged downstream and separately, so that a bad execution overlay cannot impeach a good signal or vice versa. Hold out and pre-register: every candidate overlay is tested against the minimal canonical control on identical windows and cost assumptions, with hypotheses registered before inspection; the null—that the minimal rule stands—is a live outcome and has, in the comparisons of Appendix A, frequently been the verdict. State the boundary: evidence of measured information content is not evidence of deployable profitability, and public claims are confined to what the measurements support. This last discipline is what permits a research program to be candid about open problems—including live-versus-backtest reconciliation—without impairing the integrity of the claims it does make.

An information coefficient is not the optimality criterion of Chapter 6.

5.5 Problems

5.1 Show that the minimal canonical gate (threshold on D , floor on M) is invariant under simultaneous long–short reflection of the return distribution, and identify the property of Theorem 4.1 responsible.

5.2 A learner is trained once on label $\tanh(\log \Omega)$ and once on raw Ω with a symmetric loss. Explain, using the mirror and the right-skew of the raw ratio, the directional bias expected of the second learner.

5.3 Show that maximizing $E[\text{reward}]$ under the asymmetric scheme ($+|r|$ above the hurdle, $-\lambda|r|$ below, $\lambda > 1$) orders policies identically to an omega-type gain-loss criterion with reweighted denominator, and interpret λ .

5.4 Design a window partition and null hypothesis for AO_s such that the statistic has a known permutation distribution, and state the rejection rule at the five-percent level.

5.5 Explain why (ρ_s, AO_s) must be computed under identical cost assumptions and windows as the reference rule, and what artifact arises if they are not.

MUTANTDeFi

Volume II

Efficiency and the Frontier

MUTANTIDEFI

IN THIS VOLUME


A risk coordinate, a dominance relation, an efficient set, a named frontier, and a selection rule on that set.

MUTANTDeFi

To the Reader of Volume II

IN THIS CHAPTER

Volume I measured. Volume II numbers the efficiency objects those measurements already used.



Mean–variance states a closed two-moment world. The volume replaces it with a hurdle-net reading and the GM! Frontier.

Volume I read the return as a hurdle net (G , L) and a ratio Ω (Definitions 4.1–4.3). It refused variance as the risk report and let the loss mass carry that job. Figure 5.1 drew an omega-efficient envelope against a mean–variance choice on a tail-aware axis. Section 5.2

and Table 6.1 already speak of a mean–omega frontier. Theorem 6.1 already characterizes an interior maximum of Ω by a security line.

Those drawings and tables are a theory of measurement and a sketch of a menu. They are not yet a theory of efficiency. What remains is to number the objects the sketch already uses: a risk coordinate, a dominance relation, an efficient set, a named frontier, and a selection rule on that set.

Volume II does that work. Chapter 6 numbers the objects and proves the algebraic bridge that makes omega a ranking primitive on the same plane whose risk axis is the loss mass. Chapter 7 states the like-for-like replacements those objects complete.


A reader who has accepted Figure 5.1 and Table 6.1 has already accepted the informal versions of Definitions S.1–S.4. The numbered statements make the informal reading checkable, and they separate the menu (the undominated set) from the selection (the max- Ω book on that set).

6

Gleim–Markowitz Efficiency and the GM! Frontier

IN THIS CHAPTER

Tail-aware risk, mean–omega dominance, and the GM! Frontier replace the Markowitz menu.



Mean–omega dominance on a tail-aware risk axis.

6.1 Setup

Fix a declared hurdle τ (Principle 6.2: no omega figure without its hurdle) and a feasible set W of books. For $w \in W$ write $R(w)$ for the book return and

$$G(w; \tau) = \mathbb{E}[\max(R(w) - \tau, 0)], L(w; \tau) = \mathbb{E}[\max(\tau - R(w), 0)], \Omega(w; \tau) = G/L \text{ when } L > 0,$$

exactly Definition 4.2. Let $\mu(w) = \mathbb{E}[R(w)]$. Proposition 4.2 (net-mean identity) is load-bearing:

$$G - L = \mu - \tau.$$

The two primitives of Principle 4.1 stay independent: (G, L) measures; Ω ranks.

W is whatever the decision problem declares: long-only simplex, long-short with leverage bounds, a lattice of conviction-weighted books, a set of compressed payoffs. The definitions below are relative to (W, τ) . Changing either changes the efficient set. That is not a defect; it is the declared-hurdle convention applied to allocation.

6.2 The ranking identity

PROPOSITION S.1 (RANKING IDENTITY)

From the net-mean identity,

$$\mu = \tau + G - L = \tau + L(\Omega - 1) \quad (L > 0).$$

Hence:

1. At fixed loss mass L , ranking books by μ is ranking them by G , and is ranking them by Ω .
2. At fixed Ω , μ is affine in L . Rays of constant omega are straight lines in the (L, G) plane; their slope is Ω itself.
3. Mean-variance's (σ, μ) plane is therefore replaced, like-for-like, by the tail-aware plane $(L(\tau), \mu)$. The omega ratio is the ranking primitive on that plane. That is why Figure 5.1 can call the omega-efficient envelope the true envelope once the risk axis prices the left tail.

Proof. Rearrange $G - L = \mu - \tau$ to $\mu = \tau + G - L$. Substitute $G = \Omega L$ on $L > 0$ to obtain $\mu = \tau + L(\Omega - 1)$. At fixed L , μ is strictly increasing in Ω (and in G). At fixed Ω , $\mu - \tau$ is linear

in L with slope $\Omega - 1$. In the (L, G) plane the locus $G = \Omega L$ is a ray from the origin of slope Ω . The map $(L, G) \mapsto (L, \tau + G - L)$ sends that ray to a line of slope $\Omega - 1$ in the (L, μ) plane. Variance does not appear. The replacement of (σ, μ) by $(L(\tau), \mu)$ is therefore like-for-like: same jobs (risk coordinate, reward coordinate, ranking of undominated books), different risk axis.

This is the closed gap. Whole-distribution omega is not a rival league-table statistic sitting beside Markowitz. It is the ranking of the full-distribution successor plane.

The identity also explains a standing confusion. Analysts sometimes treat Ω as a third coordinate, as if one could plot (σ, μ, Ω) and keep variance as risk. Proposition S.1 forbids that reading. Once the risk axis is L , Ω is already the ranking of that plane at the declared hurdle. Adding variance back in is not a refinement; it is a change of problem.

A second reading is computational. Because $\mu = \tau + G - L$, every statement about Pareto improvement in (G, L) at fixed τ is a statement about mean-omega improvement in (L, μ) once the risk axis is L . The dual picture of Section 6.4 uses that translation.

6.3 Numbered definitions

DEFINITION S.1 (TAIL-AWARE RISK)

Tail-aware risk at a declared hurdle τ is the loss mass $\rho_{-\tau}(w) := L(w; \tau) = E[\max(\tau - R(w), 0)]$. Any strictly increasing transformation of L that still prices the left tail is an equivalent risk axis. No omega figure is published without its hurdle.

The first sentence names the coordinate. The second permits strictly increasing reparameterizations— $\log(1 + L)$, a scaled put, a monotone transform used for display—so long as the left tail remains the priced object. Variance is not such a transform. The third sentence is Principle 6.2 restated at the point of definition: a loss mass without a hurdle is an incomplete report, and an omega figure without a hurdle is undefined as a ranking.

Expected shortfall at a stated level is a sibling, not a substitute. Rockafellar–Uryasev CVaR at level α prices a quantile slice [15]. Loss mass prices everything below a declared hurdle. When τ is chosen to be a VaR level the two objects communicate; they are not identical. A CVaR frontier in $(CVaR_\alpha, \mu)$ is another tail-aware menu. It is not the GM! Frontier unless the hurdle and the quantile are locked together and the ranking primitive is still Ω .

DEFINITION S.2 (MEAN-OMEGA DOMINANCE)

Book w mean-omega dominates w' at τ when $\rho_{-\tau}(w) \leq \rho_{-\tau}(w')$ and $\mu(w) \geq \mu(w')$, with at least one inequality strict. Variance is not an admissible risk coordinate for this relation.

The relation is the Markowitz dominance relation with σ replaced by ρ_{τ} . It is a partial order on W at the declared hurdle. Crossing omega curves (Figure 7.1, Chapter 7) mean the order can reverse when τ moves; that is information, not a defect, and it is why Principle 6.2 forbids hurdle shopping. The last sentence is load-bearing. A book that is mean-variance efficient can be mean-omega dominated, which is exactly the configuration Figure 5.1 draws: A sits on the (σ, μ) frontier and inside the tail-aware envelope.

Dominance is not optimality. A book can be undominated without maximizing Ω . The next two definitions separate the menu from the selection.

DEFINITION S.3 (OMEGA-EFFICIENCY / GLEIM-MARKOWITZ EFFICIENCY)

A feasible book w^* is omega-efficient (Gleim-Markowitz efficient) at τ on W when no w in W mean-omega dominates it. This is the full-distribution replacement of Markowitz mean-variance efficiency.

The parenthetical is the name used in this volume for the same object. “Gleim-Markowitz” records the lineage (endnote). It does not claim Harry Markowitz’s co-authorship, review, or endorsement. Efficiency is a set property: the collection of undominated books at (W, τ) . Every undominated book is on the menu. The trustee rule of Table 6.1—“hold the omega-efficient mix at the declared τ ladder”—is S.3 applied rung by rung. It does not by itself pick a single book.

DEFINITION S.4 (GM! FRONTIER)

The GM! Frontier (mean–omega frontier, omega-efficient frontier) at (W, τ) is the set of omega-efficient books, or its image in (ρ_τ, μ) . It is the Table 6.1 successor to the (σ, μ) allocation menu. It is not an index, feed, or product.

The GM! Frontier is the image of the omega-efficient set in the tail-aware plane.

Three names, one object. “GM! Frontier” is the name used in this volume; “mean–omega frontier” and “omega-efficient frontier” are the descriptive names already used in Volume I. The object is either the efficient set in W or its plot in the tail-aware plane. Table 6.1 listed that successor without numbering it. This definition is that numbering.

The last sentence is a plane separation, not a slogan. The frontier is an allocation envelope over books. An index publishes constituents. A feed publishes readings. A product delivers a service. None of those objects is the set of undominated books at (W, τ) .

DEFINITION S.5 (GLEIM–MARKOWITZ OPTIMALITY)

A feasible book w^* is Gleim–Markowitz optimal at (W, τ) when it is omega-efficient and maximises $\Omega(w; \tau) = G/L$ on $\{w \text{ in } W : L(w; \tau) > 0\}$. An interior optimum satisfies the omega security line $g_i - \Omega_p \ell_i = \kappa$ (Theorem 6.1).

S.3 is the menu (every undominated book). S.5 is the selection on that menu—the full-distribution analogue of a max-Sharpe / tangency choice, now characterized by the omega security line rather than by a covariance beta. Table 6.1’s trustee rule is S.3 applied rung by rung; it does not by itself pick the global S.5 book.

The restriction $L > 0$ is the domain of Ω . A book with $L = 0$ and $G > 0$ has no finite omega; it is a free lunch relative to the hurdle and is either infeasible under a serious W or a signal that the hurdle has been set below the support of every feasible return. Volume I already treated Ω as defined only for $L > 0$. S.5 inherits that domain.

Theorem 6.1 is proved in Chapter 7, where it sits in the CAPM-replacement argument that first produced it. The statement is repeated here because S.5 names the book to which the security line applies. The hypotheses travel with the theorem: interior maximum, continuously distributed returns, budget constraint, no atoms. Compression (Definition 3.6) creates an atom at the loss bound; a subgradient version covering atoms and binding constraints is future work (Appendix C, item R7).

6.4 Menu and selection

The distinction between S.3 and S.5 is the distinction Markowitz already required between the efficient frontier and the tangency portfolio. The frontier is the menu a trustee can hold under a declared preference or a declared hurdle ladder. The tangency book is a selection on that menu once a riskless rate, or here a ranking primitive, is named.

Work the distinction on a three-book menu at $\tau = 0$.

Book	L	G	$\mu = G - L$	$\Omega = G/L$
A	1.0	1.2	0.20	1.20
B	0.6	0.84	0.24	1.40
C	1.2	1.2	0.00	1.00

B has lower loss mass and higher mean than A and than C, so B mean- Ω dominates both (S.2). The Ω -efficient set is {B} (S.3). Gleim-Markowitz optimality (S.5) selects B. Variance is not reported. If A had a smaller variance than B, a mean-variance menu could prefer A. That preference is exactly the error Figure 5.1 was drawn to show.

A second menu, with two undominated books, makes the trustee rule visible.

Book	L	G	μ	Ω
P	0.4	0.48	0.08	1.20
Q	0.8	1.20	0.40	1.50
R	0.9	0.90	0.00	1.00

R is dominated by both P and Q. Neither P nor Q dominates the other: P has less loss mass, Q has higher mean. The GM! Frontier is {P, Q} (or the two points (0.4, 0.08) and (0.8, 0.40) in (ρ_τ, μ)). S.5 selects Q, the max- Ω book on that set. A trustee who has declared a

conservative rung and will not accept L above 0.5 holds P : that is $S.3$ at a constrained W , not a contradiction of $S.5$ on the unconstrained set. The two books answer different questions. Confusing them is how a hurdle ladder becomes a hidden leverage choice.

The ranking identity reads both tables. At fixed L , higher Ω is higher μ . At fixed Ω , higher L is higher $|\mu - \tau|$ along the ray. Q is steeper than P in (L, G) ; that is why it is the $S.5$ book. P is nearer the origin on a shallower ray; that is why a loss-mass constraint can select it without leaving the frontier.

6.5 Dual picture: the (G, L) Omega frontier

Kapsos, Zymler, Christofides, and Rustem write Omega maximization as a linear program in gain–loss coordinates [34]. Related treatments use the Charnes–Cooper transformation of a linear fractional program [35]. This section cites that literature for computation. It does not claim the discovery that Omega can be optimized. This volume’s contribution is the two-primitive split (net versus ratio), the canonical triad and the mirror (Theorem 4.1), the declared-hurdle convention, the security-line characterization (Theorem 6.1), and the time-inherent polarity (Principle 1.1).

Let returns be discrete on scenarios $s = 1, \dots, S$ with probabilities p_s , and let book returns be linear in weights, $R(w)_s = w \cdot r_s$. Then

$$G(w; \tau) = \sum_s p_s \max(w \cdot r_s - \tau, 0), L(w; \tau) = \sum_s p_s \max(\tau - w \cdot r_s, 0).$$

Maximizing $\Omega = G/L$ on $L > 0$ is a linear fractional program. The Charnes–Cooper scaling $t = 1/L$, $z = tw$ produces an equivalent linear program: maximize $tG(z/t; \tau)$ subject to $tL(z/t; \tau) = 1$, $t > 0$, and the image of W under the same scaling. In the (L, G) plane the feasible image is a polytope when W is a polytope. No attainable point lies above the upper-left Pareto face of that polytope. A ray from the origin is a constant- Ω locus. The steepest feasible ray is max Ω . Kapsos et al. call that Pareto set the Omega frontier.

Proposition S.1 maps that set onto the GM! Frontier. Because $\mu = \tau + G - L$, Pareto improvement in (G, L) at fixed τ is mean–omega dominance in (L, μ) once the risk axis is L . The upper-left face in (L, G) is the northwest face in (ρ_τ, μ) . The steepest ray is the $S.5$ book. Interior points of the face that are not on that ray are $S.3$ books that are not $S.5$. The computational Omega frontier and the GM! Frontier are the same set under two coordinate charts.

Cite the LP literature for algorithms. Cite this chapter for the named efficiency, frontier, and optimality objects.

The linear program is a picture, not a product. It does not compute an index. Hurdle-net budgeting by marginal contributions is a first-order heuristic; it is not a solution of the Kapsos program over W , and it does not certify S.5.

6.6 Relation to Theorem 6.1 and Table 6.1

Theorem 6.1 states that at an interior maximum of $\Omega_p = G/L$ over budget-constrained weights, with continuously distributed returns, there is a constant κ such that

$$g_i - \Omega_p \ell_i = \kappa \text{ for every asset } i,$$

where $g_i = \partial G / \partial w_i = \mathbb{E}[R_i 1_{R_p > \tau}]$ and $\ell_i = \partial L / \partial w_i = -\mathbb{E}[R_i 1_{R_p < \tau}]$. The book that satisfies the line, if it is omega-efficient, is Gleim–Markowitz optimal. The line is the first-order face of S.5, not a second definition of S.3. A book can lie on the GM! Frontier without equalizing $g_i - \Omega_p \ell_i$ across names; it is then an efficient mix at a binding constraint or at a declared loss-mass budget, analogous to a mean–variance book that is efficient but not tangency.

Table 6.1’s replacement map listed “mean–omega frontier on a tail-aware risk axis” as the successor of the (σ, μ) allocation menu, and “hold the omega-efficient mix at the declared ladder” as the successor of the policy portfolio. Those two rows are now S.4 and S.3. The row that allocates by marginal (g_i, ℓ_i) is hurdle-net budgeting: a first-order heuristic that uses the ingredients of Theorem 6.1 without solving S.5. That heuristic is not the envelope.

6.7 What this chapter defines

The object numbered here is an allocation envelope over books. An index of published constituents is a different object. An information coefficient is not the optimality criterion. Numbering S.1–S.5 does not create a service, and it does not claim Harry Markowitz’s co-authorship, review, or endorsement.

6.8 What Volume I already supplied

Classical	Volume I	Numbered in this chapter
Return distribution $\rightarrow (\mu, \sigma)$	Whole distribution $\rightarrow$ hurdle net (G, L) and Ω (Def 4.1–4.3)	—
Variance as risk	Tail-aware / downside axis; loss mass L as the risk report (Ch 7)	S.1, a numbered risk coordinate
Mean–variance dominance	Fig 5.1 schematic: A^* omega-efficient vs A mean–variance	S.2, a numbered dominance relation
Efficient frontier (σ, μ)	§5.2 and Table 6.1: “mean–omega frontier on a tail-aware risk axis”	S.3–S.4, a numbered frontier
Investor / tangency choice	Theorem 6.1 omega security line; Table 6.1 “omega-efficient mix”	S.5, a numbered optimality definition

Volume I’s mathematics is the measurement. Volume II’s contribution is the numbering, the ranking identity, and the explicit separation of menu from selection. A reader who already accepted Figure 5.1 and Table 6.1 has accepted S.1–S.4 informally. S.5 is the one object Volume I characterized (Theorem 6.1) without naming.

6.9 Problems

1. Prove Proposition S.1 from Proposition 4.2, and exhibit the image of a constant- Ω ray in the (ρ_τ, μ) plane.
1. Using the first three-book table, verify S.2–S.5 by hand. Then add a fourth book D with $L = 0.6, G = 0.72$. State whether the efficient set changes and whether the S.5 book changes.
1. Show that if F second-order stochastically dominates G at equal mean, then $L_F(\tau) \leq L_G(\tau)$

for every τ . Conclude that SSD at equal mean implies that G cannot mean- ω dominate F . Show the converse is false by a single-hurdle counterexample.

1. Write the Charnes–Cooper form of $\max G/L$ for discrete scenarios and identify the steepest-ray book with S.5.
1. Explain, in the language of S.3 versus S.5, why a trustee mix at a declared τ ladder can be ω -efficient at each rung and still fail to be Gleim–Markowitz optimal on the union of those feasible sets.

MUTANTDeFi

7

Replacement Proposals: Omega Successors for the Classical Infrastructure

IN THIS CHAPTER

Replace the job, not the brand of the statistic. Dual-report, then cut over instrument by instrument.

Job	Classical statistic	Successor
Performance	Sharpe / Sortino	$D(\tau^+)$ and $\Omega(\cdot)$
Risk report	VaR quantile	Loss mass $L(\tau)$ + bound
Systematic risk	$\beta = \text{Cov}/\text{Var}$	Omega beta ℓ_i

Like-for-like successors: the job stays, the statistic changes.

This chapter carries Volume II to its institutional conclusion. Chapter 1 presented the classical composition; this chapter proposes, instrument by instrument, its omega successor. Each proposal is like-for-like: it preserves the job the classical instrument performs—ranking, reporting, pricing, budgeting—while replacing the statistic that performs it. The proposals are labeled MD-P1 through MD-P5 and are offered to the academic and practitioner communities for scrutiny.


Chapter 6 numbered the efficiency objects that Table 6.1 already used. This chapter states the replacement map those objects complete. Theorem 6.1 is proved here, where the security line is used as a replacement instrument.

7.1 The replacement principle

PRINCIPLE 6.1 (LIKE-FOR-LIKE REPLACEMENT)

A classical instrument is a statistic attached to a job and a communication surface. A replacement proposal must (i) name the job, (ii) supply a full-distribution statistic that performs it, (iii) preserve or improve the communication surface (one number where one number is expected, a curve where due diligence permits), and (iv) specify a migration path—a dual-reporting period during which classical and omega figures are published side by side with parity checks—before any cutover.

A replacement must name the job, supply the full-distribution statistic, keep the communication surface, and specify the migration path.



Like-for-like successors: the job stays, the statistic changes.

One governance convention underlies every proposal and is stated first because the omega ratio, unlike the variance, is a family indexed by the hurdle.

PRINCIPLE 6.2 (THE DECLARED-HURDLE CONVENTION)

No omega figure is published without its hurdle. Standard reporting declares a small hurdle ladder—zero, the riskless rate, and the stated investment target—and headline figures are quoted at the declared hurdle appropriate to the audience. Figure 7.1 shows why the convention is load-bearing: omega curves of two assets can cross, so a ranking is meaningful only at a declared hurdle, and the full curve is the due-diligence object.

7.2 Replacing Sharpe and Sortino: performance measurement

The Sharpe ratio $(\mu - r_f)/\sigma$ [24] is the second-moment orthodoxy's performance instrument, and it inherits the orthodoxy's blindness wholesale: the two densities of Example 2.1, tied in mean and variance, are tied in Sharpe by construction. Sortino's repair—replacing σ with the downside semi-deviation [22]—breaks that particular tie but remains a two-number summary: it reads the loss side through a single second-order moment and is silent about tail shape below the hurdle. Both are members of a family the omega ratio completes.




Figure 7.1: Why the declared-hurdle convention is load-bearing. Two assets with nearly identical means: A is steady and thin-tailed; B carries a higher upside with a fat left tail. Panel (a): their omega curves cross. A dominates at conservative hurdles (its loss mass is small), B at aggressive ones. A single-point ranking is meaningful only at a declared τ ; the curve is the due-diligence object. Panel (b): the same information in the bounded direction coordinate $D(\tau) = \tanh(1/2 \log \Omega(\tau))$, the recommended headline scale.

PROPOSITION 6.1 (OMEGA WITHIN THE KAPPA FAMILY)

Write the lower partial moment of order n at hurdle τ as $\text{LPM}_n(\tau) = \mathbb{E}[\max(\tau - R, 0)^n]$, and the Kaplan–Knowles measures $\kappa_n(\tau) = (\mu - \tau)/\text{LPM}_n(\tau)^{1/n}$ [23], of which Sortino is κ_2 . Then, by the net-mean identity of Proposition 4.2,

$$\Omega(\tau) = 1 + (\mu - \tau)/\text{LPM}_1(\tau) = 1 + \kappa_1(\tau): \quad (7.1)$$

the omega ratio is, up to the unit shift, the first-order member of the family [21]—the member whose denominator reads the entire loss-side mass rather than a single moment of it, and whose value characterizes the full distribution as τ varies (Proposition 4.1).

PRINCIPLE 6.3 (MD-P1: PERFORMANCE REPORTING)

Replace point Sharpe/Sortino reporting with: (i) a headline direction figure $D(\tau^*) = \tanh(1/2 \log \Omega(\tau^*))$ at the declared hurdle τ^* —bounded in $(-1, 1)$, mirror-symmetric between long and short (Theorem 4.1), and therefore fit for league tables; and (ii) the omega curve $\Omega(\cdot)$ over the declared ladder as the due-diligence exhibit. Rankings are by D at the declared hurdle; disagreements between hurdles are disclosed, not averaged away.

7.3 Replacing VaR: risk reporting and capital

VaR's job is to report exposure and anchor capital; its failures (Chapter 1) are quantile blindness and incoherence. The hurdle net supplies a reporter with neither defect, and the compression overlay of Chapter 3 supplies something no estimator can: a bound that is engineered rather than estimated.

PROPOSITION 6.2 (THE LOSS MASS NEVER PENALIZES DIVERSIFICATION)

The hurdle-net loss mass is subadditive with additive hurdles: for any positions X, Y and hurdles τ_X, τ_Y ,

$$\mathbb{E}[\max(\tau_X + \tau_Y - (X + Y), 0)] \leq \mathbb{E}[\max(\tau_X - X, 0)] + \mathbb{E}[\max(\tau_Y - Y, 0)], \quad (7.2)$$


since $\max(a+b, 0) \leq \max(a, 0) + \max(b, 0)$ pointwise. The combined book's reported risk never exceeds the sum of its parts—precisely the coherence property whose failure condemned the quantile [14].

PRINCIPLE 6.4 (MD-P2: THE TWO-LINE RISK REPORT)

Replace the VaR line with two lines. The distribution line reports the loss mass $L(\tau)$ at the declared ladder together with expected shortfall at a stated level [15]: coherent, tail-magnitude-aware, and comparable across books. The engineering line reports the enforced loss bound of the compression overlay where payoff surgery is in force and auditable: “maximum session loss is a design parameter, not an estimate.” Capital attaches to the engineering line where it is auditable and to the distribution line where it is not; the two-line format makes the difference between estimated and enforced risk visible to every reader—a quantifiable improvement over single-number best practice, since the gap between the distribution line and the engineering line is itself a reported figure, position by position, and its narrowing is an auditable objective.

Definition S.1 is the numbered risk coordinate that this two-line report already used. Loss mass is ρ_τ . Expected shortfall is the sibling on the distribution line. The engineering bound is not a risk measure; it is a designed atom.

7.4 Replacing CAPM: pricing and systematic exposure



The omega security line of Theorem 6.1.

The CAPM prices covariance because, in a mean–variance world, covariance is the only risk that survives aggregation. In an omega world the analogous question—what does an asset do to the portfolio’s hurdle net?—has an exact answer, and it yields both a first-order condition and a pricing line. Let $R_p = \sum_j w_j R_j$, with portfolio gain and loss masses G , L at hurdle τ , and define each asset’s marginal hurdle-net contributions

$$g_i = \partial G / \partial w_i = \mathbb{E}[R_i 1_{R_p > \tau}], \quad \ell_i = \partial L / \partial w_i = -\mathbb{E}[R_i 1_{R_p < \tau}]: \quad (7.3)$$

g_i is the asset’s average payment in the portfolio’s gain states; ℓ_i is positive when the asset deepens the portfolio’s loss mass—downside co-movement, the omega-native generalization of beta.

THEOREM 6.1 (OMEGA FIRST-ORDER CONDITION AND SECURITY LINE)

At an interior maximum of $\Omega_p = G/L$ over budget-constrained weights (returns continuously distributed), there is a constant κ such that

$$g_i - \Omega_p \ell_i = \kappa \text{ for every asset } i \quad (7.4)$$

Proof. With multiplier λ for $\sum_j w_j = 1$, stationarity of G/L gives $(g_i L - G \ell_i)/L^2 = \lambda$ for all i ; multiply through by L and set $\kappa = \lambda L$. Numerical optimization can check an implementation of this first-order condition, but cannot replace its hypotheses or proof.

Equation (6.4) is the omega security line: in the cross-section of an omega-efficient portfolio, gain-state contribution net of Ω_p -weighted loss-state contribution is equalized. An asset that deepens the loss tail (large ℓ_i) must pay commensurately more in gain states to be held—the coskewness premium of Harvey and Siddique [18] generalized from the third moment to the whole loss side, and the natural completion of the lower-partial-moment equilibria of Bawa–Lindenberg and Harlow–Rao [25, 26].

If the interior maximum is attained at a feasible book, that book is Gleim–Markowitz optimal (Definition S.5) and therefore omega-efficient (Definition S.3). The theorem does not conversely certify that every omega-efficient book satisfies the line: books on a binding loss-mass or position constraint can be S.3 without being S.5.

PRINCIPLE 6.5 (MD-P3: EXPOSURE AND PRICING)

Replace $\beta_i = \text{Cov}(R_i, R_M)/\text{Var}(R_M)$ with the pair (g_i, ℓ_i) measured against the declared market portfolio and hurdle, with ℓ_i —the omega beta—as the headline systematic-risk figure: it reads downside co-movement directly, exists under tail indices where the covariance does not, and prices through (6.4) rather than through a second moment. Equilibrium prior construction (in the sense of Black–Litterman [27]) proceeds from (6.4) exactly as classical priors proceed from the CAPM.

7.5 Replacing the load-bearing mean–variance infrastructure

The deepest dependency is not any single statistic but the presumptive infrastructure—mandates, budgets, benchmarks, attributions—built on the premise that (μ, σ) suffices. Table 6.1 states the like-for-like map; three entries deserve commentary.

Table 6.1. The replacement map (MD-P4). Each classical instrument, its job, and its omega successor. Every successor quotes its declared hurdle (Principle 6.2).

Classical instrument	Job	Omega successor
Efficient frontier (σ, μ)	allocation menu	mean–omega frontier on a tail-aware risk axis (Chapter 5); GM! Frontier, Definitions S.3–S.4 (Chapter 6)
Sharpe / Sortino ratio	ranking, league tables	headline direction $D(\tau^*)$ plus the omega curve (MD-P1)
VaR line	exposure report, capital	two-line report: loss mass $L(\tau)$ with expected shortfall; enforced bound (MD-P2); tail-aware risk ρ_τ (S.1)
CAPM β	systematic exposure, pricing	marginal contributions (g_i, l_i) ; omega security line (6.4) (MD-P3); Gleim–Markowitz optimality (S.5)
Covariance risk budgeting / risk parity	risk allocation	hurdle-net budgeting: allocate by marginal (g_i, l_i) ; portfolio slope protection as the drawdown governor
Tracking error / information ratio	benchmark-relative mandate	hurdle net of active returns at $\tau = 0$: report Ω and D of $R_p - R_b$
Performance attribution	explain results	attribute changes in G and L by position: gain-side and loss-side contributions separately
Policy portfolio (60/40 and kin)	governance default	hurdle-declared policy: trustees declare the τ ladder; hold the omega-efficient mix at the declared ladder (S.3, not automatically S.5)
Gaussian stress and scenario weights	tail preparedness	direct loss-mass stress: shift the hurdle, reprice $L(\tau)$; no

Risk budgeting. Classical risk parity equalizes covariance contributions—a quantity that is unstable exactly when it matters. Hurdle-net budgeting allocates by the marginal contributions of Theorem 6.1, which are conditional first moments and remain finite and estimable under any tail index above one; the portfolio slope-protection rule of Chapter 3 then governs realized drawdown without any covariance estimate at all. That budgeting row is a first-order heuristic. It is not the GM! Frontier and does not certify S.5.

Benchmark-relative mandates. The information ratio is the Sharpe ratio of active returns and fails with it. Its successor is immediate: form the hurdle net of $R_p - R_b$ at $\tau = 0$ and report its Ω and D —a bounded, mirror-symmetric statement of whether the manager’s active distribution earns its keep, with the full curve available to the mandate’s due diligence.

Attribution. Because G and L are expectations of position-additive hinge payoffs, attribution decomposes exactly: each position’s contribution to the period’s gain mass and loss mass is computed directly, and “where did the risk come from” receives a loss-side answer rather than a variance-share answer.

PRINCIPLE 6.6 (MD-P5: MIGRATION DOCTRINE)

Adoption proceeds by dual reporting, never by decree: classical and omega figures are published side by side over a declared parallel period; parity checks reconcile the two where they should agree (Gaussian books) and disclose divergences where they should not (fat-tailed books, asymmetric books); hurdles are declared in advance and never moved mid-period; and cutover occurs instrument by instrument only when the omega figure has carried the instrument’s job through a full reporting cycle. The evidence discipline of Chapter 5 applies to the infrastructure itself: a replacement is a hypothesis, and the null—that the classical figure suffices for a given book—is a live outcome where returns are demonstrably near-Gaussian over the declared period.

7.6 What the proposals do not claim

Consistent with the evidence boundary maintained throughout this volume, the proposals of this chapter are measurement and governance proposals. They assert that the omega family reports, ranks, prices, and budgets on the full distribution where the classical family reads

two moments; they do not assert that any portfolio, fund, or strategy governed by them is or will be profitable. Empirical support for the measurement claims, in limited exhibit form, appears in Appendix A; complete validation records are available through the channel of Appendix D.

7.7 The GM futures canon

The chapter closes by answering a compositional question directly: of the theorems, propositions, principles, and definitions of this volume, which set best represents GM futures theory—the study, under Principle 1.1, of time-bound, distribution-wide information movement, as distinguished from portfolio diversification and optimization theory? The selection criterion is deliberately strict: an item enters the canon only if it is an original, unique contribution of this program and load-bearing for the theory; results that situate the work in the literature, however useful, are excluded. Five items qualify, and they are collected in Table 6.2 as an author-named convention.

Table 6.2. The GM futures canon (GM-1 through GM-5).

Item	Source	Contribution
GM-1: The Polarity Principle	Principle 1.1	time absent versus time inherent: the axis that defines the post-Sharpe class and names the theory's subject
GM-2: The Two-Primitive Principle	Principle 4.1; Definitions 4.2 and 4.3	the hurdle net as measurement primitive and the omega ratio as ranking primitive, independent and jointly generating the canonical triad of mass, direction, and increment
GM-3: The Mirror Convention	Theorem 4.1	exact long-short evenhandedness as an adoptable coordinate convention: $D(1/\Omega) = -D(\Omega)$ and $(G - L)/(G + L) = \tanh(1/2 \log \Omega)$
GM-4: The Subsumption Theorem	Theorem 4.2	the kinematic triad identified, member by member, as a heuristic estimator of the canonical triad—the bridge from working practice to theory
GM-5: The Activity Instruments	Theorem 6.1; Definition 5.1 with Principle 5.2	the omega security line as the pricing face of the theory, and the time-resolved activity–opportunity diagnostics as its measurement face

Definitions S.1–S.5 are Volume II’s numbered optimality layer. They complete Table 6.1; they do not enter the canon as a sixth item, because the canon’s criterion is Volume I’s measurement and polarity contribution. The ranking identity (Proposition S.1) is a corollary of Proposition 4.2 and is claimed as a situating result for the successor plane, not as a sixth canonical object.

Exclusions are as instructive as inclusions. The net-mean identity (Proposition 4.2) and the subadditivity of the loss mass (Proposition 6.2) are elementary consequences of the hinge function and are claimed as situating results, not original mathematics; the placement of omega within the Kappa family (Proposition 6.1) belongs to the prior literature [21, 23]; and the declared-hurdle convention (Principle 6.2), indispensable as governance, is procedure rather than theory. GM-3 deserves the same candor: as mathematics the mirror identity is elementary; its contribution is the convention—the adoption of an exactly symmetric, bounded coordinate as the common reporting and learning scale—and the canon labels it accordingly. What remains is compact: one polarity, two primitives, one convention, one bridge theorem, and two working instruments. That set is GM futures theory in its present form.

7.8 Interpretive benchmark: a full-distribution instrument and the VIX

VIX is a model-free, options-implied expectation of equity-index variance over a stated horizon [31, 32]. A full-distribution activity instrument is a different design class: it measures gain and loss mass around declared hurdles, preserves direction, and can report readings by instrument and horizon. The comparison is pedagogical, not promotional.

Table 6.3. Two design classes with different objects and epistemic sources.

Axis	VIX	Full-distribution activity instrument
Object	risk-neutral expected variance	physical-measure gain and loss mass around declared hurdles
Direction	undirected	signed through gain/loss imbalance
Horizon	stated standard horizon	declared horizon set; like-for-like comparisons
Time	snapshot expectation	reading plus explicitly defined increments and windows
Epistemic source	traded option prices	estimated research measurement
Validation	settlement and market adoption	uncertainty, calibration, holdouts, and reproducibility
Governance	published methodology	declared hurdles, methods, revisions, and publication boundary

The epistemic difference is decisive. VIX is formed from traded prices under a risk-neutral measure and benefits from settlement and a mature derivatives ecosystem. A physical-measure research statistic must earn confidence through methodology, uncertainty estimates, holdout validation, and governance. Design richness does not erase that distinction.

7.9 Problems

6.1 Prove Proposition 6.1 from the net-mean identity, and show that $\kappa_1(\tau)$, $D(\tau)$, and $\Omega(\tau)$ induce the same ranking at a common declared hurdle.

6.2 Construct distributions tied in Sharpe and Sortino that an Omega curve separates.

6.3 Prove loss-mass subadditivity and contrast it with the digital-option VaR example.

6.4 Derive Theorem 6.1 and state the assumptions required for its Gaussian relation to beta.

6.5 Explain why ranking can reverse across a hurdle ladder and draft a disclosure.

6.6 Classify each axis of Table 6.3 as favoring one design, the other, or neither; defend the epistemic-source classification.

Conclusion: Durable Ideas

The classical system uses mean and variance to allocate, price, and report risk. It is coherent under its assumptions, but it cannot distinguish distributions whose asymmetry, tail mass, or path structure differ outside the first two moments.

GM futures theory responds in stages. Higher moments expose the missing information. Market structure identifies when that information is economically load-bearing. Kinematic readings provide a practical, time-inherent heuristic. The hurdle net names the gain and loss masses being estimated; the Omega curve organizes them across economic hurdles; and Nguyen's formal results establish when that curve characterizes the underlying distribution. Volume II then numbers the successor plane: tail-aware risk, mean-omega dominance, omega-efficiency, the GM! Frontier, and Gleim-Markowitz optimality.

The reader should leave with five durable ideas:

- Distributions, not isolated moments. A moment is a projection; a decision can depend on the shape discarded by that projection.
- Heuristics are hypotheses about canonical objects. A useful engineering reading becomes more auditable when its target is formally defined.
- One declared objective improves comparability. Rules, learners, policies, and portfolios can be evaluated on common gain/loss geometry without claiming that they are interchangeable.
- Time is a coordinate. Increments, windows, clustering, and path-dependent mechanisms cannot be recovered after event order is erased.
- Menu is not selection. Omega-efficiency is the undominated set on the tail-aware plane. Gleim-Markowitz optimality is the max- Ω book on that set. Table 6.1's trustee mix is the

first; the security line characterizes the second.

Boundary of the book. GM futures theory is a framework for measurement, comparison, and disciplined questions. It is not an index, product, trading recommendation, or guarantee that an implementation will be profitable. The GM! Frontier is a theoretical envelope. It is not GM! Index.

The result is a continuation of the portfolio problem on a larger information set: the whole distribution, observed through time, with hurdles and uncertainty explicit.

MUTANTDeFi

A

Limited Empirical Exhibits

Consistent with the character of this volume, empirical material is confined to the program-level summaries below. Complete protocols, datasets, and validation records are available as described in Appendix D. No exhibit constitutes a claim of deployable profitability.

A.1 A.1 Exhibit 1: Loss-magnitude control (controlled A/B validation, 2025)

A controlled A/B study ($n = 481$ trades) compared a treatment arm imposing the full volatility-compression overlay of Chapter 3 against an otherwise identical control. The treatment arm was profitable (+\$888) at a 28.6% win rate; the control lost \$30,049 at a 54.5% win rate. The decisive lever was a 79% reduction in mean loss magnitude (-0.38% versus -2.3% per losing trade). The exhibit supports Proposition 3.1 and the Chapter 3 signature claim: hit rate is not the driver; the left tail is.

A.2 A.2 Exhibit 2: Gate quality from kinematic alignment (program validation, 2026)

Replacing naïve momentum cascades with the volume-aware kinematic alignment gates of Definition 3.4 reduced false entries—entries followed immediately by a protective stop—by approximately 34% in program backtesting, supporting the Chapter 2 claim that volume-coupled readings separate conviction from noise.

A.3 A.3 Exhibit 3: The minimal canonical rule and out-of-sample discipline (pre-registered battery, 2026)

In a pre-registered comparison of gating designs on six instruments spanning crypto and macro futures (hourly sampling, daily horizon), the pooled measurement profile cleared its holdout with positive out-of-sample rank correlation to forward returns (+0.032) and a positive active increment (+14.5 basis points). The minimal canonical rule—a threshold on the signed direction with a confidence floor, exactly the gate of Chapter 5—was the only arm with material positive holdout expectancy (+0.85% over an eighteen-day holdout window; 23 trades; 65.2% win rate; 0.83% maximum drawdown), and no richer overlay applied in the comparison beat it out of sample. Reconciliation of live results with the backtest remains an open problem, and it gates any advancement of performance claims.

These facts are program-level. They are not profitability claims, not evidence for Definitions S.1–S.5, and not a foundation of the theorems.

B

Peer-Review Colloquy: A Devil's Advocate Reading

Good theory survives its best critics. This appendix subjects the volume's mathematics and assertions to a respectful adversarial reading—the questions a rigorous referee should ask—together with replies and, where the honest answer is “more evidence is required,” a pointer to the roadmap of Appendix C.

B.1 On the polarity (GM-1)

Objection. Dynamic portfolio theory exists: Merton's intertemporal capital asset pricing model [7] and the multiperiod literature treat time explicitly, so “time-absent” overstates the case against the classical school. Response. Accepted in part, and the principle is stated precisely for that reason: the polarity concerns the canonical Sharpe-grade instruments as used—the single-period 1952 program [1, 6], the point-statistic Sharpe ratio, the fixed-horizon VaR—in which time is absent by intent. Intertemporal extensions restore dynamics to the investor's problem while retaining the moment premise in the risk object; the polarity claims novelty for time-inherent primitives, not for the observation that time exists.

B.2 On hurdle dependence

Objection. Ω is a curve, not a number; any headline figure invites hurdle shopping, and crossing curves (Figure 7.1) mean the theory cannot even deliver a total order on assets.

Response. Conceded as a governance risk and answered as governance: the declared-hurdle convention (Principle 6.2) requires pre-declared ladders and forbids mid-period movement. The absence of a total order is treated as information—two assets genuinely serve different hurdles—but the referee is right that league-table integrity depends on enforcement, not mathematics.

B.3 On the mirror theorem (GM-3)

Objection. $\tanh(-x) = -\tanh(x)$ is a one-line identity; presenting it as a theorem inflates elementary mathematics. Response. Agreed as mathematics, and the canon says so explicitly (Section 7.7): the contribution is the convention—the adoption of an exactly symmetric, bounded coordinate as the common reporting and learning scale—and the accompanying claim that asymmetric raw-ratio scales bias learners is a testable statement (Appendix C, item R6).

B.4 On the subsumption theorem (GM-4)

Objection. The theorem establishes shared properties—boundedness, sign structure, change structure—not estimator consistency: no result bounds the error between the kinematic channels and the canonical quantities they are said to estimate. Response. Correct, and this is the volume’s most important open mathematical edge. The correspondence is structural, and the estimator-tracking claim—that speed, velocity, and acceleration track mass, direction, and increment under stated conditions—is exactly the kind of statement that requires data (Appendix C, item R5).

B.5 On the omega security line (GM-5)

Objection. Theorem 6.1 assumes an interior optimum, differentiability (hence continuously distributed returns), and no position constraints; the compression overlay introduces atoms at the loss bound, creating exactly the kinks the proof assumes away. The Gaussian-limit reduction to classical beta is assigned as an exercise, not proven. Response. All accepted. The theorem is stated under its hypotheses; a subgradient version covering atoms and constraint sets is future work, and the Gaussian-limit reduction is flagged for formal treatment

(Appendix C, item R7). The numerical verification reported in the proof is a check, not a substitute. Definition S.5 inherits the same interior hypotheses when it invokes the security line.

B.6 On subadditivity and the combined hurdle

Objection. Proposition 6.2 adds hurdles ($\tau_X + \tau_Y$); this is natural for capital aggregation but not the only convention, and a fixed portfolio-level hurdle changes the statement. Response. Conceded. The hinge convexity that drives the result also delivers diversification-friendliness under the fixed-hurdle convention, but the proposition as stated is the additive-hurdle case and should be read as such.

B.7 On estimation under fat tails

Objection. LPM_1 is finite for tail index above one, but its sample estimator is itself heavy-tailed there; loss-mass reports and the marginal contributions (g_i, ℓ_i) may carry substantial sampling error precisely in the regimes the book cares about. Response. Fully conceded and, in this volume's view, the single most important statistical caveat: every hurdle-net figure should ship with resampled confidence intervals, and until it does, point figures are provisional (Appendix C, items R1–R3). Definition S.1 does not remove that caveat. Naming the risk coordinate does not make its estimator regular.

B.8 On the empirical exhibits

Objection. The exhibits of Appendix A are program-level: one A/B study ($n = 481$), one small holdout arm (23 trades), single-program provenance, no cross-regime replication; the false-entry reduction carries no standard error. Response. Conceded without reservation; this is why the exhibits are labeled limited, why no profitability claim is made, and why the roadmap exists (Appendix C, items R1, R2, and R4).

B.9 On compression's upside cost

Objection. Proposition 3.1 celebrates left-tail amputation, but profit capping truncates the right tail as well, which taxes the omega numerator; the net distributional effect is asserted

from one A/B study. Response. A fair point: the decomposition of compression's effect into loss-bounding benefit and cap cost is measurable and should be measured (Appendix C, item R4).

B.10 On comparing instruments of different epistemic kinds

Objection. An options-implied volatility index is a traded consensus under a risk-neutral measure. A full-distribution activity reading is an estimate under a physical measure. A table risks implying substitutability. Response. Accepted. Section 7.8 compares design classes and jobs, not market status. Settlement-based price formation and a mature ecosystem are advantages that a richer measurement surface cannot reproduce by definition.

B.11 On the reciprocal Omega convention

Objection. The literature uses both gain-over-loss and loss-over-gain orientations; monotonicity statements can appear contradictory. Response. Accepted and repaired in the formal Omega interlude. The book uses gain over loss and labels Nguyen's reciprocal orientation Ω_N .

B.12 On private-sample constants

Objection. Thresholds or effect sizes inferred from a private data set cannot be promoted to theoretical constants. Response. Accepted. The main exposition states mechanisms and assumptions. Sample-specific estimates remain in Appendix A, with limitations and replication debt in Appendix C.

B.13 On the named frontier

Objection. Linear-programming treatments already optimize Omega and already speak of an Omega frontier. Numbering a GM! Frontier risks claiming a computational discovery. Response. Accepted as a priority claim and refused as a naming claim. Chapter 6 cites Kapsos et al. [34] and Charnes–Cooper [35] for the dual picture. The objects of this volume are the two-primitive split, the declared hurdle, the security line, the polarity, and the named efficiency relation on (ρ_τ, μ) . The computational Pareto set in (G, L) is prior art; its identification with mean–omega efficiency on the tail-aware plane is the ranking identity, a corollary of

Proposition 4.2.

MUTANTDeFi

C

Statistical Significance Roadmap

This appendix flags, item by item, where further statistical significance is required to support the volume's claims. Subsequent versions of this textbook will provide the corresponding data and analyses as framed here: the item definitions below are the pre-registration.

Table C.1. Significance roadmap. Each item names the claim it supports, the evidence currently in hand, and the evidence a subsequent version must supply.

MUTANZAPHE

Item	Claim supported	Current evidence	Required evidence
R1	loss-magnitude control (Proposition 3.1; Exhibit A.1)	one A/B study, $n = 481$	replication across regimes and venues; bootstrap confidence intervals on the loss-magnitude reduction; effect stability by period
R2	minimal-rule holdout (Exhibit A.3)	23 trades, one window	rolling multi-window holdouts; exact intervals on expectancy and win rate; multiplicity control across arms
R3	hurdle-net estimation error	none reported	block-bootstrap standard errors for G, L, D, M and for (g_i, ℓ_i) under heavy tails; coverage studies
R4	compression decomposition	net A/B effect only	separate attribution of loss-bound benefit versus profit-cap cost on matched samples
R5	kinematic-to-canonical tracking (Theorem 4.2)	structural correspondence only	empirical correlation of (S, V, A) with $(M, D, \Delta\Omega)$ across instruments and granularities; tracking-error bounds
R6	raw ratio label bias	argued, not	controlled learner

No claim listed above is advanced beyond its current evidence class elsewhere in this volume; where a chapter states a result, this roadmap states its debt. Definitions S.1–S.5 are definitions, not empirical claims; they inherit the estimation debt of L and Ω (R3) and the interior hypotheses of Theorem 6.1 (R7).

MUTANTDeFi

D

Data Availability and Contact

The datasets, experimental protocols, validation records, and rolling research reports underlying Appendix A—and the broader research lineage summarized in this volume—are maintained by the MutantDeFi Research Lab of ConfigSecretAI, Inc.

Qualified academic and institutional parties may direct inquiries, data-access requests, and collaboration proposals to the Lab. X is the preferred contact channel.

Research unit	MutantDeFi Research Lab
Parent	ConfigSecretAI, Inc.
X	@mutantdefi · @billgleim
Web	mutantdefi.com
Address	4225 Executive Square, Suite 600, La Jolla, CA 92037, United States
House journal	https://apps.mutantdefi.com/research/journal/2026-08-25-full-distribution-financial-science-2e

D.1 Availability boundary

Requests are evaluated for scholarly purpose, confidentiality, licensing, privacy, market-data rights, and reproducibility. Availability is not guaranteed merely because a result is cited.

Construction details below the public methodology boundary are not disclosed in this volume or through the ordinary data-access process.

Reproducibility statement. The inability to publish a proprietary input narrows the claim that may be made from it. Theoretical statements must stand independently; empirical statements remain limited to evidence qualified readers can audit.

MUTANTDeFi

Endnote on the Name

1. “GM” records the Gleim–Markowitz research lineage: “Gleim” identifies the present full-distribution, time-inherent synthesis, while “Markowitz” identifies the portfolio-selection problem it extends. “Futures conviction” names the object of study—directional, distribution-aware assessment of market state through time—not a product line. “Futures” includes exchange-traded futures and the optional study of future states, but does not make horizon-bound modeling a theory requirement. This is a scholarly naming convention. It does not imply that Harry Markowitz authored, reviewed, or endorsed this volume.
1. In memoriam. Harry Markowitz (1927–2023) remains the root of modern portfolio selection. A brief memorial line is reserved for this endnote rather than the title page.
1. The author collaborated personally with Harry Markowitz prior to his passing. That collaboration is gratefully acknowledged; the present text is an independent scholarly development and does not claim his co-authorship, review, or endorsement.

Bibliography

- [1] H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.
- [2] J. Tobin. Liquidity preference as behavior towards risk. *The Review of Economic Studies*, 25(2):65–86, 1958.
- [3] W. F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, 1964.
- [4] J. Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, 1965.
- [5] J. Mossin. Equilibrium in a capital asset market. *Econometrica*, 34(4):768–783, 1966.
- [6] H. Markowitz. *Portfolio Selection: Efficient Diversification of Investments*. John Wiley & Sons, New York, 1959.
- [7] R. C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887, 1973.
- [8] B. Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.
- [9] E. F. Fama. The behavior of stock-market prices. *The Journal of Business*, 38(1):34–105, 1965.
- [10] S. A. Berkowitz, D. E. Logue, and E. A. Noser, Jr. The total cost of transactions on the NYSE. *The Journal of Finance*, 43(1):97–112, 1988.
- [11] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.
- [12] J.P. Morgan / Reuters. *RiskMetrics—Technical Document*. Fourth edition, New York, 1996.
- [13] P. Jorion. *Value at Risk: The New Benchmark for Managing Financial Risk*. McGraw-Hill, 3rd edition, 2007.

- [14] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [15] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.
- [16] R. C. Scott and P. A. Horvath. On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, 35(4):915–919, 1980.
- [17] A. Kraus and R. H. Litzenberger. Skewness preference and the valuation of risk assets. *The Journal of Finance*, 31(4):1085–1100, 1976.
- [18] C. R. Harvey and A. Siddique. Conditional skewness in asset pricing tests. *The Journal of Finance*, 55(3):1263–1295, 2000.
- [19] C. Keating and W. F. Shadwick. A universal performance measure. *Journal of Performance Measurement*, 6(3):59–84, 2002.
- [20] C. Keating and W. F. Shadwick. An introduction to Omega. The Finance Development Centre, London, 2002.
- [21] H. Kazemi, T. Schneeweis, and R. Gupta. Omega as a performance measure. *Journal of Performance Measurement*, 8(3):16–25, 2004.
- [22] F. A. Sortino and L. N. Price. Performance measurement in a downside risk framework. *The Journal of Investing*, 3(3):59–64, 1994.
- [23] P. D. Kaplan and J. A. Knowles. Kappa: A generalized downside risk-adjusted performance measure. *Journal of Performance Measurement*, 8(3):42–54, 2004.
- [24] W. F. Sharpe. Mutual fund performance. *The Journal of Business*, 39(1):119–138, 1966.
- [25] V. S. Bawa and E. B. Lindenberg. Capital market equilibrium in a mean–lower partial moment framework. *Journal of Financial Economics*, 5(2):189–200, 1977.
- [26] W. V. Harlow and R. K. S. Rao. Asset pricing in a generalized mean–lower partial moment framework: Theory and evidence. *Journal of Financial and Quantitative Analysis*, 24(3):285–311, 1989.
- [27] F. Black and R. Litterman. Global portfolio optimization. *Financial Analysts Journal*, 48(5):28–43, 1992.
- [28] V. N. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, New York, 1995.
- [29] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2nd edition, 2018.

- [30] M. López de Prado. *Advances in Financial Machine Learning*. John Wiley & Sons, 2018.
- [31] R. E. Whaley. The investor fear gauge. *The Journal of Portfolio Management*, 26(3):12–17, 2000.
- [32] Cboe Exchange, Inc. *The Cboe Volatility Index—VIX*. White paper, 2019.
- [33] V. N. Nguyen. *Omega Function: A Theoretical Introduction*. Plan B paper, M.A. Mathematics, University of Hawai‘i at Mānoa, 2009. <http://hdl.handle.net/10125/25918>.
- [34] M. Kapsos, S. Zymler, N. Christofides, and B. Rustem. Optimizing the Omega ratio using linear programming. *Journal of Computational Finance*, 17(4):49–57, 2014.
- [35] A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3–4):181–186, 1962.
- [36] W. L. Gleim. *GM Futures Conviction Theory*. MutantDeFi Graduate Texts in Quantitative Finance, July 2026. Prior volume.
- House journal for this book: <https://apps.mutantdefi.com/research/journal/2026-08-25-full-distribution-financial-science-2e>