Supporting information for "Climate-Conditioned Catastrophe Modeling for Dynamic Risk Assessment"

Francesco Comola[1], Ian Bolliger[2], Yann Boulben Meyer[1], Bernhard Märtl[1], Jamie Rodney[2], and Hilary Paul[1]

[1]*LGT ILS Partners Ltd, Zurich, Switzerland*
[2]*Reask, London, United Kingdom*

This supporting material provides additional information on:

- The predictive skills of seasonal landfall forecasts (Section S1).
- The composition of the ILW portfolio and construction of ILW price index (Section S2).
- The investment performance based on April hurricane forecasts (Section S3).
- The sensitivity analysis to the risk-aversion parameter (Section S4).
- The sensitivity analysis to the starting year and investment duration (Section S5).
- The sensitivity analysis to historical-sequence bootstrapping (Section S6).
- A workflow diagram of the methodology implemented in this study (Section S7).
- The climate-conditioned hurricane generative model (Section S8).
- The event count adjustment implemented in the CBRA tool (Section S9).
- The conditioning and convergence analysis on the loss covariance matrix (Section S10).

S1 Results: Predictive skill of probabilistic seasonal landfall forecasts

This study assesses the benefit of incorporating probabilistic seasonal hurricane landfall forecasts into investment decisions. This application requires a model that exhibits skill in predicting physical hurricane outcomes, such as landfalls, before considering the noisier relationship between hurricanes and insured losses. The Reask Unified Tropical Cyclone model has previously been evaluated when forced by reanalysis fields and large-ensemble climate-model output Loridan and Bruneau [2024], and in near-real-time tropical-cyclone forecasting applications Loridan et al. [2026]. A comprehensive validation of its seasonal forecast application is currently in preparation. Here, we provide a focused evaluation of the model's ability to represent interannual variability in U.S. landfall counts, which is a necessary, but not sufficient, condition for skillfully representing interannual variability in insured losses.

To assess skill in predicting landfall counts, we adopt a likelihood ratio approach. Our test asks the question, "how many times more or less likely would the observed pattern have been under a data-generating process of seasonally-varying distributions than it would have been assuming a static distribution?" We evaluate the relative likelihood for landfalls of Tropical Storm strength and greater, Category 1 strength and greater, and Category 3 strength and greater. Simultaneously, we group landfalls across the entire U.S. into seven distinct regions, or use the original set of 72 landfall gates used in the main text.

To construct a static distribution, we aggregate seasonally-forced UTC output across many years. We assess the performance of both the SEAS5-forced forward-looking forecast, which contains a larger event set size (51,000 stochastic simulations per year) but also contains forecast uncertainty over each season's meteorological and oceanographic conditions, and the ERA5-forced event sets, which are forced by a reanalysis of these conditions but contain only 2,500 simulations per season. For the static model, to compute the likelihood of observed landfalls in a given year, we leave that year out of the event set to simulate applying a static model view to a new, unseen season.

Across all metrics analyzed, and using both the SEAS5- and ERA5-forced event sets for the seasonal landfall forecast, the observed record was more likely to have been produced by the model that allows seasonal variation than by the static model (Figure S1). The likelihood ratio ranges from approximately 3 to 800,000, depending on the choice of evaluation metric and definition of the static and dynamic landfall rates. The null hypothesis of equal expected log scores is rejected for most evaluated configurations at conventional significance levels, with one-sided p-values ranging from below 0.001 (high confidence) to 0.37 (low confidence). These results indicate that the improvement in expected log score under the seasonal model is statistically distinguishable from the static model for most evaluated configurations. The seasonal model's excess performance increases with the granularity of landfall aggregation, suggesting that its skill comes not only from variation in the number of landfalls but also from their spatial distribution. Finally, the April initialization appears to perform roughly as well as the June initialization, with some metrics indicating marginally better or worse performance. This is consistent with our findings that the impact on relative investment returns between the April and June initializations is minimal.

S2 Methods: Insurance-linked Portfolio

The ILW risk premia are sourced from two reinsurance broker reports as of January 2024 (markers in Figure S2a). The average quotes (Table S1 and dashed lines in Figure S2a) are used as market prices for 2024 in our analysis.

Nationwide		Florida		Gulf		Northeast	
Strike	Price	Strike	Price	Strike	Price	Strike	Price
25	39.5%	10	39.4%	10	28.8%	15	10.8%
30	35.5%	15	31.8%	15	22.0%	20	8.0%
35	31.9%	20	27.3%	20	17.0%	25	7.1%
40	25.6%	25	22.9%	25	13.4%	30	6.2%
50	18.8%	30	19.4%	30	10.8%	35	5.8%
60	15.0%	35	16.3%	35	9.0%	40	5.2%
70	12.5%	40	13.7%	40	8.2%	50	4.0%
80	10.7%	50	10.8%	50	6.7%		
90	9.4%	60	8.2%				
100	8.0%	70	7.1%				
		80	6.1%				

Table S1: **ILW portfolio.** List of strike levels, risk premia as of 2024, and geographic scope of the 36 ILWs in the investment portfolio. Northeast includes coastal states from Virginia to Maine. Gulf includes coastal states from Texas to Alabama. Nationwide refers to the contiguous United States.

To account for the historical hardening and softening phases of reinsurance risk premia, we multiply the ILW prices $\mathbf{p}$ as of 2024 (Figure S2a) by an annual price index I_y, such that $\mathbf{p}_y = \mathbf{p}_{2024}\, I_y$ for any year y in the period 1985–2024. This index reflects the market dynamics of supply and demand for reinsurance capital, leading to price increases in the aftermath of destructive events (e.g., 1992 Hurricane Andrew, 2005 Hurricane Katrina, and 2022 Hurricane Ian) and price decreases in periods of abundant reinsurance capital.

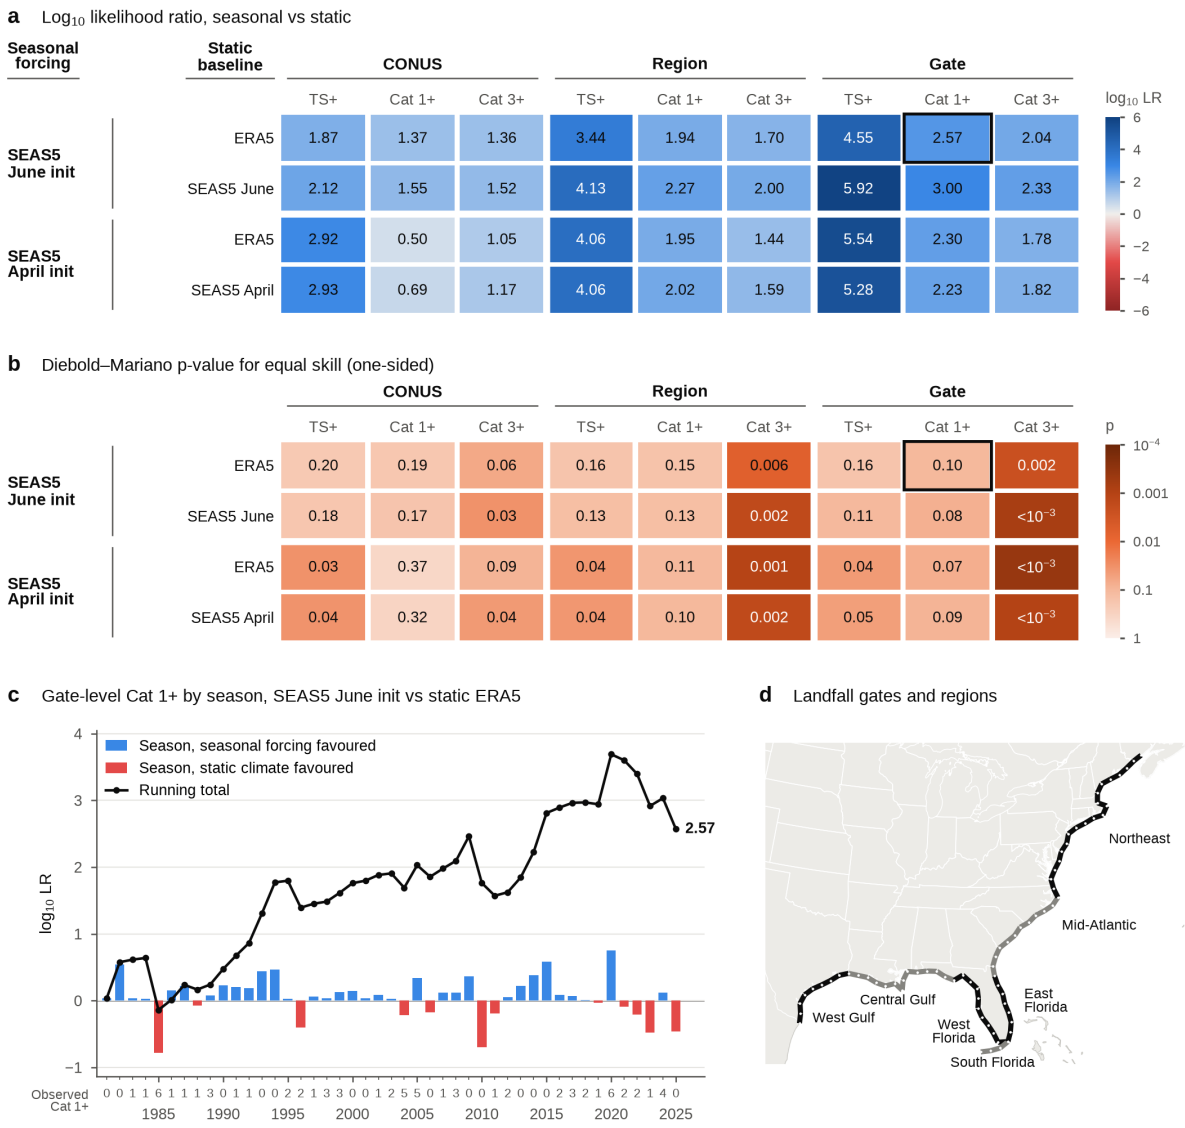


Figure S1: **Seasonal forecasts of US landfall improve on static climatology.** The figure compares the likelihood of the observed record of CONUS landfalls from 1981 to 2025 under seasonally varying and static landfall distributions. One is the seasonally forced UTC event sets using both the June and April initializations. The other is a static UTC event set created by pooling over all 1981–2025 seasons except the season being scored, using either the reanalysis-forced event sets or the same seasonal initialization. **(a)** Total log_{10}-likelihood ratio over all 45 seasons. Positive values (blue) indicate the observed record for the given metric is more likely under the dynamic empirical distributions of the seasonally forced model. Intensity groups are tropical storm and stronger (TS+), Cat 1+ and Cat 3+. Ratios are shown for annual counts aggregated to all of CONUS, to 7 coastal regions, and to 72 landfall gates, each shown in (d). **(b)** One-sided Diebold–Mariano p-value for equal expected log score per season, using a Newey–West variance with a 3-season lag and the Harvey–Leybourne–Newbold small-sample correction (t distribution, 44 degrees of freedom). **(c)** Season-by-season log_{10}-likelihood ratio (bars) and its running total (line) for the outlined cell in (a) and (b). Numbers below the axis are observed Cat 1+ CONUS landfalls in each season. **(d)** The 72 landfall gates along the CONUS coastline, grouped into the seven regions used for the regional aggregation; adjacent regions alternate in tone.

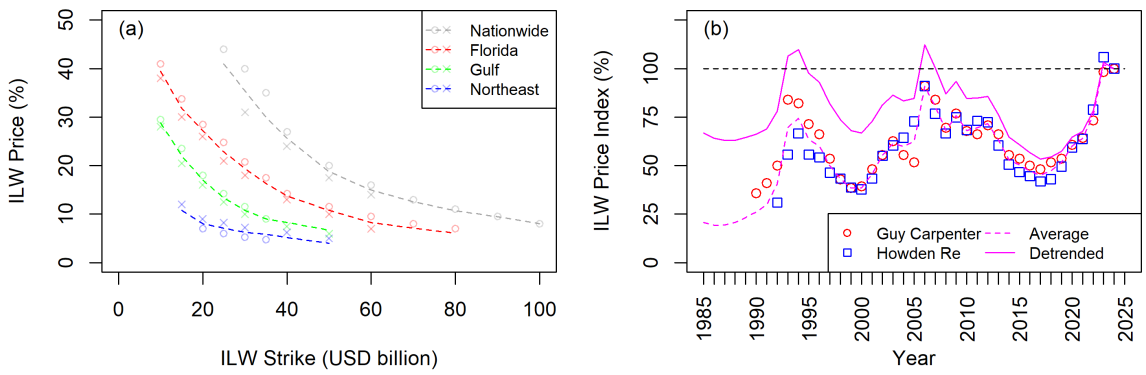


Figure S2: **ILW Prices.** (a) ILW risk premia as of 2024 quoted by two reinsurance brokers (Aon and Howden, circles) and average values used as market prices (dashed lines). (b) ILW Price Index used in our investment analysis (solid magenta line) obtained from the US Property Catastrophe Rate-on-Line Index published by reinsurance brokers Guy Carpenter (1990–2024, red circles) and Howden Re (1992–2024, blue squares). The RoL indices are normalized to their respective 2024 values, averaged (dashed magenta line), and detrended.

To construct the price index for our analysis, we utilize the US Property Catastrophe Rate-on-Line (RoL) Index published by leading reinsurance brokers Guy Carpenter and Howden Re (red and blue markers in Figure S2b). The RoL index is a measure of the change in dollars paid for coverage year on year on a consistent reinsurance program base, reflecting the pricing impact of a growing exposure base, changes in buying habits, as well as changes in market conditions.

We normalize the RoL indices to their 2024 values, average them from 1990 to 2024, and fill the values from 1985 to 1989 using proprietary reinsurance submission data (dashed magenta line in Figure S2b). We then detrend the index (solid magenta line in Figure S2b) to remove the impact of growing exposure base from the signal. This last step allows us to account for reinsurance market phases while ensuring that ILW prices are commensurate with catastrophe model results based on the 2024 industry exposure.

S3 Results: April Hurricane Forecasts

Section 2 of the main manuscript shows the benefit of a climate-aware investment strategy based on seasonal hurricane forecasts released in June. Here, we show the results of the investment strategy initialized with the April forecasts. These forecasts provide information at a longer lead time than the June forecasts and can therefore support earlier investment decisions. The timing element can be a critical advantage in reinsurance applications, given that a large number of hurricane-exposed reinsurance programs renew in June and July.

The results obtained using April forecasts are broadly consistent with those obtained using June forecasts (Figures 3 and 4 of the main manuscript). In both cases, the climate-conditioned strategy outperforms the long-term benchmark across a wide range of risk-aversion values during the 1985–2024 evaluation period. Differences between the April and June initializations are generally modest and depend on the specific performance metric considered. The April forecasts provide information at a longer lead time and can therefore support earlier investment decisions.

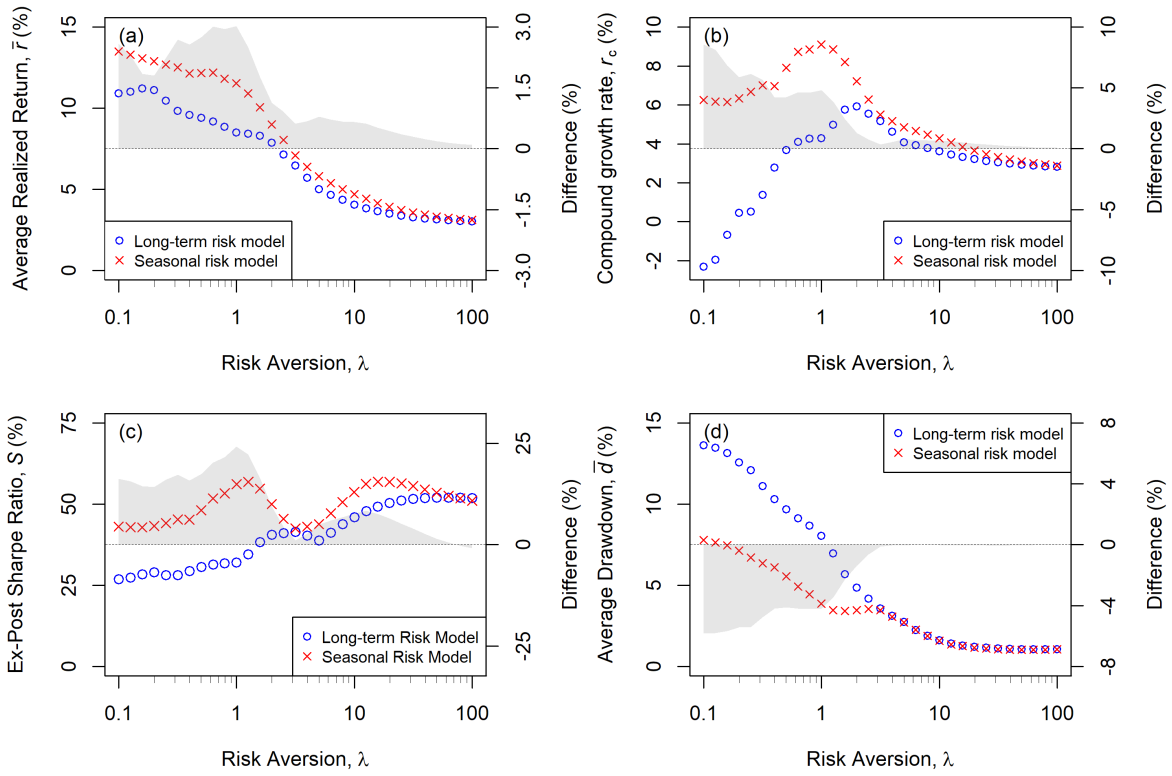


Figure S3: **Investment performance with April Forecasts.** (a) Average realized returns, (b) compound growth rates, (c) Sharpe ratios, and (d) average drawdowns generated by the climate-aware investment strategy forced with April forecasts (red crosses) and the benchmark (blue circles) in the historical period 1985–2024. Shaded gray areas indicate the difference between the climate-aware strategy and the benchmark (vertical axis on the right side of each plot).

S4 Results: Sensitivity to Risk-aversion Parameter

The main manuscript shows the capital gains generated in the 40-year investment period (1985–2024) when adopting an aggressive ($\lambda = 1$) and a conservative ($\lambda = 10$) strategy. Here, we show results for a broader range of values of the risk-aversion parameter λ, from 0.5 to 20.

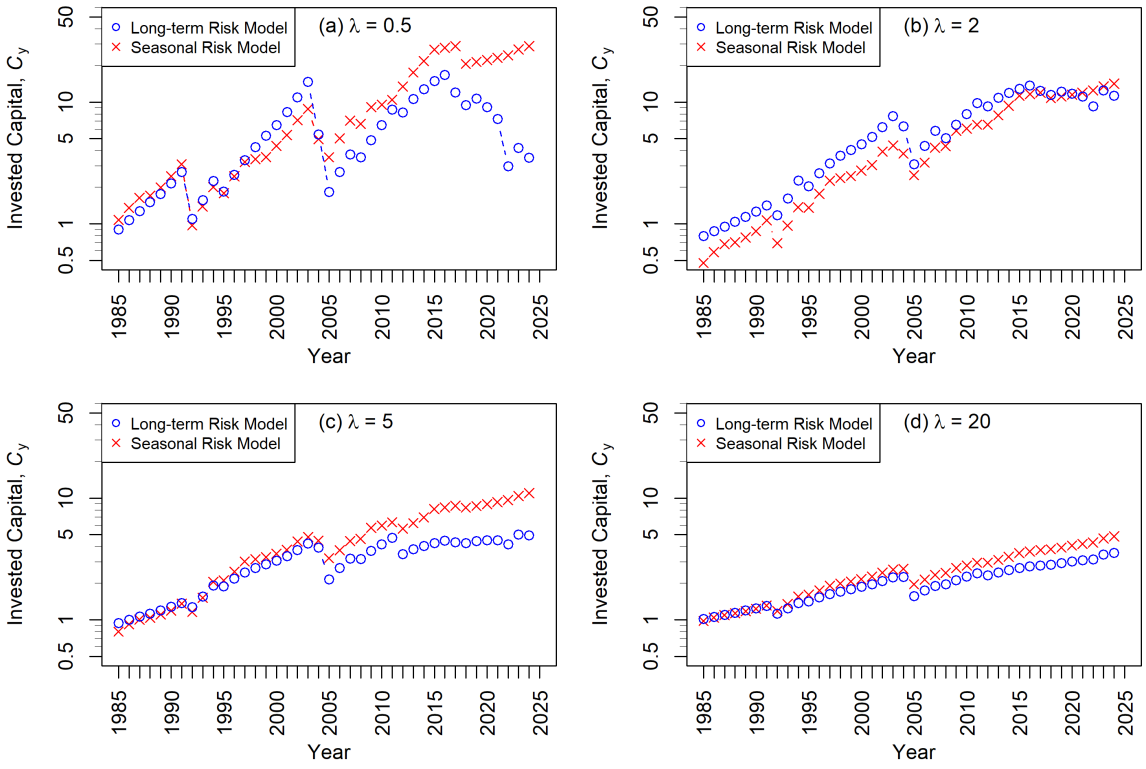


Figure S4: **Sensitivity of capital Gains to risk-aversion.** Evolution of invested capital (equation 8 in the main manuscript) during the 1985–2024 period for different values of the risk-aversion parameter λ. Red crosses represent the results obtained with the seasonal risk model, blue circles those from the long-term risk model.

The capital generated by the seasonal strategy over the 40-year evaluation period exceeds that generated by the long-term strategy for all values of risk aversion examined (Figure S4). Although the seasonal strategy performs less well in some individual years, including 1985 and 1992 for $\lambda = 2$ (Figure S4c), it compensates for these periods of relative underperformance over the longer evaluation horizon.

S5 Results: Sensitivity to Investment Horizon and Start Date

The primary analysis evaluates investment performance over the full 1985–2024 period. To assess the sensitivity of the results to the selected evaluation period, we repeated the analysis for multiple investment durations and starting years.

Investment durations ranging from 5 to 40 years were considered. For each duration, all possible starting years within the historical record were evaluated. The compound annual growth rate was calculated separately for the seasonal and long-term strategies, and the difference between the two was recorded.

Figure S5 shows the resulting differences in compound annual growth rate as a function of investment

duration and starting year. Positive values indicate superior performance of the seasonal strategy.

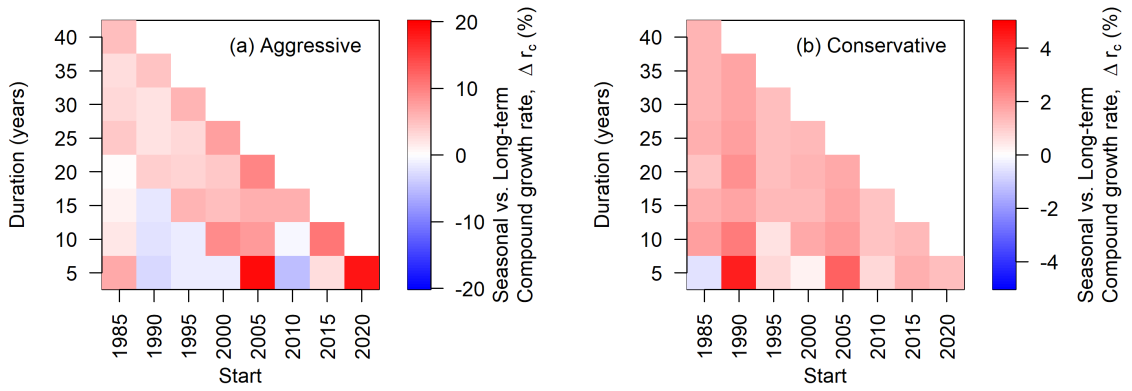


Figure S5: **Sensitivity of investment performance to starting year and investment duration.** Difference in compound annual growth rate (Δr_c) between the seasonal and long-term strategies for (a) the aggressive profile ($\lambda = 1$) and (b) the conservative profile ($\lambda = 10$). Each grid cell corresponds to a unique combination of investment starting year and investment duration. Positive values indicate superior performance of the seasonal strategy. White areas indicate combinations of starting year and investment duration that extend beyond the available 1985–2024 evaluation period.

The seasonal strategy outperforms the long-term benchmark in 78% of the evaluated investment windows for the aggressive profile and 97% of the evaluated windows for the conservative profile. Negative differences occur primarily for shorter investment horizons, where the influence of individual hurricane seasons on compound performance is greater. As investment duration increases, the relative performance becomes progressively more stable and positive.

These results indicate that the advantage of the climate-conditioned strategy is not limited to the specific 1985–2024 evaluation period and persists across a broad range of alternative start dates and investment horizons.

S6 Results: Historical-sequence Bootstrap

The historical evaluation period from 1985 to 2024 contains only a limited number of seasons with severe insured hurricane losses. Consequently, the relative performance of the seasonal and long-term strategies could be disproportionately influenced by the particular composition and frequency of high-loss years in the observed 40-year record. We therefore performed a non-parametric bootstrap analysis to evaluate the sensitivity of the compound investment performance to the years included in the historical sample.

For each bootstrap realization, we constructed an alternative 40-year investment sequence by sampling years from the 1985 to 2024 period independently and with replacement. Each sampled year retained the paired annual returns generated by the seasonal and long-term strategies. Sampling the two strategy outcomes jointly preserves the comparison under the same realized hurricane-loss conditions. Because sampling was performed with replacement, individual years could occur multiple times or not at all within a given bootstrap realization. Some realizations therefore place greater weight on high-loss years such as 1992, 2005, and 2017, whereas others contain fewer such years.

For each bootstrap realization, we calculated the compound annual growth rate of each strategy using equation 5 of the main manuscript. The same sequence of sampled years was used for both strategies in each realization. We generated 5,000 bootstrap realizations for both the aggressive strategy, with $\lambda = 1$, and the conservative strategy, with $\lambda = 10$.

Figure S6 shows the resulting distributions of compound annual growth rates. The seasonal strategy produces a higher compound annual growth rate than the long-term strategy in 95% (aggressive profile) and 99% (conservative profile) of the bootstrap realizations, as shown by the predominantly positive distribution of the paired difference $\Delta r_c = r_{c,seasonal} - r_{c,long-term}$. These results indicate that the relative performance advantage of the seasonal strategy is not attributable solely to the exact composition and frequency of loss years in the observed 40-year sample.

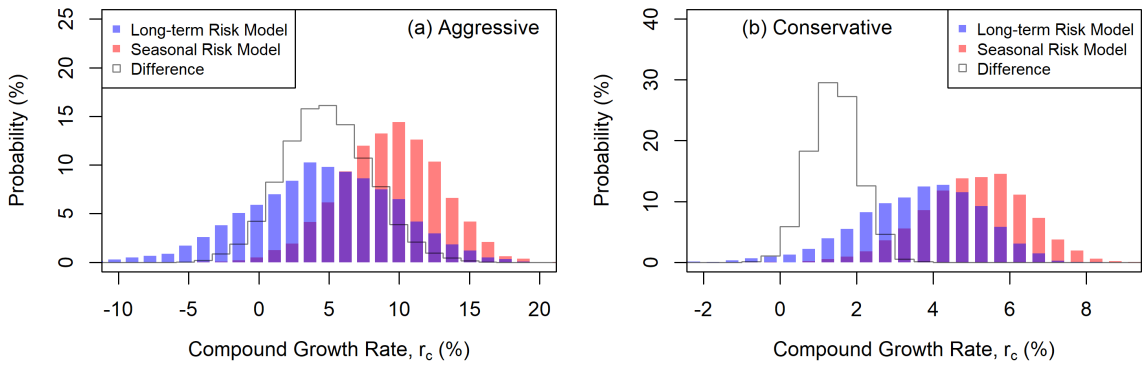


Figure S6: **Historical-sequence bootstrap of investment performance.** Probability distributions of compound annual growth rates obtained by resampling the paired annual investment outcomes from the 1985 to 2024 historical evaluation period with replacement. Results are shown for (a) the aggressive risk profile, with risk-aversion $\lambda = 1$, and (b) the conservative risk profile, with risk-aversion $\lambda = 10$. Red bars represent the seasonal strategy and blue bars represent the long-term strategy. The black curve shows the distribution of the paired difference between the compound annual growth rates of the seasonal and long-term strategies. Positive differences indicate higher compound growth under the seasonal strategy.

This analysis resamples years from the observed historical period and therefore does not generate previously unobserved hurricane seasons or loss events beyond the historical record. It should consequently be interpreted as a test of sensitivity to the composition of the available historical sample, rather than as evidence that the 40-year record provides a complete characterization of long-term tail risk.

S7 Methods: Workflow Diagram

Figure S7 illustrates how different data sources and software are leveraged to produce a seasonal-adjusted view of risk and inform a climate-aware investment strategy. Each step of the workflow is described in detail in the Section 4 of the main manuscript.

S8 Methods: Dynamic Hurricane Simulations

The UTC is trained by first extracting key climate modes through principal component analysis (PCA) of gridded fields – such as sea surface temperature, wind shear, mean sea level pressure, and steering flow – and selecting the "climate connectors" that correlate meaningfully with TC activity. The ERA5 reanalysis is used to define the gridded fields, while the International Best Tracks Archive for Climate Stewardship (IBTrACS) is used to define basin-specific TC activity. These connectors feed into hierarchical Bayesian (Poisson) models per basin to simulate the annual number of cyclones. Genesis location and timing are adjusted using empirical static climatologies modulated by local SST and wind shear anomalies, while event trajectories and intensity evolution are handled via quantile regression forests trained using the same reanalysis and best-track data to estimate conditional distributions of hourly track displacements and intensity changes.

Once trained, the UTC model proceeds with multi-stage synthetic generation: for each climate state, it

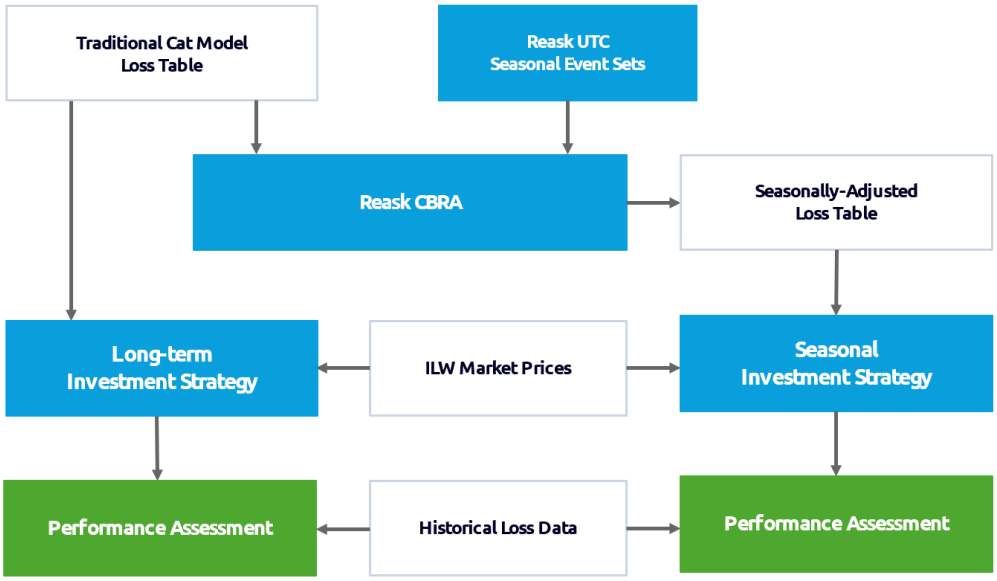


Figure S7: **Schematic representation of the methodology.** The diagram illustrates the workflow described in Section 4 of the main manuscript. White boxes represent data sources, blue boxes represent software, and green boxes represent the results shown in Figures 3 and 4 of the main manuscript.

probabilistically samples annual TC counts by basin, then for each event, it sequentially simulates genesis (location, date, intensity), tracks, intensification or weakening, and lysis until termination, finally translating central pressures into maximum sustained wind speeds. By aggregating stochastically simulated hurricane seasons, one can generate a large ensemble "event set" that reflects the impacts of both climate variability and intrinsic stochasticity on tropical cyclogenesis.

S9 Methods: Climate-Based Risk Adjustment

After calculating the event count adjustment factors by landfall gate and intensity category, CBRA takes additional steps to ensure that correlations between event types and scaling of total landfalls are preserved.

To preserve correlation between event types, when additional events need to be added for a particular event type, CBRA first calculates the mean landfall rate for all other event types across all simulated seasons in which the target event type occurs. Next, for all simulated seasons in which the target type does *not* occur, CBRA selects the one for which the landfall log-likelihood is maximized across all event types, assuming a Poisson likelihood function with parameters defined by the previously calculated means.

To ensure appropriate scaling of total landfalls at a given gate, CBRA rescales the total count of events at each gate hierarchically across several groups of Saffir-Simpson categories. After scaling the counts of events within each category, and each gate, separately, CBRA then rescales the total number of events in the major (Cat 3+) and minor (Cat 1-2) hurricane intensity groups. Finally, the counts of events are rescaled at the all hurricane (Cat 1+) level. This stepwise adjustment addresses a case where the baseline UTC event set differs in intensity distribution from the input event set, which can yield an inconsistent scaling at the "all hurricane" grouping.

S10 Methods: Covariance Conditioning and Sampling Stability

The portfolio optimization depends on the covariance matrix of losses across the 36 ILWs. We therefore evaluated both the numerical conditioning of the covariance matrices and their sampling stability with respect to the length of the stochastic catalog.

We assessed numerical conditioning using the condition number

$$\kappa(\Sigma) = \frac{\lambda_{max}(\Sigma)}{\lambda_{min}(\Sigma)}, \tag{S1}$$

where λ_{max} and λ_{min} denote the largest and smallest eigenvalues of the covariance matrix Σ, respectively. The resulting values (Figure S8a) indicate that the matrices are sufficiently well conditioned for the numerical optimization and that matrix conditioning is unlikely to be a material source of instability in the calculated portfolio weights.

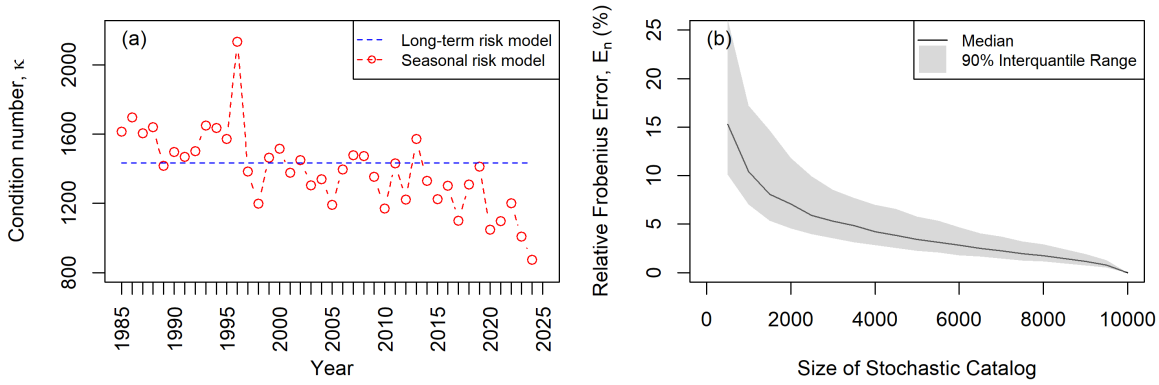


Figure S8: **Numerical conditioning and sampling stability of the covariance matrix.** (a) Condition number κ (equation S1) of the covariance matrices used in the portfolio optimization for the long-term and seasonal investment strategies. (b) Relative Frobenius error (equation S2) between covariance matrices estimated from stochastic catalogs of increasing size and the covariance matrix estimated from the full 10,000-year catalog. The solid line represents the median error across 500 random subsets for each catalog size and the shaded region indicates the 5th–95th percentile range.

We additionally evaluated the sampling stability of the covariance estimates as a function of stochastic catalog size. For each catalog size n, we generated 500 random subsets of the full 10,000-year catastrophe-model catalog and calculated the covariance matrix Σ_n. Each subset estimate was compared with the covariance matrix calculated from the full catalog Σ, using the relative Frobenius error:

$$E_n = \frac{\|\Sigma_n - \Sigma\|_F}{\|\Sigma\|_F}. \tag{S2}$$

The median relative error decreases progressively as the catalog size increases, falling below 5% at approximately 3,000 stochastic years and below 2% at approximately 8,000 stochastic years. The 5th–95th percentile interval across the 500 random subsets also narrows with increasing catalog size. These results indicate that, conditional on the catastrophe-model catalog used in this study, the estimated covariance structure stabilizes as the number of stochastic years increases and is not primarily driven by sampling variability associated with a small number of simulated extreme-loss events.

This convergence test assesses sampling stability within the modeled stochastic catalog. It does not eliminate model uncertainty concerning the representation of the extreme-loss tail or demonstrate that alternative catastrophe models would produce the same covariance structure.

References

Thomas Loridan and Nicolas Bruneau. Reask UTC: a machine learning modeling framework to generate climate connected tropical cyclone event sets globally. *EGUsphere*, pages 1–37, November 2024. doi: 10.5194/egusphere-2024-3253. URL https://egusphere.copernicus.org/preprints/2024/egusphere-2024-3253/. Publisher: Copernicus GmbH.

T. Loridan, N. Bruneau, B. Mani, C. Sampson, and J. Knaff. Local-scale wind forecasting for tropical cyclone early warnings. *Bulletin of the American Meteorological Society*, -1 (aop), August 2026. ISSN 0003-0007, 1520-0477. doi: 10.1175/BAMS-D-25-0182.1. URL https://journals.ametsoc.org/view/journals/bams/aop/BAMS-D-25-0182.1/BAMS-D-25-0182.1.xml.