Assessing the presence of physical risk with a structural credit risk model

Elia Smaniotto

joint work with Davide Radi, Alessandro Sbuelz, Andrea Tarelli

Università degli Studi di Pavia

Seminar
May 18^{th}, 2026

Introduction

- We propose a structural credit-risk model with downward jumps, where the total jump intensity is affine in a stochastic physical-risk factor.
- The physical-risk factor is modeled as a CIR process, estimated from a media-based proxy for climate concerns, and then used to identify the corresponding market price of risk from US equity returns.
- To define the model under the risk-neutral measure $\mathcal{Q}$, we introduce a stochastic discount factor with two priced sources of risk: aggregate market risk and systematic physical risk.
- To price CDS, we derive a first-order perturbation approximation of the Laplace transform of the first-passage probability, which yields a tractable semi-closed-form pricing formula.
- The model is estimated by Sequential Monte Carlo (SMC^2) on monthly US CDS spreads, combining an outer sampler for static parameters with an inner particle filter for the latent firm-value process.

The model: firm-value dynamics

We consider a **jump structural credit-risk model** where $V_{j,t}$ denotes the value of firm j. Under the real-world measure $\mathcal{P}$, its dynamics is:

$$\frac{dV_{j,t}}{V_{j,t}} = \mu_{j,t}\,dt + \sigma_{j,M}\,dW_{M,t} + \sigma_j\,dW_{j,t} + d\left(\sum_{i=1}^{N_{j,t}} (e^{Z_{j,i}} - 1)\right) \tag{1}$$

- $\sigma_{j,M}$: market-risk exposure of firm j, loading on the aggregate shock $W_{M,t}$;
- σ_j: idiosyncratic volatility of firm j, loading on the firm-specific shock $W_{j,t}$.

The jump component collects all negative jumps affecting firm value, where $N_{j,t}$ is a Poisson process and $Z_{j,i} < 0$ is the size of the i-th downward jump, in particular:

$$Z_{j,i} \sim -\mathrm{Exp}(\eta_j)$$

with density

$$f_{Z_j}(z) = \eta_j e^{\eta_j z}\mathbf{1}_{\{z<0\}}$$

The model: jump intensity and physical risk

The jump intensity is stochastic and is affine on an aggregate physical-risk factor:

$$\lambda_{j,t} = a_j + b_j \lambda_{E,t} \tag{2}$$

having dynamics:

$$d\lambda_{E,t} = k_E(\theta_E - \lambda_{E,t})\,dt + \psi_E\sqrt{\lambda_{E,t}}\,dW_{E,t} \tag{3}$$

The jump intensity can be therefore decomposed as:

- a_j: baseline jump intensity of firm j;
- $b_j\lambda_{E,t}$: physical-risk-driven jump intensity;
- b_j: sensitivity of firm j's jump intensity to physical risk.

Risk-neutral measure $\mathcal{Q}$

To move from $\mathcal{P}$ to $\mathcal{Q}$, we introduce the stochastic discount factor:

$$\frac{\mathrm{d}M_t}{M_t} = -r\,\mathrm{d}t - b_M\,\mathrm{d}W_{M,t} - b_E\sqrt{\lambda_{E,t}}\,\mathrm{d}W_{E,t}$$

$$\begin{aligned} b_M : &\quad \text{market price of aggregate market risk} \\ b_E\sqrt{\lambda_{E,t}} : &\quad \text{market price of systematic physical risk} \end{aligned}$$

Defining $X_{j,t} := \log(V_{j,t})$, under the risk-neutral measure $\mathcal{Q}$ the model becomes

$$\begin{cases} \mathrm{d}X_{j,t} = \phi^{\mathcal{Q}}_{j,t}\,\mathrm{d}t + \sigma_{j,tot}\,\mathrm{d}W^{\mathcal{Q}}_{j,t} + \mathrm{d}\left(\sum_{i=1}^{N_{j,t}} Z_{j,i}\right) \\ \mathrm{d}\lambda_{E,t} = k^{\mathcal{Q}}_E\left(\theta^{\mathcal{Q}}_E - \lambda_{E,t}\right)\mathrm{d}t + \psi_E\sqrt{\lambda_{E,t}}\,\mathrm{d}W_{E,t} \end{cases}$$

with

$$\phi^{\mathcal{Q}}_{j,t} = r - \frac{1}{2}\sigma^2_{j,\mathrm{tot}} - \lambda^J_{j,t}\,\xi_j \qquad \lambda^J_{j,t} = a_j + b_j\lambda_{E,t}$$

$$\sigma^2_{j,\mathrm{tot}} = \sigma^2_{j,M} + \sigma^2_j \qquad k^{\mathcal{Q}}_E = k_E + \psi_E b_E \qquad \theta^{\mathcal{Q}}_E = \frac{k_E\theta_E}{k_E + \psi_E b_E} \qquad \xi_j = \frac{\eta_j}{\eta_j + 1} - 1$$

Default time and survival probability

In the structural framework, default occurs when the firm value hits the default barrier D from above. Defining:

$$X_{j,t} := \log\left(\frac{V_{j,t}}{D_j}\right)$$

the default time is the **first-passage time**:

$$\tau_d := \inf\{t \geq 0 : \; X_t \leq 0\}$$

Hence, the *undiscounted* risk-neutral default probability over the interval $(t, T]$, for a firm j, is:

$$H(t, X_{j,t}, \lambda_E; T) = \mathbb{E}^{\mathcal{Q}}\left[\mathbf{1}_{\{t \leq \tau_d \leq T\}} | X_{j,t}, \lambda_{E,t}\right]$$

The two-dimensional boundary-value problem

Let[1]:

$$\Phi(t,X,\lambda_E;\omega) = \omega\,\mathcal{L}(H(t,X,\lambda_E;T))\,(\omega) = \mathbb{E}^{\mathcal{Q}}\left[e^{-\omega(\tau_d - t)}\,\middle|\,\mathcal{F}_t\right], \qquad \omega > 0$$

then Φ solves

$$\begin{cases} (\mathcal{A}-\omega)\,\Phi(t,X,\lambda_E;\omega) = 0 & X>0,\ \lambda_E>0, \\ \Phi(t,0,\lambda_E;\omega) = 1 & X\leq 0,\ \lambda_E>0, \\ \Phi(t,X,\lambda_E;\omega) \to 0 & \text{as } X\to\infty \end{cases}$$

where the model infinitesimal generator under $\mathcal{Q}$ reads $\mathcal{A} = \mathcal{A}_0 + \mathcal{A}_{\lambda_E} + (a+b\lambda_E)\,\mathcal{J}$, with:

$$\mathcal{A}_0\Phi = \phi^{\mathcal{Q}}\,\Phi_X + \tfrac{1}{2}\sigma_{\text{tot}}^2\,\Phi_{XX}$$

$$\mathcal{A}_{\lambda_E}\Phi = k_E^{\mathcal{Q}}(\theta_E^{\mathcal{Q}} - \lambda_E)\,\Phi_{\lambda_E} + \tfrac{1}{2}\psi_E^2\,\lambda_E\,\Phi_{\lambda_E\lambda_E}$$

$$\mathcal{J}\Phi = \int_{-X}^{0}\big(\Phi(X+z,\lambda_E) - \Phi(X,\lambda_E)\big)\,\eta e^{\eta z}\,dz + \int_{-\infty}^{-X}\big(1-\Phi(X,\lambda_E)\big)\,\eta e^{\eta z}\,dz$$

so that the second integral accounts for jumps that overshoot the default barrier.

[1] we omit references with the respect to firm j and time t

Setup: a singularly perturbed Laplace-domain problem

Under the hypothesis that the CIR process evolves in a faster scale compared to the dynamics of the firm, we can apply multiscale methods. We consider a family of *fast mean reverting* CIR processes:

$$d\lambda_{E,t} = \frac{k_E}{\varepsilon}\left(\theta_E - \lambda_{E,t}\right) dt + \frac{\psi_E}{\sqrt{\varepsilon}} \sqrt{\lambda_{E,t}}\, dW^{\mathcal{Q}}_{E,t}, \qquad 0 < \varepsilon < 1.$$

where the 2D boundary value problem becomes singularly perturbed [Fouque et al., 2011, Pavliotis and Stuart, 2008]:

$$(\omega - \mathcal{A}^{\varepsilon})\Phi^{\varepsilon} = 0, \qquad \mathcal{A}^{\varepsilon} = \frac{1}{\varepsilon}\mathcal{A}_{\lambda_E} + \mathcal{A}_X(\lambda_E)$$

$\mathcal{A}_{\lambda_E}$	infinitesimal generator of the CIR factor (*fast*)
$\mathcal{A}_X(\lambda_E) = \mathcal{A}_0 + (a + b\lambda_E)\,\mathcal{J}$	generator of the *slow* X-dynamics
$\mathcal{J}$	jump-integral operator in X
π	invariant distribution of λ_E

$$\mathbb{E}_{\pi}[\lambda_E] = \theta_E \qquad \mathrm{Var}_{\pi}(\lambda_E) = \frac{\theta_E \psi_E^2}{2\kappa_E} \tag{4}$$

The small parameter multiplies the fast generator, so it sits in front of the highest derivatives in $\lambda_E \to$ *singular* perturbation

Asymptotic ansatz and the order-by-order hierarchy

We look for a solution *(singular perturbed expansion)* in the form:

$$\Phi^{\varepsilon}(t, X, \lambda_E; \omega) = \Phi_0 + \varepsilon\, \Phi_1 + \varepsilon^2\, \Phi_2 + O(\varepsilon^3)$$

Substituting in $(\omega - \mathcal{A}^{\varepsilon})\Phi^{\varepsilon} = 0$ and matching powers of ε:

$$\begin{aligned} O(\varepsilon^{-1}): &\quad \mathcal{A}_{\lambda_E}\, \Phi_0 = 0 \\ O(1): &\quad \mathcal{A}_{\lambda_E}\, \Phi_1 = \big(\omega - \mathcal{A}_X(\lambda_E)\big)\Phi_0 \\ O(\varepsilon): &\quad \mathcal{A}_{\lambda_E}\, \Phi_2 = \big(\omega - \mathcal{A}_X(\lambda_E)\big)\Phi_1 \end{aligned}$$

Where the default boundary condition is imposed pointwise (at $X = 0$, for every λ_E):

$$\Phi_0(\omega, 0, \lambda_E) = 1 \qquad \Phi_1(\omega, 0, \lambda_E) = 0 \qquad \Phi_2(\omega, 0, \lambda_E) = 0$$

Reading the hierarchy

Each line is a Poisson equation in the fast variable, forced by the previous order. A Poisson equation $\mathcal{A}_{\lambda_E} u = h(\lambda_E)$ is solvable **iff** the forcing is centered under π, $\mathbb{E}_\pi[h(\lambda_E)] = 0$. This Fredholm-type condition is exactly what produces *averaging*.

Leading order: averaging out the fast factor

- At order $O(\varepsilon^{-1})$, the equation

$$\mathcal{A}_{\lambda_E}\,\Phi_0 = 0$$

forces $\Phi_0 \in \ker(\mathcal{A}_{\lambda_E})$. Hence:

$$\Phi_0(\omega, X, \lambda_E) = \Phi_0(\omega, X), \qquad \Longrightarrow \qquad \Phi_0 \text{ does not depend on the fast factor.}$$

- At order $O(1)$, the solvability condition requires that the forcing is centered:

$$\mathbb{E}_\pi\big[(\omega - \mathcal{A}_X(\lambda_E))\Phi_0\big] = 0$$

With $\mathcal{A}_X(\lambda_E) = \mathcal{A}_0 + (a + b\lambda_E)\mathcal{J}$ and $\mathbb{E}_\pi[\lambda_E] = \theta_E$, define the **averaged generator**:[2]

$$\overline{\mathcal{A}}_X \;\equiv\; \mathbb{E}_\pi\big[\mathcal{A}_X(\lambda_E)\big] = \mathcal{A}_0 + (a + b\theta_E)\,\mathcal{J} \qquad (\omega - \overline{\mathcal{A}}_X)\Phi_0 = 0$$

At leading order, the fast factor enters *only through its mean* θ_E. All fluctuations of λ_E are averaged out — the slow dynamics see a constant climate intensity $a + b\theta_E$.

[2]We note that the averaged leading-order problem coincides with the one-dimensional Kou first-passage jump-diffusion case [Kou and Wang, 2003]

First-order corrector (I): the centered piece Φ_1^c

The first-order corrector solves a non-homogeneous Poisson equation and therefore decomposes into a centered particular solution plus a homogeneous component:

$$\Phi_1(\omega, X, \lambda_E) = \underbrace{\Phi_1^c(\omega, X, \lambda_E)}_{\text{centered particular solution}} + \underbrace{\Phi_{1,0}(\omega, X)}_{\in \ker(\mathcal{A}_{\lambda_E})}$$

Subtracting the average from the $O(1)$ equation, yields:

$$\mathcal{A}_{\lambda_E} \Phi_1^c = -b\,(\lambda_E - \theta_E)\,\mathcal{J}\,\Phi_0$$

Hence, the auxiliary Poisson equation $\mathcal{A}_{\lambda_E}\Theta_1 = -(\lambda_E - \theta_E)$ has explicit solution $\Theta_1(\lambda_E) = (\lambda_E - \theta_E)/k_E$, hence

$$\Phi_1^c = \frac{b}{k_E}\,(\lambda_E - \theta_E)\,\mathcal{J}\,\Phi_0 \tag{5}$$

Φ_1^c is the *instantaneous response* to deviations of the climate factor from its mean: it vanishes at $\lambda_E = \theta_E$, grows linearly with the sensitivity b, and is damped by the mean-reversion speed k_E — strong mean reversion suppresses the correction.

First-order corrector (II): the averaged piece $\Phi_{1,0}$

$\Phi_{1,0}$ is independent of λ_E and is fixed by solvability at the *next* order:

$$\mathbb{E}_\pi\big[(\omega - \mathcal{A}_X(\lambda_E))\,\Phi_1\big] = 0$$

Inserting $\Phi_1 = \Phi_1^c + \Phi_{1,0}$, the centered piece Φ_1^c generates a $\mathcal{J}^2\Phi_0$ forcing, weighted by the *quadratic* fluctuations $\mathbb{E}_\pi[(\lambda_E - \theta_E)^2]$. Specifically, for the CIR factor, we have:

$$\mathbb{E}_\pi[(\lambda_E - \theta_E)^2] = \frac{\theta_E \psi_E^2}{2k_E}$$

Hence $\Phi_{1,0}$ solves an averaged Poisson equation in X:

$$(\omega - \overline{\mathcal{A}}_X)\,\Phi_{1,0} \;=\; m_2\,\mathcal{J}^2\,\Phi_0, \qquad m_2 \;\equiv\; \frac{b^2\,\theta_E\,\psi_E^2}{2\,k_E^2}$$

The *variance* of the fast factor moves into the slow dynamics as a second-order forcing of strength m_2. After averaging, climate volatility leaves a footprint on default through the operator $\mathcal{J}^2\Phi_0$.

Outer approximation

The final expression reads:

$$\Phi^{\varepsilon}_{\text{out}}(\omega, X, \lambda_E) \;=\; \Phi_0(\omega, X) \;+\; \varepsilon\Big[\, \tfrac{b}{k_E}(\lambda_E - \theta_E)\,\mathcal{J}\,\Phi_0 \;+\; \Phi_{1,0}(\omega, X) \,\Big] + O(\varepsilon^2).$$

Leading Φ_0	**Centered** Φ_1^c	**Averaged** $\Phi_{1,0}$
$(\omega - \overline{\mathcal{A}}_X)\Phi_0 = 0$	$\dfrac{b}{k_E}(\lambda_E - \theta_E)\,\mathcal{J}\,\Phi_0$	$(\omega - \overline{\mathcal{A}}_X)\Phi_{1,0} = m_2\mathcal{J}^2\Phi_0$
sees only the *mean* of λ_E	fluctuation $\lambda_E - \theta_E$ around the mean	*variance* $\propto \psi_E^2$ feeds back into X

Boundary condition

At $X = 0$, $\Phi_1^c = \frac{b}{k_E}(\lambda_E - \theta_E)\,\mathcal{J}\Phi_0(\omega, 0) \neq 0$, hence $\Phi^{\varepsilon}_{\text{out}}$ does **not** match the default condition $\Phi^{\varepsilon}(\omega, 0, \lambda_E) = 1$ pointwise in λ_E.

The residual is killed by a *boundary layer corrector*.

Boundary-layer corrector at the default barrier

The outer solution misses the BC by an $O(\varepsilon)$ residual that still depends on λ_E:

$$\Phi^{\varepsilon}_{\text{out}}(\omega, 0, \lambda_E) - 1 \;=\; \varepsilon \, \tfrac{b}{k_E} (\lambda_E - \theta_E) \, \mathcal{J} \, \Phi_0(\omega, 0) + O(\varepsilon^2)$$

To account this problem, standard approach is to stretch near the barrier with the inner variable $\rho = X/\sqrt{\varepsilon}$ and add a localized corrector B_1:

$$\Phi^{\varepsilon}_{\text{app}} \;=\; \underbrace{\Phi_0 + \varepsilon \, \Phi_1}_{\text{outer}} + \underbrace{\varepsilon \, B_1}_{\text{boundary layer}}$$

B_1 kills the residual at $X = 0$ and decays exponentially away from it:

$$B_1 = -\tfrac{b}{k_E} (\lambda_E - \theta_E) \, \mathcal{J} \, \Phi_0(\omega, 0) \, \exp\left(-\frac{\sqrt{2k_E}}{\sigma_{\text{tot}}} \frac{X}{\sqrt{\varepsilon}} \right)$$

The boundary layer has thickness $\sqrt{\varepsilon}\,\sigma_{\text{tot}}/\sqrt{2k_E}$ — it shrinks with ε and with stronger mean reversion k_E.

Spatial picture: where the approximation lives

boundary layer **outer (averaged) region**

width $\sqrt{\varepsilon}\,\sigma_{\text{tot}}/\sqrt{2k_E}$

$\Phi^{\varepsilon}_{\text{out}} = \Phi_0 + \varepsilon\Phi_1$, error $O(\varepsilon^2)$

$X = 0$ (default)

- Away from the boundary, the solution is the outer expansion:

$$\Phi^{\varepsilon}_{\text{out}}(\omega, X, \lambda_E) = \Phi_0(\omega, X) + \varepsilon\left[\frac{b}{k_E}(\lambda_E - \theta_E)\mathcal{J}\Phi_0(\omega, X) + \Phi_{1,0}(\omega, X)\right]$$

- Near default, when $X = O(\sqrt{\varepsilon})$, the outer approximation is corrected by a boundary-layer term:

$$\Phi^{\varepsilon}_{\text{app}}(\omega, X, \lambda_E) = \Phi^{\varepsilon}_{\text{out}}(\omega, X, \lambda_E) + \varepsilon B_1\left(\omega, \frac{X}{\sqrt{\varepsilon}}, \lambda_E\right)$$

Hence, the correction B_1 restores the boundary condition $\Phi^{\varepsilon}_{\text{app}}(\omega, 0, \lambda_E) = 1$

Validation of the survival approximation

• To validate the pricing method, we compute the survival probability $S^\varepsilon(t, X, \lambda_E; T)$ from the full two-dimensional model, where ($\tau = T - t$):

$$\partial_\tau S^\varepsilon = \mathcal{L}^\varepsilon S^\varepsilon \qquad \mathcal{L}^\varepsilon = \mathcal{A}_X(\lambda_E) + \frac{1}{\varepsilon}\mathcal{A}_{\lambda_E} + \mathcal{J}(\lambda_E)$$

On the grid $(X_i, \lambda_{E,j})$, the benchmark is computed with an ADI / Crank–Nicolson scheme:

$$\frac{S^{n+1} - S^n}{\Delta\tau} \approx \frac{1}{2}\left[\mathcal{L}_X(S^{n+1} + S^n) + \mathcal{L}_{\lambda_E}(S^{n+1} + S^n)\right] + \mathcal{J}S^n$$

• Second, the approximated survival probability is derived from the singular expansion by numerically inverting the Laplace transform $\widehat{S}^\varepsilon(\omega, X, \lambda_E)$

$$S^{\mathrm{Lap}}(t, X, \lambda_E; T) = \mathcal{L}^{-1}\left[\widehat{S}^\varepsilon(\omega, X, \lambda_E)\right](\tau)$$

by using the Gaver–Stehfest algorithm.

Validation idea

We compare S^{Lap} with the 2D ADI benchmark S^{FDM}. As ε decreases, the fast physical-risk factor approaches its averaged regime, so the perturbation approximation should converge to the finite-difference benchmark.

Relative error across values of ε

We validate the first-order approximation by comparing it with the 2D finite-difference solution:

$$\mathrm{RelErr}(T,\varepsilon) = 100 \times \left| \frac{S^{\mathrm{Lap}}(T,\varepsilon) - S^{\mathrm{FDM}}(T,\varepsilon)}{S^{\mathrm{FDM}}(T,\varepsilon)} \right|$$

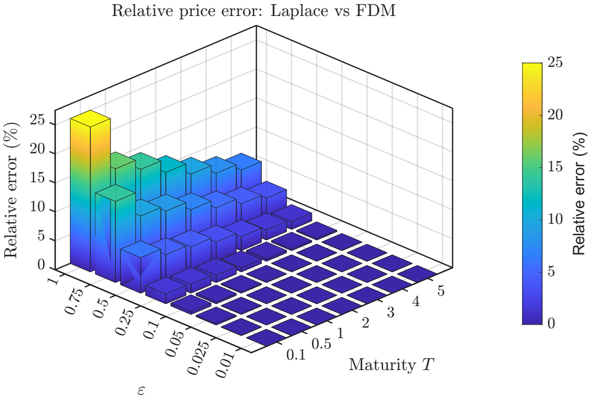

Average relative error

ε	Error (%)
1	12.943
0.75	7.557
0.5	3.439
0.25	0.872
0.1	0.156
0.05	0.041
0.025	0.014
0.01	0.008

The approximation error decreases as the separation-of-scales parameter ε becomes smaller.

Average comp. time FDM $= 3000$ sec

Average comp. time Laplace $= 0.25$ sec

Parameter	Value	Parameter	Value	Parameter	Value	Parameter	Value
a	0.2	b	0.2	σ	0.25	η	3
λ_0^E	0.1	θ_E	1	k_E	0.2	ψ_E	0.1
X_0	1	M_{Gaver}	8				

From survival Laplace transform to CDS spreads

Credit Default Swaps are contracts that protect the buyer against the default of a firm:

- the buyer pays a continuous premium until maturity or default
- if default occurs, the seller compensates the buyer by the loss-given-default $1 - \Upsilon$

The **fair CDS spread** equates the present values of the two legs, $\mathrm{PV}_{\mathrm{Prem}} = \mathrm{PV}_{\mathrm{Prot}}$, yielding

$$\Pi(t, X, \lambda_E; T) = (1 - \Upsilon)\, \frac{\mathbb{E}^{\mathcal{Q}}\left[e^{-r(\tau_d - t)} \mathbf{1}_{\{\tau_d \leq T\}}\right]}{\mathbb{E}^{\mathcal{Q}}\left[\int_t^T e^{-r(z-t)} \mathbf{1}_{\{\tau_d > z\}}\, dz\right]}.$$

Both expectations are linear functionals of survival probabilities, hence of Φ. Applying a Laplace transform in T and inverting, we get:

$$\Pi(t, X, \lambda_E; T) = \frac{\mathcal{L}^{-1}\big(F_{\mathrm{Prot}}(\omega)\big)}{\mathcal{L}^{-1}\big(F_{\mathrm{Prem}}(\omega)\big)}$$

with

$$F_{\mathrm{Prot}}(\omega) = \Phi^{\varepsilon}(t, X, \lambda_E; \omega + r)\, \frac{1 - \Upsilon}{\omega}$$

$$F_{\mathrm{Prem}}(\omega) = \big(1 - \Phi^{\varepsilon}(t, X, \lambda_E; \omega + r)\big) \left(\frac{1}{r\omega} - \frac{1}{\omega(r + \omega)}\right)$$

CDS pricing validation: ADI vs. Laplace inversion

ADI benchmark for CDS. From the ADI survival curve $S(u)$, the CDS benchmark is recovered by integrating the discounted survival probability:

$$\Pi^{\mathrm{ADI}}(t, X, \lambda_E; T) = (1 - \Upsilon)\left(\frac{1 - e^{-r(T-t)}S(T)}{\int_t^T e^{-r(u-t)}S(u)\,du} - r\right)$$

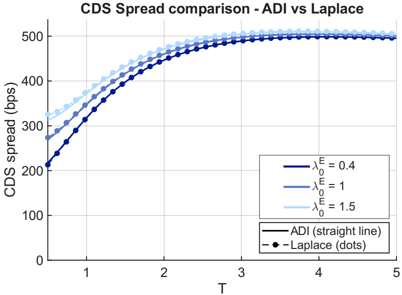

(a) $X_0 = 0.6$

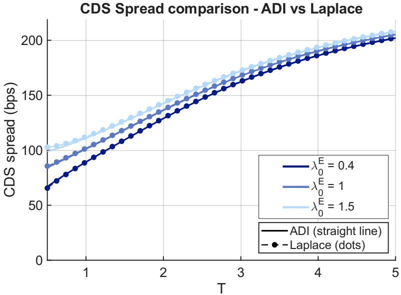

(b) $X_0 = 1$

The pricing has been validated by setting $\varepsilon = 1$, with parameters of the CIR process aligned with empirical findings *(fast-mean reverting)*.

Physical-risk intensity: a CIR specification

We model the aggregate physical-risk intensity $\lambda_{E,t}$ as a Cox–Ingersoll–Ross (CIR) process:

$$\mathrm{d}\lambda_{E,t} = k_E(\theta_E - \lambda_{E,t})\,\mathrm{d}t + \psi_E\sqrt{\lambda_{E,t}}\,\mathrm{d}W_{E,t} \tag{6}$$

Data. Monthly *Media Climate Change Concerns – Environmental Impact* index from *Sentometrics*, January 2003 – July 2025. The series is de-trended, preserving its sample mean.

Estimation. Maximum likelihood, using the exact CIR transition density (a scaled non-central χ^2).

Parameter	Estimate	Std. Error
k_E	7.5364	0.8954
θ_E	0.6583	0.0199
ψ_E	0.8487	0.0179

The fitted process delivers the standardized physical-risk innovations $\varepsilon_{E,t+1}$ used in the next step.

Pricing framework: SDF and market index

We recall the stochastic discount factor introduced above, which prices two systematic sources of risk – aggregate market risk and systematic physical risk – as:

$$\frac{\mathrm{d}M_t}{M_t} = -r\,\mathrm{d}t - b_M\,\mathrm{d}W_{M,t} - b_E\sqrt{\lambda_{E,t}}\,\mathrm{d}W_{E,t}$$

where b_M and b_E are the market prices of the two risks.

We consider a market index I_t loading on both shocks:

$$\frac{\mathrm{d}I_t}{I_t} = \mu_t\,\mathrm{d}t + \sigma_{IM}\,\mathrm{d}W_{M,t} + \sigma_{IE}\sqrt{\lambda_{E,t}}\,\mathrm{d}W_{E,t}$$

where the physical-risk loading is scaled by $\sqrt{\lambda_{E,t}}$, so its contribution to the index variance is time-varying and increases with the level of climate concern.

From theory to empirical specification

By no-arbitrage condition $\mathbb{E}_t[\mathrm{d}(M_t I_t)] = 0$, we derive:

$$\frac{\Delta I_{t+1}}{I_t} = (r + b_M \sigma_{IM} + b_E \sigma_{IE} \lambda_{E,t})\,\Delta t + \sigma_{IM}\sqrt{\Delta t}\,\varepsilon_{M,t+1} + \sigma_{IE}\sqrt{\lambda_{E,t}\,\Delta t}\,\varepsilon_{E,t+1} \quad (7)$$

which maps directly into the regression:

$$\frac{\Delta I_{t+1}}{I_t} - r\,\Delta t = \alpha + \beta_1\,\lambda_{E,t} + \beta_2\,\sqrt{\lambda_{E,t}}\,\varepsilon_{E,t+1} + u_{t+1} \quad (8)$$

having coefficients mapping:

$$\alpha = \sigma_{IM} b_M \Delta t, \qquad \beta_1 = \sigma_{IE} b_E \Delta t, \qquad \beta_2 = \sigma_{IE}\sqrt{\Delta t}, \qquad u_{t+1} = \sigma_{IM}\,\varepsilon_{M,t+1}.$$

and $\lambda_{E,t}$ and $\varepsilon_{E,t+1}$ come from the CIR estimation, while excess market returns are taken from the Kenneth French data library.

Risk-premia estimation: results

We estimate (8) on U.S. monthly excess returns over January 2010 – July 2025. The four parameters of interest are:

- b_M: market price of aggregate market risk
- b_E: market price of physical risk
- σ_{IM}: index loading on market shocks
- σ_{IE}: index loading on physical-risk shocks

Parameter	**Estimate**	**Std. Error**
σ_{IM}	0.1473	0.0076
σ_{IE}	-0.0237	0.0073
b_M	1.4370	0.7122
b_E	-1.4233	3.3848

Aggregate market risk carries a positive and significant price. Physical risk is priced *negatively* – consistent with investors accepting lower expected returns on assets that hedge climate concerns – but the estimate is imprecise.

Model estimation

The model is estimated by *Sequential Monte Carlo* (SMC^2), combining an outer sampler for static parameters with an inner particle filter for the latent state. In the empirical implementation, we condition on the observed paths $\{\lambda_{E,t}, \varepsilon_{M,t}\}_{t=1}^{T}$, so that the only latent process is X_t.

The state-space specification is:

$$\begin{cases} \lambda_{E,t+1} = \lambda_{E,t} + k_E\left(\theta_E - \lambda_{E,t}\right)\Delta t + \psi_E\sqrt{\lambda_{E,t}}\sqrt{\Delta t}\,\varepsilon_{E,t+1}, \\ \lambda_{j,t} = a_j + b_j\lambda_{E,t}, \\ \phi_{j,t} = r + b_M\sigma_{jM} - \dfrac{1}{2}\left(\sigma_{jM}^2 + \sigma_j^2\right) - \xi_j\lambda_{j,t}, \\ X_{j,t+1} = X_{j,t} + \phi_{j,t}\Delta t + \sigma_{jM}\sqrt{\Delta t}\,\varepsilon_{M,t+1} + \sigma_j\sqrt{\Delta t}\,\varepsilon_{j,t+1} + \sum_{i=1}^{N_{j,t+1}} Z_{j,i}, \\ \boldsymbol{\Pi}_j^{\text{obs}}(t,\boldsymbol{\tau}) = \boldsymbol{\Pi}_j\big(t, X_{j,t}, \lambda_{E,t}; \boldsymbol{\tau}\big) + \boldsymbol{\varepsilon}_{j,t}^{\text{err}}(\boldsymbol{\tau}), \qquad \boldsymbol{\varepsilon}_{j,t}^{\text{err}}(\boldsymbol{\tau}) \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{j,\Pi}) \end{cases} \tag{9}$$

The observation covariance is parameterized as

$$\Sigma_\Pi = \sigma_{id}^2 I_K + \sigma_{lev}^2 \mathbf{1}\mathbf{1}^\top$$

so as to capture both tenor-specific noise and a common level component across the CDS term structure.

The vector of firm-specific static parameters is

$$\Theta = (X_0, \sigma, \sigma_M, \eta, a_\lambda, b_\lambda, \sigma_{lev}, \sigma_{id})$$

Sequential Monte Carlo estimation: inner filter

Since the model is nonlinear and non-Gaussian, the likelihood is not available in closed form. The procedure is constituted as:

Inner filter: guided auxiliary particle filter
Given Θ, the latent state X_t is filtered with a guided APF:

$$X_t^{(i)} \sim q_t\Big(\cdot \mid X_{t-1}^{(i)}, \Pi_t^{mkt}, \lambda_{E,t}, \varepsilon_{M,t}, \Theta\Big).$$

The transition density

$$p(X_t \mid X_{t-1}, \lambda_{E,t}, \varepsilon_{M,t}, \Theta)$$

is computed from the model characteristic function using FFT.

The proposal density $q_t(\cdot)$ and the look-ahead score are obtained using an Unscented Kalman Filter (UKF).

The measurement density is:

$$g_t(\Pi_t^{mkt} \mid X_t, \lambda_{E,t}, \Theta) = \mathcal{N}\big(\Pi_t^{mkt}; \Pi_t(X_t, \lambda_{E,t}), \Sigma_\Pi\big)$$

Sequential Monte Carlo estimation: outer sampler

Outer sampler: tempered SMC over static parameters

The outer SMC sampler targets a sequence of tempered distributions:

$$\pi_\beta(\Theta) \propto p(\Theta)\, p(\Pi_{1:T}^{mkt} \mid \Theta)^{\beta}, \qquad \beta \in [0,1]$$

The tempering parameter β moves smoothly from the prior to the posterior:

$$\beta = 0 \quad \Rightarrow \quad \pi_0(\Theta) = p(\Theta),$$

$$\beta = 1 \quad \Rightarrow \quad \pi_1(\Theta) \propto p(\Theta)\, p(\Pi_{1:T}^{mkt} \mid \Theta)$$

Parameter particles are sequentially:

- reweighted using the particle-filter likelihood estimate;
- resampled when the ESS falls below a threshold;
- rejuvenated by Metropolis–Hastings moves.

The output is the posterior distribution of Θ, together with filtered and smoothed estimates of the latent firm-value process $X_{1:T}$.

Example of preliminary results: CenterPoint Energy - 5y CDS

Prior and posterior quantiles for the primitive parameters. Sample: January 2010–July 2025.

	Prior					Posterior				
Parameter	5%	25%	Median	75%	95%	5%	25%	Median	75%	95%
σ_j	0.0494	0.1440	0.2482	0.3670	0.5535	0.2497	0.2971	0.2971	0.3213	0.3894
σ_{jM}	0.0509	0.1374	0.2371	0.3558	0.5113	0.1155	0.1372	0.1372	0.1451	0.1841
η_j	1.3277	2.4816	3.7784	5.3164	7.1511	1.1386	1.2349	1.3135	1.3135	1.4496
a_j	0.1500	0.3040	0.4646	0.6444	0.9106	0.1437	0.2409	0.2764	0.2764	0.3459
b_j	0.1465	0.2938	0.4424	0.6289	0.9225	0.1090	0.1215	0.1654	0.1654	0.2094
σ_j^{err}	0.4186	1.4941	2.7075	4.0286	6.2960	0.6303	0.6867	0.6867	0.8556	1.0592
$X_{j,0}$	1.0009	2.1415	3.0805	4.0770	5.5087	2.9430	3.4598	3.4765	3.5784	4.4755

Prior and posterior quantiles for the derived ratios. Sample: January 2010–July 2025.

	Prior					Posterior				
Parameter	5%	25%	Median	75%	95%	5%	25%	Median	75%	95%
$b_j/(a_j+b_j)$	0.2121	0.3668	0.4882	0.6193	0.7798	0.2587	0.3261	0.3743	0.3743	0.5280
a_j/η_j	0.0345	0.0738	0.1214	0.1997	0.4280	0.1260	0.1951	0.2105	0.2177	0.2587
b_j/η_j	0.0332	0.0709	0.1188	0.1966	0.4050	0.0794	0.0944	0.1259	0.1259	0.1839
$(a_j+b_j)/\eta_j$	0.0966	0.1629	0.2488	0.3908	0.7679	0.2509	0.2895	0.3364	0.3364	0.4088

Conclusions

- We develop a structural credit-risk model with downward jumps, where the total jump intensity is affine in a stochastic physical-risk factor, capturing the effect of climate-related adverse events on firm value.
- The pricing framework is based on Laplace transforms and a multiscale perturbation approach, which reduces the original two-dimensional first-passage problem to a tractable one-dimensional approximation and delivers a semi-closed-form CDS pricing formula.
- The risk-neutral measure is introduced through a stochastic discount factor with two priced sources of risk: aggregate market risk and systematic physical risk; the physical-risk factor is modeled as a CIR process and estimated from a media-based proxy of climate concerns.
- In the empirical analysis, the model is estimated on monthly U.S. CDS spreads by Sequential Monte Carlo, combining filtering of the latent firm-value process with inference on firm-specific static parameters.

Thanks for your attention!

References

Jean-Pierre Fouque, Sebastian Jaimungal, and Matthew J Lorig. Spectral decomposition of option prices in fast mean-reverting stochastic volatility models. *SIAM Journal on Financial Mathematics*, 2(1):665–691, 2011.

S. G. Kou and H. Wang. First passage times of a jump diffusion process. *Advances in Applied Probability*, 35(2):504–531, 2003.

Grigorios A. Pavliotis and Andrew M. Stuart. *Multiscale Methods: Averaging and Homogenization*, volume 53 of *Texts in Applied Mathematics*. Springer, New York, NY, 2008. ISBN 978-0-387-73828-4. doi: 10.1007/978-0-387-73829-1.