



Silfio: Quantifying Armed-Conflict Risk

A Modelling Framework for Geopolitical Risk

Moez Abeidi, FIA

Silfium Analytics

Model version 1.1 · October 2026

Abstract

Armed conflict is a live risk on insurance and investment portfolios, but no established quantitative model measures it. In its place are country risk scores and scenario narratives, which rank exposure rather than quantify it. Silfio is a conflict risk model built on the catastrophe-modelling framework. It simulates whether and when conflict begins, how severe it becomes, where inside a country the violence falls, and the financial impact on each insured risk and investment holding. Results are expected annual loss, tail risk and exceedance-probability curves, with insurance losses reported gross and net of cover. This paper sets out how the model is built, how it is validated, and what it produces; the fitted specification and full validation diagnostics are available on request.

Contents

Executive Summary	3
1 Introduction	5
2 Modelling Scope	5
3 Architecture	6
4 Event Module: Conflict Frequency	7
4.1 News-based adjustment	8
5 Hazard Module: Conflict Severity	10
6 Exposure Module: Geographic Exposure	12
6.1 Exposure to conflict-onset hazard	12
6.2 Subthreshold modelling within exposure	14
7 Vulnerability Module: Damage if Affected	15
8 Investment Module: Financial Asset Shocks	17
9 Financial Module: Insurance Loss Waterfall	19
10 Validation	21
11 Known Limitations	23
12 Data Sources	25
13 Scenario Testing and Expert Overrides	25

Purpose and Basis of Use

This whitepaper is written for insurance, reinsurance, investment and risk professionals evaluating the Silfio framework. It sets out the methodology, validation evidence and limitations of model version 1.1. It does not constitute advice on any particular portfolio, transaction or decision.

Changes in 1.1. Version 1.1 models ongoing conflicts as well as new ones, and refines how conflict severity translates into physical damage.

Executive Summary

What Silfio does. Silfio is a catastrophe model for armed conflict, putting a loss figure on conflict risk for insurance and investment portfolios. It simulates thousands of possible years, answering four questions in each:

1. **Does conflict break out?** For every country, the model estimates the chance that a new armed conflict begins over the next year, based on its economy, political conditions and conflict history. The estimate is then updated with AI-scored evidence from news reporting, so the view keeps pace with events.
2. **How severe is it?** Each simulated conflict is given a severity that sets its death toll, the share of buildings damaged, and the shock to the economy and financial markets.
3. **Where does the damage fall?** Fighting concentrates in some parts of a country and misses others. The model places damage where conflict activity has historically clustered, so a site in a contested border region and one in a stable city are treated differently. Smaller-scale organised violence that falls short of a full conflict is modelled alongside.
4. **What does it cost?** For each property, the model calculates the damage accumulated over the course of the conflict and converts it into a financial loss, before and after insurance and reinsurance. Investment holdings in affected countries are revalued for the conflict's shocks to financial markets.


Figure 1 shows the result in a simple case: the same \$10m office placed in five capital cities. The differences in loss cost come from how likely a new conflict is to break out in each country, the conflicts already under way there, and how exposed the office's location is within it.

What it produces. For any portfolio of insured properties and investments, Silfio reports expected annual loss, the loss at return periods such as 1-in-100 and 1-in-200 years, and the full distribution of possible losses. Results are shown before and after insurance and reinsurance, from the point of view of the insured, the insurer and the reinsurer, and for property and investments separately and combined. Users can also run their own scenarios, for example making conflict in a country certain or fixing the level of damage or market shock, and measure the effect on the portfolio. Silfio models both the outbreak of new conflicts and the continuation of conflicts already under way, so exposure to current wars is included in every run.

Same Office, Five Countries

As at 12 October 2026

Identical \$10m office in each capital's centre.



Relativity 1x (dashed) is a typical property in the country. Loss cost is annual, in basis points of property value.

Subthreshold share: proportion of hits from subthreshold violence.

* Affected in fewer than 30 simulated years, so damage if affected is imprecise.

Figure 1: An identical \$10m office in five capital cities. Read left to right, the columns follow the model's steps through to the annual loss cost.

Why it can be relied on.

- Tested on years it had not seen.** For each year from 2020 to 2025, the model was rebuilt using only information available at the start of that year, and its forecasts were compared with what then happened. The 10% of countries it rated riskiest accounted for 60% of the 52 conflicts that actually began, and the news-based update improved statistical forecast accuracy in 10 of 13 test years.
- Realistic severity.** Across 59 conflicts with death tolls ranging from 25 to over 124,000, 92% of the realised tolls fell within the model's 90% range.
- Location matters.** Tested on later events, the model's conflict risk map typically placed about seven times more probability where fighting actually occurred than a uniform map would, and about three times more in countries it had never been trained on. It also reflects how individual conflicts concentrate into theatres: for conflicts it had not seen, its range contained the share of damage in the worst-hit area 88% of the time.
- Honest about uncertainty.** Rather than a single best estimate for each assumption, the model carries a range, so results reflect what is not known as well as what is.

Known limitations and planned improvements are set out in Section 11.

1 Introduction

Geopolitical risk analysis remains dominated by qualitative techniques, drawing on expert judgement, country assessments, and scenario narratives. These approaches provide essential political context, but do not by themselves yield calibrated event probabilities or portfolio loss distributions. For underwriting, capital allocation, and investment decisions, the practical requirement is to translate assessments of instability into financial impact. Catastrophe models exist for adjacent man-made perils; state-based armed conflict has not been treated in the same way.

Silfio aims to close that gap with a quantitative framework, complementing qualitative analysis with probabilistic estimates of conflict occurrence, severity, and financial loss. Its focus is armed conflict and associated physical and financial losses. By linking conflict occurrence and severity to portfolio exposures, it supports consistent estimation of expected loss, tail risk, and the effects of insurance and reinsurance.

2 Modelling Scope

Definition of conflict. The main event pathway adopts the Uppsala Conflict Data Program (UCDP) definition of *state-based armed conflict*. This requires a dispute over government or territory, armed confrontation between opposing parties with at least one state government involved, and a minimum of 25 battle-related fatalities during a calendar year [11]. The definition covers interstate and internal conflicts; it does not require a formal declaration of war. Silfio models qualifying episode onsets, including renewed episodes.

Violence outside the main onset pathway. Property damage can also arise from organised violence that does not satisfy the state-based conflict definition. Conflicts can fall outside the main onset pathway when the death threshold is not met and/or the opposing parties are non-state actors (e.g., cartel violence). Silfio therefore includes a separately modelled *subthreshold* input within the exposure module.

3 Architecture

Silfio composes six modules in a directed pipeline, mirroring the standard catastrophe-model decomposition:

$$\underbrace{\text{Event}}_{\text{conflict onsets}} \rightarrow \underbrace{\text{Hazard}}_{\text{severity shocks}} \rightarrow \begin{cases} \text{Exposure} \rightarrow \text{Vulnerability} \rightarrow \text{Financial} & (\text{insurance}) \\ \text{Investments} & (\text{investments}) \end{cases}$$

Each component’s distribution is conditioned on the realised output of its predecessor, so the pipeline is a chain of conditional distributions. This is the decomposition used by established natural-catastrophe models [16, 17], and it is adopted here for the same reasons: each component can be fitted to the data that identifies it, validated on its own terms, and revised without rebuilding the chain.

The statistical components are estimated within a Bayesian framework. The simulation propagates epistemic uncertainty in model parameters alongside aleatory variation in conflict occurrence and losses. Each modelling component exports 10,000 posterior draws, and a synthetic year uses one assembled parameter draw, conditional on which the event and loss processes are simulated. The resulting annual loss distribution incorporates both sources of uncertainty.

The predictive distribution integrates over parameter uncertainty, with a Monte Carlo approximation based on the exported draws:

$$F_Y(y \mid \mathcal{D}) = \int F_Y(y \mid \theta) \Pi(d\theta \mid \mathcal{D}) \approx \frac{1}{M} \sum_{m=1}^M F_Y(y \mid \theta^{(m)}). \quad (1)$$

Here Y is a scalar model output, y a threshold value, and $\mathcal{D}$ the fitting data. The posterior predictive cumulative distribution function $F_Y(y \mid \mathcal{D})$ gives the probability that $Y \leq y$, accounting for parameter uncertainty. The conditional distribution $F_Y(y \mid \theta)$ gives that probability for a fixed parameter vector θ , while $\Pi(d\theta \mid \mathcal{D})$ describes parameter uncertainty given the data. The approximation averages these conditional probabilities over $M = 10,000$ parameter draws $\theta^{(m)}$, indexed by m .

Components are fitted separately, on largely disjoint datasets, and combined under a product-posterior approximation: posterior dependence is retained within each component but not across them. Dependence between simulated *outcomes* is unaffected, being induced by the shared events and shocks the pipeline passes forward rather than by the parameter draws. Catastrophe models commonly represent damage uncertainty through a conditional distribution around a mean damage ratio [18, 19], with model parameters

otherwise held fixed within a given model view; propagating posterior draws for every fitted component is a stronger treatment of parameter uncertainty, of which independence across components is the limiting assumption.

4 Event Module: Conflict Frequency

Data. Conflict onsets are taken from the UCDP/PRIO Armed Conflict Dataset (v25.1), restricted to the period from 1985 onwards to avoid fitting to the idiosyncratic decolonisation era.

Model. Onset probability is modelled with a Bayesian logistic regression at the country–day level, using the established rare-events (case–control) sampling design of King and Zeng [1] with an importance-weighted likelihood. Covariates span economic indicators, conflict history, measures of political repression, and geographic region.

Conditional on onset, the conflict type (interstate vs. internal) is assigned from its historic share.

Simulation. Conflict onset is simulated on the annual level and follows a *hurdle* specification, separating occurrence from the count conditional on occurrence. The fitted daily onset probability q is conditional on the covariates and information available at the valuation date. For each posterior draw, the annual projection holds this probability constant over a 365-day model year, consistent with the model’s country–day response scale. The resulting annual occurrence probability π is

$$\pi = 1 - (1 - q)^{365}, \quad 0 \leq q \leq 1. \quad (2)$$

This projection does not assert that risk will remain unchanged as new information arrives; the treatment of within-year variation is discussed in Section 11.

The number of onsets in an affected year is estimated separately from occurrence, using a hierarchical count model conditioned on structural conflict risk. Writing Z for the occurrence indicator, K for the positive count and N for the annual count,

$$Z \sim \text{Bernoulli}(\pi), \quad N = \begin{cases} 0, & Z = 0, \\ K, & Z = 1, \end{cases} \quad \mathbb{E}[N] = \pi \mathbb{E}[K]. \quad (3)$$

Onset dates and conflict types are sampled subsequently.

Count model. A hurdle count model estimates the number of onsets conditional on at least one occurring, with the mean allowed to vary with structural risk. This avoids

mechanically tying conditional counts to annual occurrence probability.

Illustration: the same office in five countries. Figure 1 compares identical offices in five capitals. Differences in country-level onset frequency drive most of the variation in loss cost.

4.1 News-based adjustment

An effective conflict model must depart from the natural-catastrophe template in one respect: it has to keep pace with events. Hurricane frequency (after allowing for seasonal effects) does not usually change materially from one month to the next; a country’s conflict risk can change within days. A structural model built on annual covariates cannot represent that.

A separate layer therefore adjusts the annual occurrence probability for evidence observable in reporting. Articles are screened for country relevance and scored by LLMs against a set criteria of political and security risk indicators. Hundreds of thousands of scored articles, 196 countries, 2005 to the present.

The relationship between indicator evidence and subsequent onset is fitted and tested out of time, and its incremental value measured against the structural model alone (Section 10).

News-Adjusted vs Structural-Only Annual Conflict-Onset Probability

As at 12 October 2026




Figure 2: Effect of the news-based adjustment on modelled onset probability.

What these probabilities represent. The quantity plotted is the probability that a qualifying episode *begins* during the year, whether a first episode or a renewal after a break, in the UCDP sense set out in Section 4. It is not the probability that a country experiences armed conflict of any kind. A country already at war is not assigned a high probability on that account, because a continuing episode is a realised state rather than a new event. Conflicts already under way are simulated separately (Section 5), and new onsets are added on top.

5 Hazard Module: Conflict Severity

Severity converts each simulated onset into a coherent vector of physical and macro-financial shocks: 12-month battle deaths, damaged share of the built stock, GDP shock, inflation shock, equity-market shock, and sovereign bond-yield shock.

Conflict intensity is unobserved. The model represents it through a latent factor linking deaths, physical damage and economic outcomes, so that their joint behaviour is carried into portfolio losses.

Anchoring on deaths. In the model, battle deaths provide the shared severity signal for damage and macro-financial outcomes. Deaths are more consistently recorded than physical damage and more directly attributable to conflict than market or economic movements.

The structure has three stages. For conflict onset i , latent intensity is

$$\eta_i = \mathbf{x}_i^\top \boldsymbol{\gamma} + u_i, \quad (4)$$

where $\mathbf{x}_i$ contains conflict and country characteristics, $\boldsymbol{\gamma}$ their intensity coefficients, and u_i unexplained intensity variation.

Intensity then informs the death count D_i :

$$\log(1 + D_i) = a_D(\mathbf{x}_i) + \lambda_D f_D(\eta_i) + \varepsilon_{iD}, \quad (5)$$

where a_D captures the baseline and covariate effects, f_D is the intensity response, λ_D its loading, and ε_{iD} residual variation.

The remaining channels condition on standardised log deaths, z_i^D :

$$g_j(Y_{ij}) = a_j(\mathbf{x}_i) + \lambda_j z_i^D + \varepsilon_{ij}, \quad j \in \{M, G, I, E\}, \quad (6)$$

where Y_{ij} denotes positive damage (M), GDP (G), inflation (I) or equity (E), and g_j is the outcome transformation. As above, a_j , λ_j and ε_{ij} give the baseline and covariate effects, severity loading and residual variation. For damage, z_i^D is replaced by deaths measured relative to population.

Distributional choices. A folded Student- t factor represents variation in conflict intensity. Deaths are modelled as Gaussian on the log scale, positive damage as logit-normal with a separate probability of zero damage, and macro-financial outcomes with Gaussian and skew-normal mixtures. These choices accommodate rare extremes, bounded dam-

age and asymmetric market responses without imposing the same distribution on every outcome.

Country size. Country area enters the damage channel only: for a given death severity, a larger country sustains a smaller *proportional* loss of its national building stock.

Satellite damage data. A remote-sensing pipeline derives conflict-damage indicators from pre-war and post-onset Sentinel-1 synthetic-aperture-radar imagery, reconciled spatially with built-environment datasets. This yields consistent, globally available damage measurements independent of ground reporting.

Expert anchoring. Ground-truth destruction data are sparse and noisy. Satellite damage estimates carry measurement error; expert-labelled UNOSAT assessments cover selected conflict theatres rather than a representative sample of national building stock; and light damage may go undetected or be recorded only above a detection threshold, introducing under-ascertainment and censoring. Several priors are therefore deliberately informative. The GDP and inflation anchors are set from documented evidence rather than from the training sample. The GDP drag follows Collier’s estimate of an approximately 2.2 percentage-point annual growth reduction [13] together with IMF evidence placing a roughly 2% decline in real GDP per capita one year after severe conflict [14]. Inflation assumptions are informed by evidence of increased inflation during civil war [15]. The equity channel is anchored differently: published conflict-equity estimates concern the markets of countries reacting to a distant war rather than the market of the country experiencing one, so they do not transfer. Its location is instead set so that the implied median loss sits near the conflict median observed in the training sample, with the downside regime carrying the tail.

Two consequences are stated explicitly rather than left implicit. First, the likelihood contribution of positive satellite damage magnitudes is deliberately tempered, so that a small number of noisy readings cannot override the expert level and tail calibration. Second, the upper tail of the country-level damaged share is calibrated against documented severe historical conflict outcomes rather than estimated from the satellite sample, combining empirical evidence with expert judgement. Because that evidence is sparse, the tail is subject to sensitivity testing. The damage calibration draws on the death distribution, fitted to UCDP records and validated out of time (Section 10). This links the simulated extremes: draws in the top 1% of damaged share align approximately with the top 1% of death tolls. Better satellite damage data would allow this relationship to be estimated more directly, reducing reliance on expert judgement.

Ongoing conflicts. Conflicts already under way at the valuation date are simulated from the first day of the year, each carried forward separately. A model fitted to the

history of conflicts gives the chance that each one continues over the coming year and, if it does, how many battle deaths it causes. The outlook depends on how intense the conflict has been recently, whether it is escalating or subsiding, and how long it has lasted. Continuing conflicts cause physical damage in the same way as new ones. Within the country, damage is directed towards where the conflict has recently been fought, balanced against the long-run pattern of risk.

Calibration and discrimination. Out-of-time tests (Section 10) support the overall death distribution, but not the ranking of individual conflicts by severity. Conflicts with similar starting conditions can develop very differently, limiting what onset characteristics alone can predict. For portfolio analysis, the model simulates many possible outcomes, making a well-calibrated range of severities more important than predicting which particular conflict becomes severe. More simulations improve numerical precision, but do not remove model uncertainty. This limitation concerns severity; the event module does distinguish onset risk between countries.

6 Exposure Module: Geographic Exposure

The exposure module combines two distinct inputs: country-level damage from the conflict-onset hazard module, and a separately modelled subthreshold violence process. Both are resolved to location-level exposure within this module before entering vulnerability and financial-loss calculations:

$$\left. \begin{array}{l} \text{Conflict-onset hazard} \\ \text{Subthreshold violence} \end{array} \right\} \longrightarrow \text{Location-level hits} \longrightarrow \text{Vulnerability}.$$

6.1 Exposure to conflict-onset hazard

The hazard’s country-level damaged share is translated to individual portfolio locations through spatial and occupancy relativities, with stochastic geographic concentration.

Spatial relativity. A continuous spatial surface is fitted per region to historical geolocated conflict activity and geographic and built-environment data, using SPDE/INLA [2, 3]. At location s in country c ,

$$\eta_c(s) = \alpha_c + \mathbf{x}(s)^\top \boldsymbol{\beta} + \mathbf{z}(s)^\top \mathbf{u}_c + \xi(s), \quad p_c(s) = \text{logit}^{-1}(\eta_c(s)), \quad (7)$$

where $\eta_c(s)$ is the linear predictor (the log-odds of the presence/background score) at location s in country c , α_c is a country intercept, and $\mathbf{x}(s)$ contains geographic and built-

environment covariates, with regional coefficients β . The subset $\mathbf{z}(s)$ allows selected effects to vary by country through slope deviations $\mathbf{u}_c$. The fitted presence/background score $p_c(s)$ describes relative spatial patterns, rather than an absolute local event probability.

The spatial field $\xi(s)$ captures persistent local differences beyond those covariates. It has Matérn covariance and is represented through

$$(\kappa^2 - \Delta)\xi(s) = \mathcal{W}(s), \quad (8)$$

where $\mathcal{W}$ is white noise, Δ the Laplacian and $\kappa > 0$ an inverse-range parameter controlling spatial correlation. The spatial estimates are regularised to reduce overfitting and limit extreme relativities. The location relativity is normalised against a property-weighted country benchmark, so a value of one represents a typical property.

Occupancy relativity. Building-use relativities are estimated from building-level conflict damage assessments joined to built-environment attributes, comparing each building type with housing in the same neighbourhood so that the relativity reflects targeting rather than location, which the spatial layers already take into account. Informative priors encode targeting doctrine and evidence on protection under international humanitarian law, and the variation in targeting between conflicts is carried into the simulation, so each simulated conflict has its own targeting pattern. Each building type’s relative risk is thus calculated.

Damage clustering. Armed conflict occupies bounded battle-zones rather than spreading evenly across national territory, and the extent of those zones varies with terrain, resources and war aims [20]. A long-run spatial risk surface describes where conflict damage is more likely to occur, but an individual conflict need not follow that average pattern. The simulation therefore represents variation in where fighting concentrates and in the local concentration of damage within affected areas. Country-level damage is allocated across plausible conflict theatres stochastically rather than in fixed proportions:

$$\boldsymbol{\theta} \sim \text{Dirichlet}(\alpha \mathbf{p}), \quad \mathbb{E}[\boldsymbol{\theta}] = \mathbf{p}, \quad (9)$$

where $\boldsymbol{\theta}$ is the vector of theatre shares realised by an individual conflict, $\mathbf{p}$ is the long-run share of country damage attributed to each theatre, and $\alpha > 0$ is a concentration parameter: a large α reproduces the average pattern, while a small one concentrates a given conflict into few theatres. Theatre shares are estimated from geolocated conflict records and checked against historical episodes.

Within a theatre, fighting concentrates further along fronts and contested districts. The


scale of a district is derived from how far building-level damage remains correlated with separation distance in the available damage surveys; it is this scale that determines whether several nearby assets in a portfolio are affected together. See Section 11 for limitations on transferability across theatres.

Illustration: the same office within one country. Figure 3 compares identical offices at three locations in India. Country-level hazard is shared; spatial relativity drives the differences in loss cost.

Same Office, Three Places in India

As at 12 October 2026

Country-level hazard is shared, so onset frequency is identical. Location relativity does the work.



Relativity 1x (dashed) is a typical property in the country. Loss cost is annual, in basis points of property value.

* Affected in fewer than 30 simulated years, so damage if affected is imprecise.

Figure 3: The same office at three locations in India. Onset frequency is shared; the location relativity drives the difference in loss cost.

6.2 Subthreshold modelling within exposure

Subthreshold modelling supplies a separate exposure input from recorded violent events outside the main conflict-onset pathway. It is estimated and simulated separately from that pathway, but remains a component of the exposure module.

Its spatial architecture is the one described in Section 6.1: a latent Gaussian field with Matérn covariance, represented through the SPDE approach and fitted by INLA on geolocated violence records. What differs is the use made of it. In the conflict-onset pathway the surface supplies a relativity that distributes a country-level damaged share. Here it is fitted as an intensity surface, and the simulation draws discrete events from it: how many occur in a country-year, where each one falls, and how severe each one is.

Both the count and the locations therefore matter. A portfolio’s outcome depends on how close its locations fall to the events simulated in that particular year. Event severity is then translated into physical-damage exposure at those locations.

This input is evaluated even when the main event module produces no qualifying onset.

The resulting location-level hit records enter the same downstream vulnerability and financial calculations as conflict-onset exposure.

“Subthreshold” denotes the modelling pathway, rather than a universal severity cutoff.

Expert judgement in the damage calibration. How often subthreshold events occur, where they fall and how intense they are is estimated from recorded data. Translating an event into property damage is where judgement enters, for two reasons: the available severity relationship is the one fitted on qualifying conflicts, applied here to far smaller events, and there is no reliable direct measurement of the area a single event of this size damages. Two judgements close that gap. The first sets the overall damage level through expert judgement, using published civil-unrest claims and losses as an order-of-magnitude guide [21]. The second concerns how far damage may extend around an event. With limited evidence for smaller incidents, this assumption draws on observations from larger conflicts. Both assumptions are used across countries without separate local calibration. Validation against insured losses remains outstanding (Section 11).

7 Vulnerability Module: Damage if Affected

What the damage ratio represents. Vulnerability in a natural-catastrophe model is ordinarily a per-event quantity: a damage ratio conditional on a stated hazard intensity at the site, for example a hurricane or an earthquake. Conflict does not present that way. A building in a contested area is not struck once at a measurable intensity; it is exposed repeatedly over months, by different weapons and at intensities that are not observed, and what can be assessed afterwards is its accumulated state. The damage ratio modelled here is correspondingly cumulative: it is the damaged share of a property over the period during which it was affected, integrating over however many times and however heavily it was hit, rather than the consequence of a single simulated event. Conflict intensity therefore enters upstream, through how likely a property is to be affected and how severe the conflict is, rather than as an intensity argument to a damage curve.

Writing D_i for the damage ratio of property i over an exposure period $[0, T]$,

$$D_i = \min \left(\sum_{k=1}^{N_i} d_{ik}, 1 \right), \quad (10)$$

where N_i is the number of damaging impacts sustained by the property during the period and d_{ik} the incremental damage ratio of the k th of them, the sum being capped at total loss. Neither N_i nor d_{ik} is observable: assessment records the state of a building at a point in time, not the sequence that produced it. Silfio therefore models D_i directly, in

two parts:

$$D_i = H_i Y_i, \quad H_i \sim \text{Bernoulli}(p_i), \quad p_i = 1 - e^{-\Lambda_i}, \quad \Lambda_i = \int_0^T \lambda_i(t) dt, \quad (11)$$

where H_i indicates whether the property was affected at any point in the period, $\lambda_i(t)$ is its instantaneous damage-arrival rate, Λ_i the intensity accumulated over the period, and $Y_i \in (0, 1]$ the damage ratio reached by the end of the period given that the property was affected, specified in equation (12) below. The severity, location and extent of the year's conflict act on $\lambda_i(t)$; the characteristics of the building act on Y_i .

Data. Vulnerability is fitted on UNOSAT building-by-building damage assessments from recent conflicts in the Middle East. Ordinal damage classes are mapped to loss ratios using a model-specific adaptation of HAZUS damage-state guidance [5]. Remote-sensing studies document difficulty identifying light and moderate damage [4]; the model addresses this by treating the moderate class as *censored* rather than observed exactly. See Section 11 on the transferability of these curves to other conflict theatres.

Model. The accumulated ratio Y_i is estimated as a censored lognormal regression on building size, capped at total loss:

$$Y_i = \min(e^{\mu_i + \sigma \varepsilon_i}, 1), \quad \mu_i = \beta_0 + \beta_1 x_i, \quad \varepsilon_i \sim \mathcal{N}(0, 1), \quad (12)$$


where x_i is a building-size covariate and β_0 , β_1 and σ are fitted parameters. The cap at 1 produces the point mass at total loss. Priors bound the size effect to engineering-plausible ranges, and the covariate effect is applied only within the support of the fitting data. Fitting treats the moderate damage class as censored, as described above. Combined with the hit process H_i , the specification reproduces the empirical shape of conflict damage distributions structurally: a point mass at zero, a continuous distribution of partial damage, and a point mass at total loss.

Illustration: three buildings on one site. Figure 4 compares a house, warehouse and office tower at effectively the same location. Occupancy affects exposure; building size affects damage if affected.

Three Buildings on One Site

As at 12 October 2026

Within 250 m of each other in Srinagar. Building type sets occupancy; volume drives damage if affected.



Relativity 1x (dashed) is a typical property in the country. Loss cost is annual, in basis points of property value.

Figure 4: Three buildings within 250 metres in Srinagar. Building size drives the difference in damage if affected.

8 Investment Module: Financial Asset Shocks

Armed conflict carries a macro-financial signature that natural catastrophes rarely produce: output, inflation, foreign exchange, equity markets and sovereign yields move together. A portfolio's loss is not confined to damaged buildings.

Holdings in an affected country are therefore marked down by asset class, driven by the shocks the hazard module simulates. Write V_i for the market value of holding i and V'_i for its value under that scenario.

Where a country experiences several onsets in the same simulated year, each onset's shocks are expressed as a multiplier on original value and compounded across onsets.

Equities. Position-level shocks apply an industry sensitivity to the simulated country equity shock, estimated from long-horizon industry return data, so that defensive and cyclical sectors respond differently in drawdowns:

$$V'_i = V_i(1 + \beta_j S_E)_+, \quad (13)$$

where S_E is the simulated country equity shock and β_j the sensitivity of industry j . Conflict-specific beta adjustments are on the development roadmap, subject to sufficient data.

Bonds. Credit-spread responses are estimated by rating from long-run spread histories with a hierarchical heavy-tailed model, partially pooled across rating buckets. The regression determines the *shape* of the rating response; its *magnitude* is benchmarked to

the Solvency II 1-in-200 spread stress [10], providing a regulatory benchmark for the scale of tail spread losses. Bonds are then repriced with duration mechanics:

$$V'_i = V_i \exp(-D_i \Delta y_r), \quad \Delta y_r = S_B + \Delta s_r, \quad (14)$$

where Δy_r is the total change in the yield at which the holding is repriced, being the simulated sovereign bond-yield shock S_B plus the credit-spread widening Δs_r for rating r , and D_i is the effective duration to policy expiry. A model-specific default approximation, motivated by reduced-form credit-risk modelling [8], then converts spread widening into a per-onset default probability,

$$\text{PD}_{ik} \approx \frac{\Delta s_r}{\text{LGD}}, \quad \text{PD}_i = 1 - \prod_k (1 - \text{PD}_{ik}), \quad (15)$$

each clipped to $[0, 1]$, where PD_{ik} is the probability implied by onset k and PD_i the annual probability across that year's onsets. Default is absorbing: one indicator $\delta_i \sim \text{Bernoulli}(\text{PD}_i)$ is drawn per instrument and year, so a holding defaults at most once however many onsets occur. Here LGD is loss-given-default, set by instrument type at industry-standard levels and referenced to the Basel foundation IRB benchmarks [7] and to long-run rating-agency recovery evidence on defaulted senior unsecured bonds [9], with the corporate assumption set to reflect the stressed conditions in which a conflict-driven default would occur.

Real estate. Property values respond to the simulated GDP shock through a property-to-GDP elasticity, estimated on a panel of national real house-price and output histories fitted hierarchically with country-level effects. Evidence of a long-run relationship between real house prices and economic activity motivates this specification [6]:

$$V'_i = V_i (1 + S_G)^\theta, \quad (16)$$

where S_G is the simulated GDP shock and θ the property-to-GDP elasticity.

Currency and cash. The model calculates and outputs in USD. Therefore, non-USD holdings receive a purchasing-power-parity-motivated currency adjustment driven by the simulated inflation shock relative to baseline.

How to read these figures. These channels are calibrated on historical market data observed during and around conflict episodes, and that data is noisy: prices move for many reasons at once, and conflict coincides with rate moves, commodity shocks, policy responses and shifts in global risk sentiment. A conflict-period drawdown is not cleanly separable from the other forces acting on the same prices, so the fitted responses carry some of that information with them. Silfio accordingly reports investment results as

losses under a conflict-consistent market scenario, rather than as an isolated causal effect of conflict on asset prices.

9 Financial Module: Insurance Loss Waterfall

The financial module converts physical damage into losses net of (re)insurance for each location, simulated year, and policy layer. Nothing in it is specific to Silfio: it applies established contract mechanics and standard risk measures, and is set out here so that the terms applied to simulated losses can be checked against a policy wording.

Direct insurance. Ground-up loss at location i is its value times the simulated damage ratio, $X_i = V_i D_i$, with $D_i = H_i Y_i$ as in Section 7. Coinsurance (average) conditions are taken into account where a policy carries them, reducing the loss presented to the layer in proportion to any underinsurance. The layer then attaches at a , under either a standard or a franchise deductible, and is capped at the occurrence limit L_o , which is shared by all locations the policy covers in the same occurrence O :

$$E_i = \begin{cases} (X_i - a)_+, & \text{standard deductible,} \\ X_i \mathbf{1}\{X_i \geq a\}, & \text{franchise deductible,} \end{cases} \quad (17)$$

$$G_i = E_i \min\left(1, \frac{L_o}{\sum_{j \in O} E_j}\right) \mathbf{1}\{\text{on risk}\}, \quad I_i = p G_i, \quad C_i = X_i - G_i, \quad (18)$$

where G_i is the gross layer loss, I_i the insurer's share of it at participation p , and C_i the loss retained by the insured, who recovers the full layer loss from the subscribing markets. The indicator $\mathbf{1}\{\text{on risk}\}$ applies the losses-occurring-during basis, with occurrences dated at conflict onset.

Reinsurance. Quota share cedes a fixed proportion q of the insurer's subscribed share,

$$R_i = q I_i. \quad (19)$$

Excess of loss instead attaches to a subject loss U : the insurer's share accumulated over the unit the treaty basis calls for (a risk, an occurrence or the year). With attachment A , limit L_t and r reinstatements, the recovery on unit U is

$$R_U = q \min(W, ((r+1)L_t - W_{\text{prior}})_+), \quad W = \min((U - A)_+, L_t), \quad (20)$$

where q is now the reinsurers' share of the layer, W the loss falling in the layer and W_{prior} the limit already consumed by earlier units in the same simulated year, so that cover is exhausted in onset order. Reinstatement premium is charged pro rata to the

limit reinstated, up to r limits. Recoveries and reinstatement premium are allocated back to the ceding layers in proportion to each layer's contribution to the subject loss, $R_i = R_U I_i / U$, so the four views reconcile:

$$\underbrace{X_i}_{\text{ground-up}} = \underbrace{C_i}_{\text{insured net}} + G_i, \quad I_i = \underbrace{(I_i - R_i + P_i)}_{\text{insurer net}} + \underbrace{(R_i - P_i)}_{\text{reinsurer net}}, \quad (21)$$

with P_i the allocated reinstatement premium, paid by the insurer to the reinsurer.

Outputs include expected loss, value at risk (VaR) and tail value at risk (TVaR) on both aggregate and occurrence bases, simulated maximum loss, and per-draw loss distributions for the property book, the investment book, and their joint book.

The measures below describe annual loss L for a specified book and financial perspective, in a common reporting currency. The distribution includes zero-loss years and satisfies $\mathbb{E}[|L|] < \infty$.

Expected loss and exceedance probability. The expected annual loss and annual aggregate exceedance probability (AEP) at threshold ℓ are

$$\text{EL} = \mathbb{E}[L], \quad \text{EP}(\ell) = \mathbb{P}(L > \ell) = 1 - F_L(\ell), \quad (22)$$

where $F_L(\ell) = \mathbb{P}(L \leq \ell)$ is the cumulative loss distribution. Losses are retained at occurrence level, so both exceedance bases are reported. On the aggregate basis L is the sum of a simulated year's occurrences; on the occurrence basis (OEP) it is the largest single occurrence in that year, being one country's onset on one date. Value at risk and tail value at risk below are computed on each basis.

Value at risk. At confidence level α , VaR is the generalised inverse of the cumulative loss distribution:

$$\text{VaR}_\alpha(L) = \inf\{\ell \in \mathbb{R} : F_L(\ell) \geq \alpha\}, \quad 0 < \alpha < 1. \quad (23)$$

Tail value at risk. TVaR is the mean of the upper $1 - \alpha$ probability mass, defined by integrating the loss quantile function:

$$\text{TVaR}_\alpha(L) = \frac{1}{1 - \alpha} \int_\alpha^1 \text{VaR}_u(L) du, \quad (24)$$

Because most simulated years carry no loss, for low-frequency exposures the tail can include zero-loss years; TVaR is then lower than the average loss in a year that has one.

A T -year return-period loss is the loss exceeded with annual probability $1/T$, that is $\text{VaR}_{1-1/T}(L)$, on either the aggregate or the occurrence basis: the 100-year loss is the

99% VaR.

10 Validation

The validation programme operates at several levels.

Statistical validation. Checks cover convergence for sampling-based models and numerical stability for approximate methods. Predictive checks assess how well the models reproduce observed patterns, while held-out data test calibration on observations excluded from fitting.

Internal consistency. End-to-end checks verify that linked modules produce consistent results, simulations reproduce specified inputs across the full parameter range, and losses and probabilities stay within valid bounds.

Rolling out-of-time backtest, 2020–2025. The frequency and severity models are validated on a strict rolling-origin protocol: for each holdout year, every time-dependent model is *refit* using only data through the prior year-end, scored on the covariate snapshot available on 1 January of the holdout year, a full year is simulated with the production engine, and the results are compared with realised UCDP-recorded outcomes. No information from the holdout year (or later) enters fitting or scoring.

Holdout year	Onsets	AUC	Onsets in top decile	Count pctl.
2020	10	0.92	7/10	0.27
2021	9	0.90	4/9	0.32
2022	8	0.90	5/8	0.31
2023	10	0.91	7/10	0.47
2024	6	0.89	2/6	0.06
2025	10	0.90	6/10	0.50

Across 53 onsets in six holdout years, mean area under the ROC curve (AUC) was 0.90 and no year fell below 0.89. Log loss improved on the climatological base-rate forecast in every year. The top risk decile captured 58% of onsets, against 10% under an uninformative ranking.

These results support discrimination between countries. Backtests are extended with each new UCDP release.

Incremental validation of the news adjustment. The table above reports the structural probabilities alone. The news-adjustment layer is evaluated separately over

2013–2025 on expanding-window out-of-time folds, each estimated only on data preceding its evaluation year and compared against the structural model. The adjustment improved the log score in 10 of 13 folds, for a pooled gain of 11.9 nats (95% confidence interval 2.7 to 21.3, $p = 0.011$; the interval is bootstrapped over countries rather than folds, since country-days within a country are not independent). Mean AUC rose from 0.893 to 0.900. Scored on the same 1 January snapshots as the table above, with both the structural model and the adjustment refit through the prior year-end, the adjusted probabilities placed 31 of 52 onsets (60%) in the top risk decile. The effect is asymmetric by construction: the median country-day is revised slightly downwards, while the upper tail of the adjustment distribution carries the gain.

The results above demonstrate the tested news component’s incremental value.

Conditional severity, 2020–2025. Severity is tested on the same rolling-origin protocol, but conditioned on onset: for each conflict that began in a holdout year, the severity model is refit through the prior year-end and its predictive distribution of 12-month battle deaths is scored against the realised toll. Across 59 onsets with realised tolls ranging from 25 to over 124,000 deaths — the mean probability integral transform was 0.495 against an ideal of 0.5, and 92% of outcomes fell inside the 90% predictive interval. On the 28 onsets of 2023–2025, 82% fell inside the 80% interval and 93% inside the 90% interval. Median absolute error is a factor of about two on a response spanning four orders of magnitude, and no cap is imposed on the upper tail.

Spatial exposure. Validation focuses on the Middle East and North Africa, Africa and Asia, which contain the most conflict-event data. Fitted through 2022 and tested on 36,606 events in 2023–2025, the current model achieves a mean country-normalised log-density gain of 1.925 nats per event over uniform land allocation, showing that it identifies where subsequent conflict activity concentrates within countries. A separate test withholding entire countries achieves a gain of 1.129 nats across 303,201 events, supporting transfer to countries without local training history.

Damage concentration. The model of how damage concentrates within a country was fitted on conflicts that began before 2017 and tested on the 88 that began later. The 90% predictive interval for the share of first-year damage falling in the most affected area contained the realised share in 88% of those conflicts, against 23% if damage were allocated in fixed proportion to the long-run risk surface. At district scale, leaving each satellite damage survey out in turn, the realised district damage rate fell inside its 90% predictive interval for 76–100% of districts, against 5–36% if every building in a theatre shared the theatre’s damage rate.

Credit spreads. For each year from 2020 to 2025, the model was fitted on earlier

data and tested on monthly spread changes by credit rating in the held-out year. Its predictive log score exceeded that of a rating-based historical benchmark in all six test years.

Additional violence pathway. Forward holdouts test where smaller violent events occur, how often they occur and their fatality totals. The spatial field improves predictive log score by 1.175 nats per event over a covariates-only benchmark. Incorporating recent activity reduces event-count prediction error by 43% across 1,479 country-years; a separate test reduces error in log-transformed fatality totals by 35%. A separate spatial-memory test, limited to MENA, finds that blending recent activity with longer-term patterns improves predictive log score by 0.086 nats per event across 1,920 held-out events. These tests assess recorded violence; validation of the resulting property losses remains outstanding.

Full model validation results for every module, including backtests and diagnostic output, are reported in the Silfio methodology document.

Model governance. Silfio operates under documented model governance covering version control, model change, validation, testing and the management of known limitations. Governance and validation documentation is maintained separately from this paper.

11 Known Limitations

Silfio maintains a documented limitation register; the principal entries are:

- **Current view of onset risk.** Annual onset probabilities are a snapshot at the valuation date, converted from a daily hazard that is held constant over the year. Risk can rise or fall within the year as events develop, and that path is not simulated; results reflect the view at the time of the run. The snapshot itself ranks countries well out of time: annual backtests beat the base-rate benchmark in every holdout year, with mean AUC 0.90 (Section 10).
- **Country dependence.** Onsets are simulated independently across countries, given each country's probability and the shared parameter draw. In reality conflict risk is correlated between countries, so the chance of losses in several countries in the same year may be understated. The news-adjusted probabilities partly capture shared developments at the valuation date, which mitigates this.
- **Damage and vulnerability.** Property damage is inferred from conflict and satellite evidence rather than a large, consistent claims database. Building construction,

occupancy and site-specific protective measures can therefore produce outcomes that differ from the modelled average.

- **Conflict intensity in the vulnerability sample.** Vulnerability is fitted on recent Middle Eastern conflicts marked by prolonged urban fighting and heavy bombardment (Section 7). Damage if affected in theatres fought at lower intensity, with different weaponry or against a different building stock may differ from the fitted curves.
- **Location and transfer.** Location relativities describe patterns in recorded conflict activity. They may be less reliable for strategically important sites, new theatres or locations outside the model’s historical coverage. Expert judgement should be exercised here.
- **Tail uncertainty.** Severe outcomes are rare and the upper tail is consequently uncertain. Tail loss measures should be stress-tested for the portfolio.
- **Subthreshold violence.** The pathway does not cover all lower-intensity political violence, and small events may go unreported. Events causing material property damage are more likely to be reported, which is expected to reduce the effect of under-reporting on loss estimates. Translating recorded activity into property damage also involves judgement.

Where the model improves most

Many of the entries above share a cause: the model has to infer property damage from evidence that was not collected to measure it. The investment channels face a similar gap, drawing on market histories that are deepest outside the countries where conflict is most likely. Three sources of data would address most of them directly.

Insurance loss data. Located claims from political-violence and property portfolios would measure damage to the insured building stock itself, rather than inferring it from conflict records and satellite imagery. They would improve every stage of the property pathway. For spatial exposure, the location relativity is fitted to where violence has been recorded (Section 6.1); claims show where insured property was actually damaged, which would calibrate how conflict intensity translates into hits at individual locations, and the occupancy relativities with them. For vulnerability, claims would extend the damage curves beyond the conflicts in the fitting sample, and the construction and occupancy attributes insurers already hold would allow construction type to be identified. For subthreshold violence, they would replace the judgement that currently sets the level of damage with a direct calibration. Finally, they would allow the model to be validated against observed insured losses, which it currently is not.

Better satellite damage data. Wider and more sensitive damage assessment would improve the physical side of the model. Consistent assessment across more conflicts, including light and moderate damage, would allow the upper tail of the country-level damaged share to be estimated rather than set by a calibration target, reduce the reliance on censoring to correct for under-detection, broaden vulnerability beyond the few conflicts it is fitted on, and test whether damage clustering at district scale transfers between them.

Financial market data across countries. Industry and security-level return data across many countries and markets, including those that have experienced conflict, would improve the investment module. Equity shocks are currently applied as a country shock scaled by industry sensitivities estimated from long-horizon industry return data (Section 8). Broader data would allow those sensitivities to be estimated by industry and by country, capturing how the composition and depth of each market shape its response, and would provide the evidence needed for the conflict-specific beta adjustments on the development roadmap. The same breadth would sharpen the credit-spread, property and currency responses in markets where the current calibration relies on pooled or proxy histories.

12 Data Sources

Silfio is fitted on public and licensed data spanning five families: conflict event and fatality records; satellite damage assessments and earth-observation imagery; built-environment and settlement data; country macro-financial series; and long-run market and credit series. The records anchoring the frequency, severity and vulnerability modules are the UCDP/PRIO Armed Conflict Dataset [12] and UNOSAT damage assessments, as cited above. The full source inventory, with vintages and licence terms, forms part of the NDA-tier documentation. All training data are either public or licensed; no client data are used in model fitting. Input datasets are subject to documented quality, consistency and reconciliation checks before model fitting, and known limitations and material sources of bias are considered in calibration and validation.

13 Scenario Testing and Expert Overrides

Silfio supports scenario analysis through user-defined overrides of selected model estimates and simulated outcomes. Overrides let users assess specified stress events, test sensitivity to alternative assumptions, quantify impact from ongoing conflicts, or substi-

tute their own expert assessment where they disagree with the model. The scenario is propagated through the portfolio calculations to quantify its effect on expected loss, tail risk and insurance recoveries.

Supported overrides. At country level, users can specify annual conflict-onset probability and conflict type. They can also fix selected hazard outcomes: the damaged share of built stock, GDP shock, inflation shock, equity-market shock, and sovereign bond-yield change. Occurrence and severity can be varied separately or together. Setting occurrence probability to 100% places the first onset on day 1 of the simulated year. Overrides currently apply at country level; selecting individual conflict theatres, so that a scenario can concentrate fighting in a chosen region, is on the development roadmap rather than in the current release.

Scenario interpretation. Supplied hazard values replace the corresponding simulated outcomes for the selected country; outcomes left unspecified keep their simulated values. An override does not condition the other hazard outcomes: fixing the damaged share, for example, leaves the simulated GDP and inflation shocks unchanged. A hazard override also does not by itself force an onset. Downstream modules then run as normal on the resulting scenario.

Comparing a scenario with the baseline on the same portfolio and contract inputs isolates the sensitivity of modelled losses to the overridden assumptions. Scenario results should be read alongside the overrides that produced them: they quantify the consequences of the user's view or stress hypothesis and do not revise the baseline calibration.

Technical Review

Silfium Analytics welcomes technical challenge and independent review of the Silfio methodology. Actuaries, quantitative analysts, statisticians, catastrophe-modelling specialists and other risk professionals with relevant experience who are interested in reviewing the methodology or validation framework are invited to contact the author.

About

Silfio is developed by Silfium Analytics, a consultancy specialising in actuarial modelling and technology. Enquiries regarding model access, portfolio analysis, integration into existing workflows, or bespoke work are welcome.

Moez Abeidi, FIA — Silfium Analytics

moez.abeidi@silfumanalytics.com

silfo.com · silfumanalytics.com

Professional designations identify the author's qualifications and do not imply endorsement of this publication by the Institute and Faculty of Actuaries.

References

- [1] King, G. and Zeng, L. (2001). Logistic regression in rare events data. *Political Analysis*, 9(2), 137–163.
- [2] Lindgren, F., Rue, H. and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *JRSS-B*, 73(4), 423–498.
- [3] Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *JRSS-B*, 71(2), 319–392.
- [4] Boccardo, P. and Giulio Tonolo, F. (2012). Haiti earthquake damage assessment: review of the remote sensing role. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XXXIX-B4, 529–532. [doi:10.5194/isprsarchives-XXXIX-B4-529-2012](https://doi.org/10.5194/isprsarchives-XXXIX-B4-529-2012).
- [5] FEMA (2010). *HAZUS-MH MR5 Multi-hazard Loss Estimation Methodology*.
- [6] Adams, Z. and Füss, R. (2010). Macroeconomic determinants of international housing markets. *Journal of Housing Economics*, 19(1), 38–50.
- [7] Basel Committee on Banking Supervision. *CRE32: IRB approach — risk components for each asset class*, paragraphs 32.6–32.7. Bank for International Settlements. [Basel Framework](#).
- [8] Jarrow, R. and Turnbull, S. (1995). Pricing derivatives on financial securities subject to credit risk. *Journal of Finance*, 50(1), 53–85.
- [9] Moody’s Investors Service. *Default and Recovery Rates of Corporate Bond Issuers, 1920–2004*. Special Comment, January 2005.
- [10] European Commission (2015). Commission Delegated Regulation (EU) 2015/35, Article 176: spread-risk sub-module (Solvency II).
- [11] Uppsala Conflict Data Program. *UCDP Definitions*. Uppsala University. [Definitions of state-based armed conflict and related concepts](#).
- [12] Pettersson, T. (2025). *UCDP/PRIO Armed Conflict Dataset Codebook v25.1*. Uppsala Conflict Data Program. [Dataset codebook](#).
- [13] Collier, P. (1999). On the economic consequences of civil war. *Oxford Economic Papers*, 51(1), 168–183.

- [14] International Monetary Fund (2024). Fragile Foundations: The Lasting Economic Scars of Conflict. *Regional Economic Outlook: Middle East and Central Asia*, Chapter 2.
- [15] Adam, C., Collier, P. and Davies, V.A.B. (2008). Postconflict monetary reconstruction. *The World Bank Economic Review*, 22(1), 87–112.
- [16] Grossi, P. and Kunreuther, H. (eds.) (2005). *Catastrophe Modeling: A New Approach to Managing Risk*. Huebner International Series on Risk, Insurance and Economic Security, vol. 25. Springer.
- [17] Mitchell-Wallace, K., Jones, M., Hillier, J. and Foote, M. (eds.) (2017). *Natural Catastrophe Risk Management and Modelling: A Practitioner’s Guide*. Wiley-Blackwell.
- [18] Oasis Loss Modelling Framework. *Modelling Methodology*. [Oasis LMF documentation](#): secondary uncertainty is “an uncertainty distribution around the mean damage ratio at each hazard point”.
- [19] Carter, J. (2016). *Uncertainty*. Oasis Loss Modelling Framework Ltd. [Technical note](#).
- [20] Buhaug, H. and Gates, S. (2002). The geography of civil war. *Journal of Peace Research*, 39(4), 417–433.
- [21] France Assureurs (2023). *Le coût des sinistres déclarés à la suite des violences urbaines de fin juin est réévalué à 730 millions d’euros*. 6 September. [Claims and loss estimates](#).