

wayfault: A Minimal Hexagonal Library for Wrong-Way Risk with Prescriptive Inverse Calibration

Dairon Aldaz*

June 3, 2026

Abstract

Wrong-way risk (WWR) is the adverse dependence between a counterparty’s exposure and its credit quality: exposure tends to be large precisely when the counterparty is likely to default. We present *wayfault*, a small, dependency-light (numpy-only core) Python library that quantifies WWR for a single counterparty from a Monte-Carlo exposure cube and a credit curve. The library computes baseline exposure metrics and unilateral CVA, a *conditional expected exposure given default* under a pluggable dependence model, the WWR-adjusted CVA, and the empirical multiplier $\alpha = \text{CVA}_{\text{WWR}}/\text{CVA}_{\text{indep}}$. We unify the Hull–White stochastic-hazard model and the Gaussian, Clayton and Frank copulas under a single *scenario re-weighting* core. Our main contribution is a *prescriptive inverse calibration*: given the proven monotonicity of α in the dependence parameter, we root-find the parameter that reproduces a *target* α or CVA, and we locate reverse-stress *breakpoints* where a risk metric crosses a threshold. All structural identities used by the implementation are machine-checked symbolically with SymPy, and we report a worked example on a five-year swap.

1 Introduction

Counterparty credit risk pricing charges a *credit valuation adjustment* (CVA) for the expected loss from a counterparty defaulting while owing a positive mark-to-market. The standard CVA computation assumes *independence* between exposure and default. Wrong-way risk (WWR) is the violation of this assumption in the adverse direction; its favourable mirror is right-way risk (RWR) [1, 2]. WWR can materially raise CVA and the regulatory exposure-at-default, yet practical, auditable, and dependency-light tooling is scarce.

wayfault targets this gap with a strict *hexagonal* (ports and adapters) design: a pure domain (standard library and numpy only), surrounded by protocol ports and swappable adapters for exposure sources, credit curves, dependence models, calibrators, and sinks. Beyond the

descriptive pipeline, we contribute an *inverse* layer that turns the tool prescriptive. Section 2 fixes notation; Sections 3–4 define the dependence models and WWR metrics; Section 5 presents the inverse calibration and its convergence; Section 6 reports numerics; and Section 7 describes the symbolic verification.

2 Exposure, credit and baseline metrics

Let $0 < t_1 < \dots < t_n$ be a tenor grid with increments $\Delta t_i = t_i - t_{i-1}$ ($t_0 := 0$). The *exposure cube* is a matrix $V \in \mathbb{R}^{S \times n}$ of netting-set values over S Monte-Carlo scenarios. Write $V_s(t_i)$ for entry (s, i) and $x^+ = \max(x, 0)$.

Exposure metrics. The expected positive/negative exposure and the potential future exposure at quantile q are

$$\text{EPE}(t_i) = \mathbb{E}[V(t_i)^+], \quad \text{ENE}(t_i) = \mathbb{E}[(-V(t_i))^+], \quad (1)$$

$$\text{PFE}_q(t_i) = \inf\{x : \Pr(V(t_i)^+ \leq x) \geq q\}, \quad (2)$$

estimated by the sample mean and empirical quantile over scenarios. The effective EPE is the time-weighted running maximum

$$\text{EEPE} = \frac{1}{t_n} \sum_{i=1}^n \left(\max_{j \leq i} \text{EPE}(t_j) \right) \Delta t_i. \quad (3)$$

Credit curve. The counterparty is described by a piecewise-constant hazard $\lambda(t) \geq 0$ (a flat curve being the single-segment case), giving the survival and marginal default probabilities

$$S(t) = \exp\left(-\int_0^t \lambda(u) du\right), \quad \text{PD}(t_{i-1}, t_i) = S(t_{i-1}) - S(t_i), \quad (4)$$

with recovery $R \in [0, 1]$ and loss-given-default $1 - R$. The marginal PDs telescope, $\sum_{i=1}^n \text{PD}(t_{i-1}, t_i) + S(t_n) = 1$.

Baseline CVA. With discount factors $DF(t_i)$ and discounted expected exposure $\text{EE}(t_i)$, unilateral CVA is

$$\text{CVA} = (1 - R) \sum_{i=1}^n DF(t_i) \text{EE}(t_i) \text{PD}(t_{i-1}, t_i). \quad (5)$$

*Independent. Working draft, not peer-reviewed. Source, tests and reproducible figures: github.com/daibeal/wayfault.

3 Dependence via scenario re-weighting

WWR is injected by replacing the unconditional exposure expectation at each tenor with a *conditional* one, computed by re-weighting scenarios toward those in which the counterparty is more likely to default. For tenor i with non-negative weights $w_{s,i}$ summing to one over s ,

$$\text{EE}^{\text{cond}}(t_i) = \sum_{s=1}^S V_s(t_i)^+ w_{s,i}. \quad (6)$$

Each dependence model supplies the weights.

Hull–White hazard. Following [1] the default intensity is coupled to portfolio value, $\lambda(t) = \exp(a(t) + bV(t))$, where $a(t)$ is fixed per tenor to match the marginal PDs and b is the wrong-way knob. The re-weighting is the numerically stabilised softmax

$$w_{s,i} = \frac{e^{bV_s(t_i)}}{\sum_{k=1}^S e^{bV_k(t_i)}}. \quad (7)$$

Gaussian copula. With a one-factor latent $X = \rho M + \sqrt{1 - \rho^2} Z$ and default when $X \leq \Phi^{-1}(p_i)$, conditioning on the (rank-normalised) market score y_s gives the per-scenario default probability

$$w_{s,i} \propto \Phi\left(\frac{\Phi^{-1}(p_i) + \rho y_s}{\sqrt{1 - \rho^2}}\right), \quad (8)$$

where $p_i = \text{PD}(t_{i-1}, t_i)$ and Φ is the standard normal CDF.

Archimedean copulas. For an Archimedean copula $C(u, v)$ the conditional default “grade” is the h -function $h(u | v) = \partial_v C(u, v)$, evaluated at $u = p_i$ and the market grade v_s . The Clayton and Frank generators yield the closed forms

$$h_{\text{Cl}}(u | v) = v^{-(\theta+1)}(u^{-\theta} + v^{-\theta} - 1)^{-\frac{\theta+1}{\theta}}, \quad (9)$$

$$h_{\text{Fr}}(u | v) = \frac{e^{-\theta v}(e^{-\theta u} - 1)}{(e^{-\theta} - 1) + (e^{-\theta u} - 1)(e^{-\theta v} - 1)}. \quad (10)$$

Clayton ($\theta > 0$) carries lower-tail dependence (defaults cluster with high exposure); Frank is symmetric and its sign selects the direction.

4 WWR metrics

Re-evaluating (5) with the conditional exposure (6) gives the WWR-adjusted CVA_{WWR} . The empirical multiplier and the exposure-at-default view are

$$\alpha = \frac{\text{CVA}_{\text{WWR}}}{\text{CVA}_{\text{indep}}}, \quad \text{EAD} = \alpha \cdot \text{EEPE}. \quad (11)$$

A counterparty is labelled **WRONG-WAY** when the dependence parameter is positive and $\alpha > 1$, **RIGHT-WAY** when negative and $\alpha < 1$, and **NEUTRAL** otherwise.

4.1 Arbitrage consistency

A calibrated model must reproduce the curve’s marginal PDs when integrated over exposure, otherwise it misprices the unconditional default.

Proposition 1 (Marginal consistency). *Define the per-scenario interval default probability $q_i(s) = p_i w_{s,i}$ with the softmax weights (7) and S scenarios. Then its scenario mean equals the target, $\frac{1}{S} \sum_s q_i(s) = p_i$.*

Proof. Since $\sum_s w_{s,i} = 1$, $\frac{1}{S} \sum_s p_i w_{s,i} S = p_i \sum_s w_{s,i} = p_i$. \square

5 Monotonicity and inverse calibration

The descriptive map $b \mapsto \alpha(b)$ admits a clean monotonicity, which is the engine of the inverse layer.

Lemma 1 (Score identity). *For softmax weights $w_s(b) = e^{bV_s} / \sum_k e^{bV_k}$ and any $f : \mathbb{R} \rightarrow \mathbb{R}$ with $f_s = f(V_s)$,*

$$\frac{d}{db} \mathbb{E}_w[f] = \text{Cov}_w(f, V), \quad (12)$$

where $\mathbb{E}_w[\cdot]$ and Cov_w are taken under $w(b)$.

Proof. Differentiating $\mathbb{E}_w[f] = \sum_s f_s e^{bV_s} / \sum_k e^{bV_k}$ in b gives $\mathbb{E}_w[fV] - \mathbb{E}_w[f]\mathbb{E}_w[V] = \text{Cov}_w(f, V)$. (Verified symbolically; Section 7.) \square

Theorem 1 (Monotone WWR). *At every tenor, $\text{EE}^{\text{cond}}(t_i)$ is nondecreasing in b . Consequently CVA_{WWR} and α are nondecreasing in b , with $\alpha \geq 1$ for $b > 0$ and $\alpha \leq 1$ for $b < 0$; at $b = 0$ the construction reduces to the independent case and $\alpha = 1$.*

Proof. Apply Lemma 1 with $f(V) = V^+$. Because V^+ is nondecreasing in V , f and the identity are comonotone, so by the FKG/Chebyshev sum inequality $\text{Cov}_w(V^+, V) \geq 0$; hence $\partial_b \text{EE}^{\text{cond}}(t_i) \geq 0$. As $\text{EE} \geq 0$ and the PDs and discount factors in (5) are non-negative, CVA_{WWR} inherits monotonicity, and so does $\alpha = \text{CVA}_{\text{WWR}} / \text{CVA}_{\text{indep}}$. At $b = 0$ the weights are uniform and (6) is the sample mean, i.e. the baseline. \square

5.1 The inverse problem

Let $g(\beta)$ denote a WWR metric (e.g. α , CVA_{WWR} , or EAD) as a function of a scalar dependence parameter β supplied through a model factory. By Theorem 1 (and its copula analogues) g is monotone, hence injective on any bracket $[\beta_{\text{lo}}, \beta_{\text{hi}}]$.

Table 1: WWR multiplier across dependence models (5y swap).

Model	α	Uplift
Hull-White ($b = 0.6$)	1.30	+30.3%
Gaussian ($\rho = 0.5$)	2.54	+153.8%
Frank ($\theta = 5$)	2.12	+111.9%
Clayton ($\theta = 2$)	4.11	+311.0%

Definition 1 (Target calibration / breakpoint). *Given a target g^* with $g(\beta_{lo}) \leq g^* \leq g(\beta_{hi})$, the inverse problem seeks β^* with $g(\beta^*) = g^*$. When g^* is a risk threshold (capital breach, internal limit) the solution is the reverse-stress breakpoint.*

Proposition 2 (Existence, uniqueness, convergence). *If g is continuous and strictly monotone on the bracket and g^* lies in $[g(\beta_{lo}), g(\beta_{hi})]$, then β^* exists and is unique, and bisection produces iterates with bracket width $|I_m| = 2^{-m}(\beta_{hi} - \beta_{lo})$, i.e. linear convergence with rate $1/2$.*

Proof. Existence and uniqueness follow from the intermediate value theorem and strict monotonicity. Each bisection step evaluates the midpoint and retains the sign-changing half, halving the bracket; after m steps $|I_m| = 2^{-m}|I_0| \rightarrow 0$, and the retained endpoint converges to the unique root. \square

Operationally, the solver evaluates g through the same service used for the forward computation, so a single deterministic pipeline backs both directions. The result reports the solved β^* , the achieved metric, the iteration count, a convergence flag, and the full WWR result at β^* .

6 Numerical illustration

We model a five-year receiver interest-rate swap: a humped exposure cube of $S = 20,000$ paths over $n = 20$ quarterly tenors, with a systematic rate factor inducing cross-tenor correlation, facing a BB-rated counterparty with an upward-sloping piecewise hazard ($1.5\% \rightarrow 3.5\%$), $R = 0.4$, discounted at a flat 3%. Figure 1 shows the exposure envelope and the Hull-White conditional exposure at $b = 0.6$; Figure 2 traces $\alpha(b)$ and $\text{CVA}_{\text{WWR}}(b)$, confirming Theorem 1.

At $b = 0.6$ the independent CVA 0.02240 rises to 0.02918, a +30.3% uplift, $\alpha = 1.30$, with $\text{EEPE} = 0.524$ and $\text{EAD} = 0.683$ (normalised units). Table 1 compares dependence families at illustrative strengths; the tail-dependent Clayton copula produces the largest uplift, consistent with defaults concentrating in the exposure tail.

Figure 3 verifies Proposition 1 numerically: the model-implied marginal PDs coincide with the curve targets to




Figure 1: Exposure profiles for the 5y swap. The shaded band is the wrong-way adjustment: the conditional EE given default (Hull-White, $b = 0.6$) dominates the independent EPE at every tenor.

a maximum absolute error of 8.5×10^{-16} , i.e. machine precision.

7 Reproducibility and symbolic verification

Every stochastic operation accepts an explicit seed, so results are bit-for-bit reproducible. Beyond numerical acceptance tests, the structural identities of Sections 2–5 are *machine-checked* with the SymPy computer-algebra system: equations (9)–(10) are confirmed to equal $\partial_v C$ for the Clayton and Frank copulas; the survival/PD telescoping, the piecewise cumulative-hazard integral, Proposition 1, the softmax normalisation, the Gaussian threshold of (8), the CVA integrand (5), and the score identity of Lemma 1 are each reduced symbolically to zero. The core attains 100% line coverage; the type checker (`mypy -strict`) and linter pass cleanly.

8 Conclusion

wayfault packages a complete, auditable WWR workflow behind a numpy-only core. The unifying re-weighting view places stochastic-hazard and copula models on a common footing, and the proven monotonicity of α underwrites a novel *prescriptive* layer: calibrating dependence to a target multiplier or CVA, and locating reverse-stress breakpoints. Natural extensions include bootstrap confidence intervals for α and a model-free, information-theoretic WWR intensity.




Figure 2: Multiplier α (left axis) and CVA_{WWR} (right axis) versus the coupling b . Both are monotone in b and cross $\alpha = 1$ at $b = 0$, as proved in Theorem 1.




Figure 3: Arbitrage consistency: the Hull–White model-implied marginal PDs sit exactly on the curve targets (max. error 8.5×10^{-16}).

References

- [1] J. Hull and A. White, “CVA and Wrong-Way Risk,” *Financial Analysts Journal*, vol. 68, no. 5, pp. 58–69, 2012.
- [2] J. Gregory, *The xVA Challenge: Counterparty Risk, Funding, Collateral, and Capital*, 3rd ed. Wiley, 2015.
- [3] Basel Committee on Banking Supervision, “Basel III: A global regulatory framework for more resilient banks and banking systems,” BIS, 2011.
- [4] R. B. Nelsen, *An Introduction to Copulas*, 2nd ed. Springer, 2006.