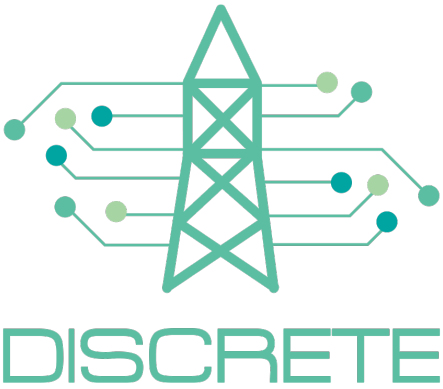

D3.3. Validation and testing of the developed models on relevant benchmarks

Authors: Kaan Yurtseven, Hakan Ergun

This deliverable is based on the publication:

Yurtseven, K., Ergun, H., & Van Hertem, D. (2026). Risk Management Under Continuous Non-Gaussian Uncertainty: Coordinating Congestion Management and Balancing in Hybrid AC/DC Grids, Under review

Table of Contents

Table of Contents 2

1. Purpose within the project 3
2. Methodological contribution 3
3. Case study results 4
4. Expanded interpretation of the complete result set 7

1. Purpose within the project

This paper is the most complete decision support layer of the project. It combines congestion management, balancing risk, N-1 security, hybrid AC/DC grid operation, and continuous non-Gaussian uncertainty in one optimization framework. The starting point is that redispatch and balancing are often optimized separately even though they rely on the same generation and grid resources. A redispatch action can reduce congestion but also limit the resources available for balancing. Conversely, a balancing action can resolve an imbalance but aggravate congestion. The paper proposes a framework that coordinates these two processes from a system level perspective.

The framework is a balancing-aware risk-based SCSOPF. It uses PCE to represent continuous non-Gaussian operational uncertainty, a contingency set to represent N-1 topological uncertainty, and a moment-based CVaR surrogate to penalize balancing risk. The model jointly optimizes generator redispatch, balancing activation, RES curtailment, and HVDC converter set-points. This directly responds to the project objective of developing probabilistic tools that are not merely academic formulations but decision support methods for transmission system operation.

Aspect	Summary
Main method	Risk based balancing aware SCSOPF with PCE and SOC representable chance and CVaR surrogate terms.
Decision timeline	Day-ahead redispatch contract volumes and a real-time look-up table conditioned on uncertainty and contingencies.
Risk metric	Balancing risk represented through a moment based CVaR surrogate in coefficient space.
Validation	5-bus studies for uncertainty and price effects, plus a Belgian transmission grid study using ENTSO-E and Elia data.

2. Methodological contribution

Figure 1 shows the operational timeline and decision structure. The framework is executed after day-ahead market clearing and before real-time operation. It takes deterministic market outcomes and uncertainty descriptions as inputs. It then produces two types of outputs: day-ahead redispatch contract volumes and a look-up table for real-time operation. Generator redispatch is modeled as a preventive here-and-now decision, while balancing activation and RES curtailment are preventive wait-and-see decisions. HVDC converter set-points are corrective wait-and-see decisions. This nested structure is important because it combines the operational uncertainty dimension with the contingency dimension rather than treating them as separate problems.

The objective combines three cost components. The first is expected redispatch cost, calculated from upward and downward redispatch relative to the market reference dispatch. The second is balancing cost risk, expressed by a CVaR based term applied to cost weighted balancing generation. The third is the expected dispatch cost of balancing generators, which prevents unnecessary curtailment and preserves economic consistency. A risk aversion parameter beta controls the trade-off between expected dispatch cost and balancing risk. The model remains tractable because the CVaR component is represented through a second order cone approximation using the first two moments of the PCE coefficients.

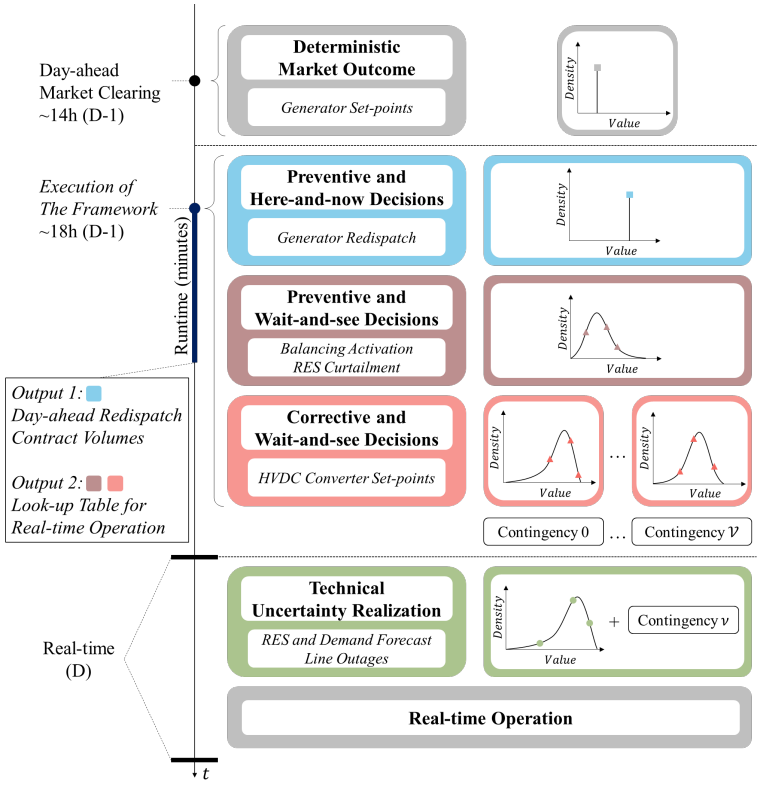

Figure 1: Timeline and decision structure of the balancing aware risk management framework

3. Case study results

Case 1 studies how redispatch adapts to different renewable uncertainty levels and risk aversion on the 5-bus hybrid AC/DC system. The system includes five AC buses, five generators, two wind farms, six AC branches, three DC buses, three HVDC branches, and converters. N-1 security is enforced against all AC and DC branch contingencies. The imbalance-to-redispatch price ratio is fixed at gamma equal to 10. Figure 2 shows two linked results. The left panel shows that as beta increases, the total balancing cost distribution becomes narrower. This confirms that the framework reduces cost volatility by accepting more conservative preventive redispatch decisions. The right panel shows the trade-off between redispatch cost and balancing risk for high, medium, low, and very low RES uncertainty. As beta increases, redispatch cost rises and balancing risk falls. The widest trade-off appears at medium renewable levels, where the system has enough renewable uncertainty to benefit from risk management but is not yet constrained by the upper installed capacity bound.

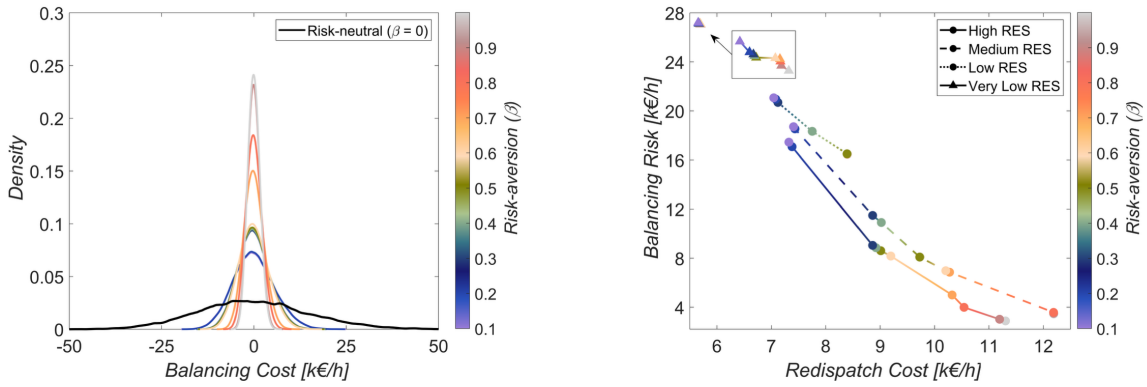

Figure 2: Balancing cost distributions (a) and redispatch cost versus balancing risk trade-off (b)

Case 2 investigates how the imbalance-to-redispatch price ratio changes the optimal risk aversion level. As illustrated in Figure 3, when gamma equals 2.5, balancing actions are only moderately more expensive than redispatch, and the minimum total cost occurs at beta equal to 0.4. When gamma equals 6, the minimum shifts to beta equal to 0.8. At very high gamma values, strong risk aversion becomes increasingly beneficial because real-time balancing becomes much more expensive than preventive redispatch. The relationship is not identical across RES conditions. The low RES case reaches beta equal to 1 at lower gamma values, because limited renewable flexibility leads to more abrupt policy changes. The medium RES case evolves more smoothly, while the high RES case requires larger gamma before high risk aversion becomes cost-effective because larger redispatch volumes are required.

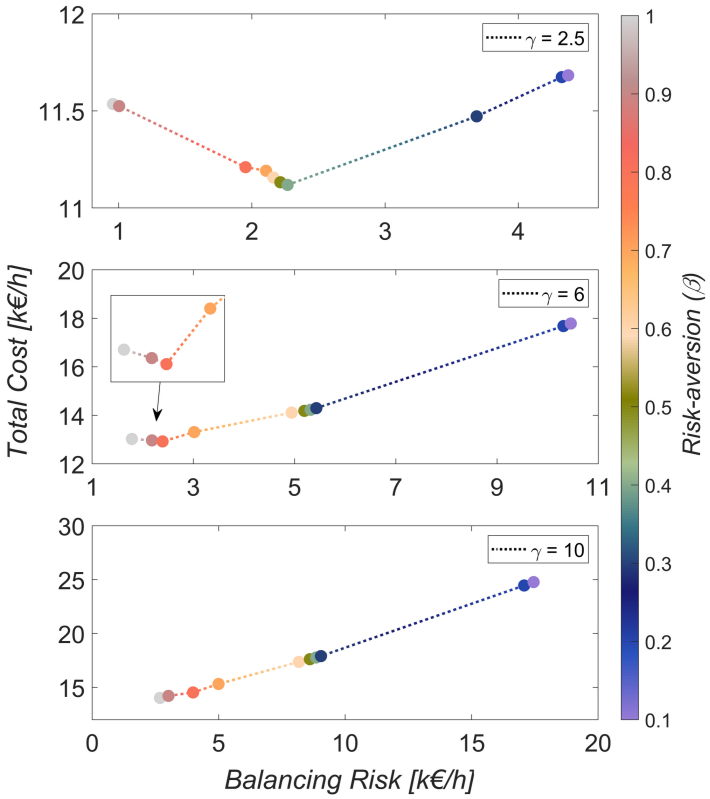

Figure 3: Trade-offs between balancing risk and total cost for varying imbalance-to-redispatch price ratios

The paper also translates this result into a practical guideline, as given in Figure 4. For instance, if gamma is expected to lie between 3 and 4, selecting beta around 0.4 to 0.5 yields near-optimal performance across the studied uncertainty conditions. This is important because it shows how the framework can be used as a decision support tool rather than simply as a solver. It allows the operator to choose a robust risk profile when exact future uncertainty and price conditions are not known.

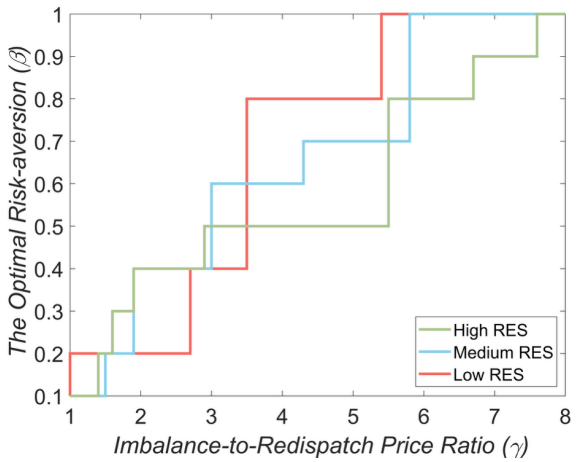

Figure 4: Relationship between the optimal risk-aversion and the imbalance-to- redispatch price ratio under different RES generation uncertainties

Case 3 applies the framework to the Belgian transmission grid. The methodology integrates Elia RES and load data, Elia imbalance prices, ENTSO-E day-ahead prices, and the Belgian TYNDP grid. For each month, 30 representative timestamps are selected with k-medoids. Forecast error samples are used to fit Beta distributions for PV, onshore wind, offshore wind, and load. The Belgian grid contingency list includes the fifteen most heavily loaded branches identified from annual operational data, and the grid includes HVDC interconnectors to the United Kingdom and Germany. Each optimization for a given beta value requires between 30 seconds and 2 minutes, which supports practical tractability.

Figure 5 summarizes the Belgian grid results. Monthly cost savings are obtained throughout the year under the optimal risk aversion level. The maximum saving is 16.8% in June, the minimum is 5.7% in February, and the annual average is approximately 12%. The correlation analysis explains what drives these savings. The Spearman correlation between gamma and total cost savings is 0.85, making the imbalance-to-redispatch price ratio the dominant driver. After separating the analysis into low, medium, and high price ratio regimes, the fitted Beta distribution characteristics also reveal clear patterns. Skewness has a consistently negative correlation with savings because right-skewed RES distributions imply that low renewable output occurs more often, which reduces the ability to economically reduce the right tail of conventional generation. Expected renewable output is positively correlated with savings because higher expected renewable generation reduces conventional day-ahead commitments and increases upward redispatch headroom. Variance is positively related to savings for PV and onshore wind, while the relationship is weaker for offshore wind because offshore wind already has high normalized output levels.

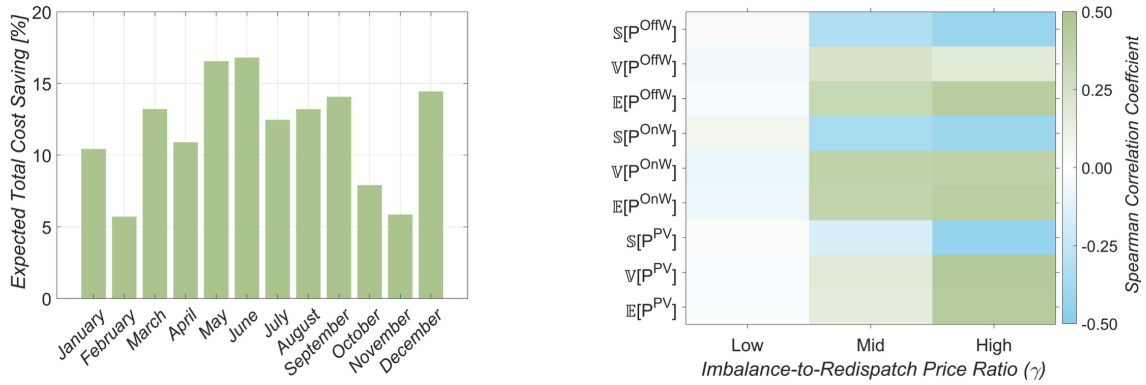

Figure 5: Belgian grid monthly savings and correlation analysis of uncertainty features

In addition, it is possible to quantify CO_2 emission benefits from the model. Although CO_2 emissions are not explicit decision variables, the optimum trades two thermal pools whose physical emissions are derivable from the publicly available ENTSO-E TYNDP 2024 data set: a redispatch pool (G_re), and a balancing pool (G_bal). For the Belgian transmission grid, G_bal comprises seven natural-gas units (4.6 GW, fleet-weighted emission factor EF_bal = 0.202 tCO_2/MWh, derived directly from the IPCC default factor for natural gas and consistent with the per-MWh CO_2 add-on embedded in the optimisation cost coefficients), while G_re comprises eight gas units, three biomass units and one oil unit (6.5 GW, EF_re = 0.176 tCO_2/MWh, lower than G_bal because of the inclusion of zero-emission biomass capacity). For every monthly representative cluster c, the optimum redispatch and balancing decisions yield a change in redispatch cost $\Delta C_{re,c}$ and a change in the balancing-cost CVaR $\Delta C_{bal,c}$ relative to the risk-neutral case ($\beta = 0$). These are converted to physical energy quantities using the cluster-specific imbalance-to-redispatch price ratio γ_c, the pool-averaged marginal generation costs, and a redispatch mark-up. Because the balancing pool is homogeneous in fuel type, the linearity of expectation and standard deviation under a scalar transformation implies that $CVaR_\alpha[CO_2\_bal] = EF_{bal} \cdot CVaR_\alpha[P\_bal]$, so the reduction in balancing-emission risk maps directly from the reduction in balancing-cost CVaR through the factor $EF_{bal}/(\gamma_c \cdot mc_{bal})$. Aggregation over the 360 representative clusters with their k-medoid probabilities and the per-month operating hours yields an annual figure. Two interpretations of the balancing pool's EF are reported: a fleet-average EF of 0.202 tCO_2/MWh and a tail-peaker EF of 0.370 tCO_2/MWh, reflecting that real-time balancing activations in the upper-quantile tail of the cost distribution physically correspond to fast-start OCGT peakers in Belgian operating practice.

Applied to the open-source Belgian TYNDP-2030 transmission grid, the framework reduces the $CVaR_{\alpha=0.9}$ of balancing emissions by approximately 750 kt CO_2/year under the fleet-average gas assumption and up to 1.4 Mt CO_2/year under the tail-peaker interpretation. Because the framework simultaneously increases preventive redispatch volume (adding about 550 kt CO_2/year), the net annual reduction in CVaR-of-CO_2 lies between 200 and 800 kt CO_2/year, with a central estimate

of approximately 400 kt CO_2/year. To place this in context, this corresponds to roughly 3 to 8% of Belgium's projected 2030 power-sector emissions, or the annual emissions of about 90,000 average passenger cars. The reduction is achieved entirely as a co-benefit of the cost and risk-driven optimisation, without any explicit emissions constraint.

While the present quantification is specific to the Belgian test case, the underlying mechanism (coordinated risk-aware substitution of high-emission tail-balancing actions by lower-emission preventive redispatch) applies to every ENTSO-E member state whose balancing fleet is dominated by fast-start thermal units, suggesting that broader adoption of such frameworks could deliver a meaningful, if country-dependent, contribution to European power-sector decarbonisation objectives.

4. Expanded interpretation of the complete result set

The numerical studies show that the framework intentionally shifts part of the operational effort from real-time balancing toward day-ahead redispatch when doing so reduces total expected cost and risk. This is an important conceptual difference from a standard redispatch minimization tool. In a purely redispatch-oriented model, higher preventive redispatch cost is undesirable. In the proposed framework, a moderate increase in redispatch can be economically rational if it reduces the exposure to expensive and volatile balancing activation. The risk-aversion parameter beta therefore controls a system-level trade-off rather than a local redispatch cost minimization.

Case 1 demonstrates how this trade-off depends on the structure of RES uncertainty. Under very low RES availability, the possible balancing deviations are limited because there is little renewable generation to deviate from the forecast. Consequently, there is limited room for risk-aware redispatch to generate savings. Under medium RES conditions, the distribution has enough spread and enough feasible generation headroom for redispatch to reduce balancing cost volatility, which creates the broadest cost-risk trade-off. Under high RES conditions, the bounded nature of the Beta distribution becomes important: although expected renewable output is high, the feasible operating region is more constrained by installed-capacity limits and grid limits. This reduces the marginal benefit of additional risk aversion.

Case 2 shows that the imbalance-to-redispatch price ratio gamma is the dominant economic driver of the optimal risk-aversion level. When gamma is low, balancing actions are not much more expensive than redispatch, so aggressive preventive redispatch is not justified. When gamma increases, avoiding balancing volatility becomes more valuable, and the optimal beta shifts upward. The reported values illustrate this clearly: at gamma equal to 2.5, the minimum total cost occurs around beta equal to 0.4, while at gamma equal to 6 it shifts to approximately 0.8. For very high gamma, the model tends toward strongly risk-averse behavior. The practical rule emerging from the study is that gamma values around 3 to 4 support beta values around 0.4 to 0.5 across several uncertainty conditions, which provides a useful operational starting point.

Case 3 is the most directly relevant to a project final report because it demonstrates the framework on an open-source Belgian transmission grid model. The study combines Belgian renewable and load data, day-ahead prices, imbalance prices, fitted monthly uncertainty distributions, and a Belgian grid representation with HVDC interconnectors. To keep the annual study tractable, each month is represented by 30 representative timestamps selected through clustering. The framework then evaluates optimal risk aversion across the year. Monthly savings are positive in every month, with a maximum of 16.8% in June, a minimum of 5.7% in February, and an annual average of approximately 12%. These numbers show that the framework is not only a theoretical formulation but a candidate decision-support tool with meaningful economic impact.

The correlation analysis explains why savings vary across months. The Spearman correlation between gamma and total cost savings is 0.85, confirming that the price ratio is the strongest driver. After separating the analysis by gamma regime, the fitted Beta-distribution features explain additional variation. Positive skewness reduces savings because it implies that low renewable production levels occur more frequently and high production levels are rarer, which leaves less opportunity to reduce expensive upward balancing exposure through redispatch. Higher expected renewable generation increases savings potential because it lowers day-ahead conventional commitments and creates more upward redispatch headroom. Variance is also generally helpful because larger uncertainty creates more room for risk-aware reallocation, although the effect is weaker for offshore wind in the Belgian data because offshore production is already high in normalized terms. These results make the framework valuable not only for operation, but also for interpreting which uncertainty characteristics make risk-aware congestion management economically attractive.