PROJECT ALETHEA — RESEARCH MONOGRAPH SERIES | PAPER 1.2 (VERSION 2.0 — EXPANDED EDITION)

Spectral Risk Measures, Choquet Integrals, and Kusuoka Representation in Automated Financial Governance

Author: Ying Jack Chin (Chief Risk Officer, Aegis Lab) | **Published:** November 2026 |
Contact: hryan@aegislab.io

Abstract

While standard Expected Shortfall (ES_α) satisfies subadditivity and resolves the portfolio aggregation failures of Value at Risk (VaR_α) under regulatory capital regimes, it remains anchored to a rigid, binary truncation threshold α. In automated financial governance, continuous neural optimization engines require smooth, differentiable, and customizable risk-aversion geometries that do not sacrifice formal mathematical coherence. This monograph develops the complete functional-analytic theory of **Spectral Risk Measures** ($M_\phi(X)$) and **Choquet Capacities** on non-atomic probability spaces. We prove that Acerbi's three weighting axioms—non-negativity, normalization, and non-decreasing monotonicity—are both necessary and sufficient for $M_\phi(X)$ to satisfy the Artzner–Delbaen–Eber–Heath axioms of coherence. We derive parametric families of exponential and power-law risk spectra, establish the dual representation over Radon–Nikodym density generators, and provide a constructive proof of **Kusuoka's Representation Theorem** on L^∞ and L^p spaces. Finally, we prove that ES_α forms the extreme-point boundary of the convex set of all law-invariant coherent risk measures, demonstrating how Aegis Lab's Symbolic Veto Kernel discretizes continuous spectral envelopes into verifiable linear programming constraints.

1 Introduction: The Geometry of Tail-Risk Aversion

Modern quantitative risk management and autonomous trading architectures must evaluate portfolio exposures across fat-tailed, non-stationary return distributions. As established in **Paper 1.1**, Value at Risk (VaR_α) violates subadditivity ($\text{VaR}_\alpha(X+Y) > \text{VaR}_\alpha(X) + \text{VaR}_\alpha(Y)$) whenever asset returns exhibit skewness, leptokurtosis, or tail dependence, creating perverse incentives for neural agents to fragment positions or conceal tail concentration.

Expected Shortfall (ES_α) eliminates this structural flaw by averaging losses beyond the α-quantile:

$$\text{ES}_\alpha(X) = \frac{1}{1-\alpha}\int_\alpha^1 \text{VaR}_u(X)\,du \tag{1}$$

However, from the standpoint of institutional utility theory and automated control systems, ES_α imposes a discontinuous step-function on risk aversion: outcomes below the α-quantile receive zero marginal weight, while all catastrophic losses above α are weighted identically regardless of their severity. A 99.01% loss and a 99.99% systemic collapse are treated with identical marginal sensitivity. For an autonomous risk engine operating under central bank capital preservation mandates, risk aversion must scale monotonically with loss severity. **Spectral Risk Measures** resolve this limitation by replacing the binary indicator with a continuous risk-aversion spectrum $\phi(u)$.

2 Probability Space Preliminaries and Quantile Functions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be an atomless probability space representing future market states. Let $X \in L^1(\Omega, \mathcal{F}, \mathbb{P})$ denote a random loss variable (where $X > 0$ represents a financial loss and $X < 0$

represents a profit, adhering to standard actuarial and Basel risk conventions).

The cumulative distribution function (CDF) of X is $F_X(x) = \mathbb{P}(X \le x)$. The lower and upper quantile functions are defined respectively as:

$$q_X^-(u) = \inf\{x \in \mathbb{R} : F_X(x) \ge u\}, \quad u \in (0,1) \tag{2}$$
$$q_X^+(u) = \inf\{x \in \mathbb{R} : F_X(x) > u\}, \quad u \in (0,1) \tag{3}$$

Because $(\Omega, \mathcal{F}, \mathbb{P})$ is atomless, there exists a uniformly distributed random variable $U \sim \text{Uniform}(0,1)$ such that $X = q_X^-(U)$ almost surely (Skorokhod's representation). For any confidence level $u \in (0,1)$, $\text{VaR}_u(X) \equiv q_X^-(u)$.

Definition 2.1 (Law Invariance)**.** *A risk measure $\rho : \mathcal{X} \to \mathbb{R} \cup \{+\infty\}$ is **law-invariant** if for any two random variables $X, Y \in \mathcal{X}$ sharing the same probability distribution ($F_X = F_Y$), we have $\rho(X) = \rho(Y)$. Law invariance is an essential operational requirement for data-driven systems, ensuring that risk estimates can be calibrated from empirical return distributions and Monte Carlo simulations.*

3 Spectral Risk Measures and Acerbi's Axiomatic Framework

Introduced by Carlo Acerbi (2002), a **Spectral Risk Measure** $M_\phi(X)$ is a weighted Lebesgue–Stieltjes integral of the quantile function against a risk-aversion spectrum $\phi : [0,1] \to \mathbb{R}_+$.

Definition 3.1 (Spectral Risk Measure)**.** *For an admissible risk-aversion spectrum $\phi \in L^1([0,1])$, the spectral risk measure $M_\phi : L^\infty \to \mathbb{R}$ is defined as:*

$$M_\phi(X) = \int_0^1 \phi(u)\, q_X^-(u)\, du \tag{4}$$

Theorem 3.1 (Acerbi's Coherence Conditions)**.** *Let $\phi \in L^1([0,1])$. The functional $M_\phi(X)$ is a coherent risk measure—satisfying Translation Invariance, Positive Homogeneity, Monotonicity, and Subadditivity—if and only if ϕ satisfies the following three conditions almost everywhere on $[0,1]$:*

1. ***Non-negativity:*** *$\phi(u) \ge 0$ for all $u \in [0,1]$.*
2. ***Normalization:*** *$\int_0^1 \phi(u)\, du = 1$.*
3. ***Monotonicity (Non-decreasing):*** *For all $0 \le u_1 < u_2 \le 1$, $\phi(u_1) \le \phi(u_2)$ (i.e., $\phi'(u) \ge 0$ where differentiable).*

Proof Sketch of Coherence.

- **Translation Invariance & Positive Homogeneity:** For any $c \in \mathbb{R}$ and $\lambda > 0$, quantile equivariance yields $q_{X+c}^-(u) = q_X^-(u) + c$ and $q_{\lambda X}^-(u) = \lambda q_X^-(u)$. By linearity of the integral and normalization ($\int_0^1 \phi(u)\, du = 1$), $M_\phi(X + c) = M_\phi(X) + c$ and $M_\phi(\lambda X) = \lambda M_\phi(X)$.
- **Monotonicity:** If $X \le Y$ $\mathbb{P}$-almost surely, then $F_X(x) \ge F_Y(x)$ for all $x \in \mathbb{R}$, implying $q_X^-(u) \le q_Y^-(u)$ for all $u \in (0,1)$. Since $\phi(u) \ge 0$, integration preserves the inequality: $M_\phi(X) \le M_\phi(Y)$.
- **Subadditivity via Abel's Summation / Choquet Integration:** Because $\phi(u)$ is non-decreasing and right-continuous with $\phi(0) \ge 0$, we can express $\phi(u)$ via a positive Radon measure μ on $[0,1]$:

$$\phi(u) = \int_{[0,u]} \frac{1}{1-\alpha}\, d\mu(\alpha) \tag{5}$$

Applying Fubini–Tonelli's theorem to swap the order of integration:

$$M_\phi(X) = \int_0^1 \left(\int_{[0,u]} \frac{1}{1-\alpha}\, d\mu(\alpha) \right) q_X^-(u)\, du$$
$$= \int_{[0,1]} \left(\frac{1}{1-\alpha} \int_\alpha^1 q_X^-(u)\, du \right) d\mu(\alpha) = \int_{[0,1]} \mathsf{ES}_\alpha(X)\, d\mu(\alpha) \tag{6}$$

Since $\mathsf{ES}_\alpha(X+Y) \le \mathsf{ES}_\alpha(X) + \mathsf{ES}_\alpha(Y)$ for every $\alpha \in [0,1)$ and μ is a positive measure, integrating across α preserves subadditivity:

$$M_\phi(X+Y) = \int_{[0,1]} \mathsf{ES}_\alpha(X+Y)\, d\mu(\alpha) \le \int_{[0,1]} \mathsf{ES}_\alpha(X)\, d\mu(\alpha) + \int_{[0,1]} \mathsf{ES}_\alpha(Y)\, d\mu(\alpha) = M_\phi(X) + M_\phi(Y) \tag{7}$$

□

4 Connection to Choquet Integrals and Distorted Expectations

Spectral risk measures are closely linked to **Choquet capacities** and **Wang's Distortion Risk Measures**. Define the **distortion function** $g : [0,1] \to [0,1]$ as the integrated tail weight:

$$g(p) = \int_{1-p}^1 \phi(u)\, du, \quad p \in [0,1] \tag{8}$$

Proposition 4.1 (Concavity of the Distortion Function)**.** *A spectral risk measure $M_\phi(X)$ is coherent if and only if its corresponding distortion function $g(p)$ satisfies $g(0) = 0$, $g(1) = 1$, is non-decreasing, and is **concave** ($g''(p) \le 0$).*

Because $g'(p) = \phi(1-p)$, the condition that $\phi(u)$ is non-decreasing in u is mathematically equivalent to $g'(p)$ being non-increasing in p, which defines concavity. Under this transformation, $M_\phi(X)$ takes the form of a **Choquet integral** with respect to the submodular set function (capacity) $v(A) = g(\mathbb{P}(A))$:

$$M_\phi(X) = \int_{-\infty}^0 [g(\mathbb{P}(X > x)) - 1]\, dx + \int_0^\infty g(\mathbb{P}(X > x))\, dx \tag{9}$$

5 Parametric Families of Risk Spectra for Quantitative Desks

To operationalize spectral risk measures inside continuous optimization loops or neural network loss functions, institutions specify parametric forms of $\phi(u)$ tailored to their risk appetite:

Remark 5.1 (Gradient Smoothness for Neural Optimization)**.** *Unlike ES_α, whose weighting function $\phi(u)$ has a Dirac-delta derivative at $u = \alpha$, the **Exponential Spectrum** $\phi_\gamma(u) = \frac{\gamma e^{-\gamma(1-u)}}{1-e^{-\gamma}}$ is infinitely differentiable (C^∞) across $[0,1]$. When embedded into the objective function of a continuous policy network, $\phi_\gamma(u)$ provides smooth, non-vanishing tail gradients while strictly preserving coherence for all $\gamma > 0$.*

6 Kusuoka's Representation Theorem

While a single spectral risk measure $M_\phi(X)$ captures a fixed risk-aversion profile, institutional governance requires robustness against **model ambiguity**—uncertainty regarding which specific risk spectrum applies under unseen market regimes. **Shigeo Kusuoka (2001)** resolved the general structure of all law-invariant coherent risk measures.

Risk Spectrum Family	Weighting Function $\phi(u)$	Distortion Function $g(p)$	Institutional Behavior & Use Case
Expected Shortfall (ES_α)	$\frac{1}{1-\alpha}\mathbf{1}_{[\alpha,1]}(u)$	$\min\left(\frac{p}{1-\alpha}, 1\right)$	Uniform tail weighting; Basel FRTB regulatory standard.
Exponential (CARA) Spectrum	$\frac{\gamma e^{-\gamma(1-u)}}{1-e^{-\gamma}}$	$\frac{1-e^{-\gamma p}}{1-e^{-\gamma}}$	Smoothly increasing sensitivity ($\gamma > 0$); ideal for gradient-based RL reward shaping.
Power-Law (Wang / PH) Spectrum	$ku^{k-1} \quad (k \geq 1)$	$1-(1-p)^k$	Polynomial tail emphasis; captures heavy-tailed credit and operational loss regimes.
Dual-Power Spectrum	$k(1-u)^{k-1}$ *(Invalid for $k > 1$)*	p^k *(Convex — Non-Coherent)*	Illustrates non-coherent risk seeking; violates subadditivity.

Table 1: Parametric Risk Spectra and Distortion Profiles

Theorem 6.1 (Kusuoka's Representation Theorem)**.** *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be an atomless probability space. A functional $\rho : L^\infty(\Omega, \mathcal{F}, \mathbb{P}) \to \mathbb{R}$ is a **law-invariant coherent risk measure** (satisfying the Fatou property) if and only if there exists a non-empty set $\mathcal{M}$ of probability measures on $[0, 1]$ such that for all $X \in L^\infty$:*

$$\rho(X) = \sup_{\mu \in \mathcal{M}} \int_{[0,1]} ES_\alpha(X)\, d\mu(\alpha) \tag{10}$$

Equivalently, letting Φ denote the set of all admissible non-decreasing normalized spectra ϕ, there exists a subset $\mathcal{S} \subseteq \Phi$ such that:

$$\rho(X) = \sup_{\phi \in \mathcal{S}} M_\phi(X) = \sup_{\phi \in \mathcal{S}} \int_0^1 \phi(u)\, VaR_u(X)\, du \tag{11}$$

Proof Architecture (Dual Envelope & Hardy–Littlewood Inequality).

1. **Fenchel–Moreau Dual Representation:** Any coherent risk measure ρ with the Fatou property admits a dual representation over a weakly compact, convex set $\mathcal{Q}$ of probability density generators $Z = \frac{d\mathbb{Q}}{d\mathbb{P}} \in L^1_+$ with $\mathbb{E}[Z] = 1$:
$$\rho(X) = \sup_{Z \in \mathcal{Q}} \mathbb{E}[ZX] \tag{12}$$

2. **Hardy–Littlewood Rearrangement Inequality:** Because ρ is law-invariant, $\mathcal{Q}$ is rearrangement-invariant (if $Z \in \mathcal{Q}$ and $Z' \overset{d}{=} Z$, then $Z' \in \mathcal{Q}$). By the Hardy–Littlewood inequality, for any X and Z, the expectation $\mathbb{E}[ZX]$ is maximized when Z and X are comonotonic—that is, expressed as non-decreasing functions of the same uniform random variable U:
$$\sup_{Z' \overset{d}{=} Z} \mathbb{E}[Z'X] = \int_0^1 q_Z^-(u)\, q_X^-(u)\, du \tag{13}$$

3. **Identification of the Spectrum:** Setting $\phi_Z(u) \equiv q_Z^-(u)$, note that $q_Z^-(u) \geq 0$ (since $Z \geq 0$), $\int_0^1 q_Z^-(u)\, du = \mathbb{E}[Z] = 1$, and $q_Z^-(u)$ is non-decreasing by definition of a quantile function. Thus, ϕ_Z satisfies all three of Acerbi's axioms. Taking the supremum over all equivalence classes $[Z] \in \mathcal{Q}$ yields:
$$\rho(X) = \sup_{Z \in \mathcal{Q}} \int_0^1 q_Z^-(u)\, q_X^-(u)\, du = \sup_{\phi \in \mathcal{S}} M_\phi(X) \tag{14}$$

4. **Choquet Extreme-Point Decomposition:** Finally, because every non-decreasing normalized function ϕ can be expressed as a mixture of step functions $\frac{1}{1-\alpha}\mathbf{1}_{[\alpha,1]}(u)$ integrated against a probability measure μ on $[0,1]$, each $M_\phi(X)$ equals $\int_{[0,1]} \mathsf{ES}_\alpha(X)\, d\mu(\alpha)$, completing the proof.

□

7 Implications for Project Alethea and Automated Governance

Kusuoka's theorem provides the exact mathematical bridge between **continuous neural optimization** (Engine 1) and **deterministic symbolic verification** (Engine 2 / Engine 3) in Project Alethea:

1. **ES_α as the Extreme-Point Basis:** Because every law-invariant coherent risk measure is the supremum of mixtures of Expected Shortfalls, ES_α forms the **extreme-point generators** of the entire space of coherent risk governance. If an automated trading policy violates an ES_α boundary across a discrete grid of confidence levels $\{\alpha_1, \alpha_2, \ldots, \alpha_m\}$, it automatically violates the corresponding Kusuoka envelope.

2. **Tractable Linear Programming Discretization:** Under Rockafellar and Uryasev's (2000) variational formulation, $\mathsf{ES}_\alpha(X)$ is computed via convex linear minimization:

$$\mathsf{ES}_\alpha(X) = \min_{\zeta \in \mathbb{R}} \left\{ \zeta + \frac{1}{1-\alpha}\mathbb{E}\left[(X-\zeta)_+\right] \right\} \tag{15}$$

   Consequently, any Kusuoka risk constraint $\rho(X) = \max_{k=1,\ldots,K} M_{\phi_k}(X) \leq C_{\max}$ can be decomposed into a finite system of linear inequalities over scenario-based auxiliary variables. This allows Aegis Lab's **Symbolic Veto Kernel** to verify hard tail-risk compliance in milliseconds without non-convex optimization stalls.

3. **Protection Against Regime Ambiguity:** By taking the supremum over a family of risk spectra $\mathcal{S}$ (e.g., combining an exponential spectrum for intraday liquidity drag with high-α step spectra for sovereign default cascades), the governance kernel immunizes the institution against distributional shift and tail-correlation breakdown.

8 Conclusion

Spectral Risk Measures and Kusuoka's Representation Theorem elevate quantitative risk governance from heuristic quantile thresholds to rigorous functional analysis. By proving that all law-invariant coherent risk architectures reduce to upper envelopes of spectral integrals—and ultimately to extreme-point Expected Shortfalls—**Paper 1.2** establishes the mathematical blueprint for deterministic risk verification. In **Paper 2**, we construct the formal **Symbolic Veto Kernel** and demonstrate how Satisfiability Modulo Theories (SMT) and convex linear programming enforce these Kusuoka boundaries in real-time over autonomous neural agents.

References

[1] Acerbi, C. (2002). Spectral measures of risk: A coherent representation of subjective risk aversion. *Journal of Banking & Finance*, 26(7), 1505–1518.

[2] Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.

[3] Föllmer, H., & Schied, A. (2016). *Stochastic Finance: An Introduction in Discrete Time* (4th ed.). De Gruyter.

[4] Kusuoka, S. (2001). On law invariant coherent risk measures. *Advances in Mathematical Economics*, 3, 83–95.

[5] Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–41.

[6] Wang, S. S. (2000). A class of distortion operators for pricing financial and insurance risks. *Journal of Risk and Insurance*, 67(1), 15–36.