Multi-Fidelity Risk-Averse Optimisation for Multi-Energy System Planning under Correlated Uncertainty

Stijn Cousin[a], Ander Martinez Alonso[a], Alex Felice[a] and Thierry Coosemans[a]

[a] Electric Vehicle and Energy Research Group (EVERGi), Department of Electrical Engineering and Energy Technology (ETEC), Flanders Alliance for Climate Technology (FACT), Vrije Universiteit Brussel (VUB), Brussels, Belgium
Corresponding Author: stijn.cousin@vub.be

Abstract:
Investing in on-site Multi-Energy Systems (MES)—integrating photovoltaics, heat pumps, batteries, and hydrogen storage—offers a promising pathway to decarbonise the energy supply of small-to-medium enterprises (SMEs). However, these capital-intensive systems must be designed under conditions that are uncertain and statistically correlated: cold winters coincide with low solar yield, elevated heat demand, and high energy prices, while the upcoming EU ETS2 framework directly links on-site gas consumption to volatile carbon allowance costs. Together, these dependencies drive compound tail risks that independence-based models systematically underestimate.
This paper aims to develop a computationally tractable framework for risk-averse MES design that explicitly accounts for these correlated uncertainties. To this end, we formulate a multi-objective optimisation problem that jointly minimises annualised capital expenditure (CAPEX) and the Conditional Value at Risk (CVaR) of annual operating expenditure (OPEX). A Multi-Fidelity Polynomial Chaos Expansion (MF-PCE) surrogate is constructed to approximate OPEX over the joint space of design variables and uncertain parameters. The surrogate is trained using a structured set of low- and high-fidelity evaluations of a Mixed-Integer Linear Programming (MILP) energy system model. To correctly represent the uncertain operational environment, a six-dimensional Student-t copula is used, and the PCE basis is orthogonalised with respect to this joint distribution. Surrogate accuracy is assessed via cross-validation, and critical tail scenarios are rigorously verified using high-fidelity model evaluations.
The framework is applied to a light manufacturing SME case study. Preliminary results demonstrate that assuming independent uncertainties significantly underestimates $\mathrm{CVaR}_{0.95}$ and leads to distorted design decisions, particularly the under-sizing of hydrogen storage. The analysis reveals that hydrogen storage builds essential resilience against uncertainty, acting with heat pump capacity as a structural hedge during compound cold, low-solar, and high-carbon scenarios. Ultimately, the proposed approach enables computationally efficient, risk-aware MES design, providing actionable insights for robust investment planning.

Keywords:
Conditional Value at Risk; Correlated Uncertainty; Hydrogen Storage; Multi-Energy Systems; Multi-Fidelity Optimisation; Polynomial Chaos Expansion; Robust Design.

1. Introduction

Achieving industrial decarbonisation targets increasingly requires small-to-medium enterprises (SMEs) to replace fossil-fuel heating and grid-dependent electricity with on-site multi-energy systems (MES). By coupling technologies such as photovoltaics, battery energy storage systems (BESS), hydrogen electrolysers, fuel cells, and heat pumps, these systems offer immediate economic value. In the Belgian context, where (distribution system operator) DSO capacity charges penalise monthly peak imports, such systems provide a direct mechanism to reduce grid dependence and operational costs.

However, the value proposition of MES is evolving from a pathway for cost optimisation into a critical hedge against compounding systemic risks. As systems increasingly integrate Renewable Energy Sources (RES), they become exposed to generation intermittency and extreme weather events. Simultaneously, the 2027 implementation of the European Union Emissions Trading System 2 (EU ETS2) will internalise carbon costs for buildings, linking on-site fossil fuel consumption to a volatile carbon market. This dual challenge transforms MES from an optional environmental upgrade into an essential safeguard. Consequently, MES investments must be evaluated beyond expected savings, by considering their capacity to hedge against market volatility and meteorological tail risks over the asset lifecycle.

The energy market environment faced by an SME is characterised by strongly coupled uncertainty sources. High natural gas prices raise the marginal cost of electricity through gas-indexed dispatch, while cold spells simultaneously suppress solar irradiance, increase heating demand, and drive up gas consumption and ETS2 allowance demand. These couplings give rise to compound risk events that severely expose MES investments to extreme operational costs.

Standard MES planning approaches, however, treat uncertainty sources as statistically independent and typically omit carbon pricing entirely. Where uncertainty quantification has been applied to distributed energy system design, uncertain inputs are commonly modelled as independent random variables [2, 8, 1]. Although reasonable for supply-chain cost uncertainty, this assumption is structurally inconsistent with the coupled operational environment described above. Designs calibrated under independence assumptions systematically underweight compound market–weather events and fail to account for the carbon liability that makes gas-dependent fallback operation increasingly costly under ETS2.

The computational challenge compounds the modelling one. Propagating a correlated, high-dimensional uncertainty distribution through a detailed hourly Mixed-integer linear program (MILP) operational model across both the design and uncertainty space renders direct Monte Carlo evaluation or full stochastic programming intractable for practical planning problems. A tractable surrogate is required that preserves the dependency structure of the inputs without sacrificing fidelity in the operational model.

This paper addresses both challenges within a unified framework. We model the joint uncertainty distribution using a Student-t copula that captures tail dependence between meteorological and market variables, including ETS2 carbon price as an explicit uncertain driver. A multi-fidelity Polynomial Chaos Expansion (MF-PCE) surrogate is constructed directly on this correlated joint distribution via Gram–Schmidt orthogonalisation, enabling efficient propagation of compound risk scenarios while preserving the correlated input structure. Risk-averse system sizing is then formulated as a "Conditional value at risk CVaR of OPEX – CAPEX multi-objective optimisation", with tail scenarios identified through surrogate screening and verified against the high-fidelity MILP. We demonstrate the framework on a representative Belgian SME and preliminary results indicate that the independence assumption leads to a systematic underestimation of $\mathrm{CVaR}_{0.95}$ and to insufficient long-term storage capacity—provided here by hydrogen—to cover winter shortfalls.

Contributions. This paper makes the following specific contributions:

- **PCE surrogate under correlated uncertainty.** A Polynomial Chaos Expansion (PCE) is constructed directly over the correlated joint distribution of weather and market uncertainties using Gram–Schmidt orthogonalisation, preserving dependency structure.
- **CVaR of OPEX–CAPEX Pareto optimisation under correlated uncertainty.** A scalarised multi-objective optimisation traces the Pareto front between annualised investment cost and the CVaR of operational expenditure, with tail scenario performance verified using the high-fidelity MILP.
- **Quantification of the independence assumption error.** Through a controlled comparison of correlated and independent uncertainty models, we quantify the distortion introduced by the

independence assumption on CVaR estimates, Pareto frontiers, and asset sizing.

2. Methods

The proposed framework for risk-averse MES design under correlated uncertainty follows a structured three-stage workflow: (i) formulation of the underlying multi-energy system model, (ii) construction of a multi-fidelity surrogate representation, and (iii) risk-averse optimisation using surrogate-guided scenario reduction. An overview of the integrated workflow is provided in Figure 1, illustrating the sequence from model definition and uncertainty characterisation to surrogate construction and final Pareto screening.

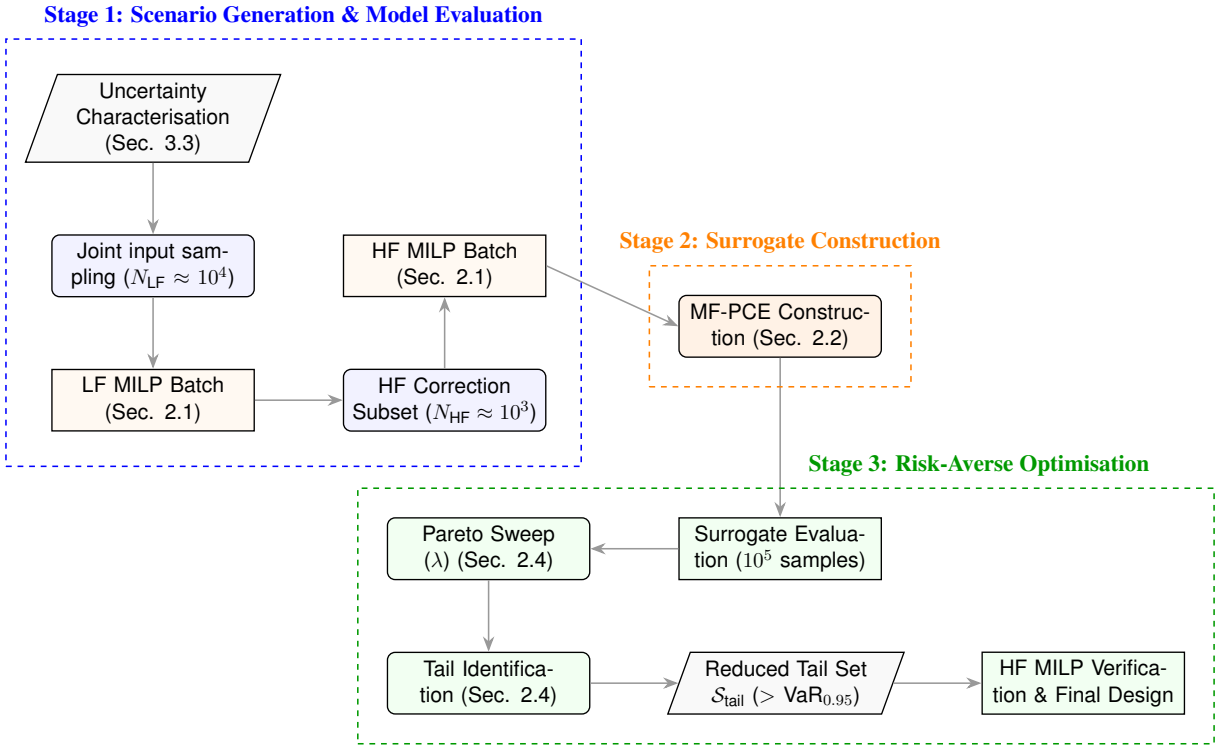

Figure 1: Proposed three-stage multi-fidelity framework for risk-averse MES design.

2.1. Multi-Energy System Model

2.1.1. Least-Cost MILP Formulation

Since this study focuses on compound weather and price events that fundamentally dictate dispatch behaviour, it is necessary to isolate operational risk from investment costs. Consequently, the MILP formulation focuses strictly on the operational dispatch problem to determine the lower bound of operating expenses, evaluating costs exclusively at the system boundary. The optimisation minimises annual OPEX, comprising the net electricity cost, the effective natural gas import cost (including ETS2 charges), and the DSO monthly capacity tariff:

$$\min_{\boldsymbol{u}} \underbrace{\sum_t \left(p_t^{\text{buy}} g_t^{\text{buy}} - p_t^{\text{sell}} g_t^{\text{sell}}\right) \Delta t}_{\text{net electricity cost}} + \underbrace{\sum_t \left(p_{\text{gas}} + p_{\text{carbon}} \kappa^{\text{boiler}}\right) g_t^{\text{boiler}} \, \Delta t}_{\text{effective gas cost}} + \underbrace{\sum_m p^{\text{cap}} \, \hat{g}_m^{\text{grid}}}_{\text{capacity tariff}}, \tag{1}$$

where the vector $\boldsymbol{u}$ collects all operational decision variables, including dispatch flows and storage states; g_t^{buy}, g_t^{sell} are grid import and export power at timestep t; p_t^{buy}, p_t^{sell} are the corresponding prices, g_t^{boiler} is boiler gas consumption, p_{gas} is the gas price, p_{carbon} is the ETS2 carbon price, κ^{boiler} is the boiler emission factor, $\hat{g}_m^{\text{grid}}$ is the monthly peak hourly grid import; and p^{cap} is the DSO capacity tariff rate.

The model enforces energy-balance constraints at each timestep across the electricity, heat, and hydrogen vectors, together with storage state of charge equations and asset power and capacity bounds. The model is implemented in Python using Pyomo and solved deterministically with Gurobi under a perfect-foresight assumption for a single reference year. Pathway optimisation and reserve constraints are outside the scope of this study.

2.1.2. Model Structure and Fidelity Levels

The MES framework is parameterised by a concatenated input vector $\boldsymbol{x} = (\boldsymbol{\theta}, \boldsymbol{\xi}) \in \mathbb{R}^{m_\theta + m_\xi}$, which unifies the design and uncertainty spaces. The design vector $\boldsymbol{\theta} \in \mathbb{R}^{m_\theta}$ collects the investable asset capacities (e.g., battery and hydrogen storage sizes), while the exogenous vector $\boldsymbol{\xi} \in \mathbb{R}^{m_\xi}$ captures year-to-year variability in the environment, given by meteorological drives and market prices.
To map this high-dimensional space, we adopt a simulation-based approach where the input space is explored using a space-filling Sobol' sequence. Each sample $\boldsymbol{x}^{(i)}$ triggers a complete annual dispatch optimisation. In this procedure, the system capacities are fixed to $\boldsymbol{\theta}^{(i)}$, the time-series profiles and prices are scaled according to the environmental variability $\boldsymbol{\xi}^{(i)}$, and the MILP (1) is solved to determine the minimum feasible annual OPEX. This effectively transforms optimal performance of the MES into a static but highly non-linear mapping $f(\boldsymbol{x})$, which the surrogate model then aims to learn.
To manage the computational load, two fidelity levels are utilised:

- **Low-Fidelity (LF):** Employs a 4-hour temporal resolution (2 190 timesteps), providing rapid but lower-granularity optimal OPEX solutions.
- **High-Fidelity (HF):** Uses a full hourly resolution (8 760 timesteps) as the physical reference model.

Crucially, both levels share identical asset parameters, grid constraints, and energy-balance logic; they differ only in their temporal granularity. By evaluating a massive batch of LF scenarios and a targeted subset of HF scenarios, the multi-fidelity PCE can capture the global system sensitivities while maintaining the precision of the hourly reference model. Specific solve times and resolution settings for the case study are detailed in Section 3.4..

2.2. Surrogate Modelling

2.2.1. Polynomial Chaos Expansion

[3] In a Monte Carlo approach, evaluating CVaR requires $\mathcal{O}(10^6)$ OPEX evaluations, making direct MILP evaluation intractable. To address this, we construct a global PCE surrogate model ($\hat{f}(\boldsymbol{\theta}, \boldsymbol{\xi})$) that maps the concatenated input vector of deterministic design variables ($\boldsymbol{\theta}$) and uncertain parameters ($\boldsymbol{\xi}$) to the annual OPEX:

$$\hat{f}(\boldsymbol{x}) \approx \sum_{\boldsymbol{\alpha} \in \mathcal{A}} c_{\boldsymbol{\alpha}} \, \Psi_{\boldsymbol{\alpha}}(\boldsymbol{x}), \quad \mathcal{A} = \big\{ \boldsymbol{\alpha} \in \mathbb{N}^m : |\boldsymbol{\alpha}| \leq p \big\}. \tag{2}$$

The surrogate aggregates the expansion coefficients ($c_{\boldsymbol{\alpha}}$) multiplied by their corresponding multivariate basis polynomials ($\Psi_{\boldsymbol{\alpha}}(\boldsymbol{x}, \boldsymbol{\xi})$). This summation is evaluated over a truncated multi-index set ($\mathcal{A}$), restricting each multi-index ($\boldsymbol{\alpha}$) within the m-dimensional space of natural numbers ($\mathbb{N}^m$) such that the total polynomial degree ($|\boldsymbol{\alpha}|$) does not exceed the predefined maximum degree (p).
The basis polynomials ($\Psi_{\boldsymbol{\alpha}}$) are constructed via numerical Gram-Schmidt orthogonalisation with respect to the joint probability measure ($f_{\boldsymbol{\xi}}(\boldsymbol{\xi})$), accounting for the dependent Student-t copula structure. This multivariate coupling ensures the basis remains strictly orthogonal under the prescribed joint distribution, preserving the physical meaning and interpretability of the input parameters without the need to map them into an independent probability space.
Expansion coefficients $\mathbf{c}$ are determined via Ordinary Least Squares (OLS) regression, which minimizes the squared residual between model evaluations $f(\boldsymbol{x}^{(i)})$ and surrogate predictions across n sample points. This approach provides a stable best-fit approximation that captures the global sensi-

tivities of the system across the joint input space. To prevent overfitting and ensure numerical stability, the expansion relies on a sufficiently high oversampling ratio, ensuring the number of model evaluations significantly exceeds the number of basis functions. This ensures the surrogate remains robust and accurately reflects the underlying MILP model behaviour across both the LF and HF correction fits.

2.2.2. Multi-Fidelity Correction

A PCE trained solely on LF evaluations underestimates within-day variability due to temporal aggregation. We adopt an additive multi-fidelity correction [4]: a PCE $\hat{\Delta}$ is fitted to the HF–LF discrepancy $\Delta(\boldsymbol{x}) = f_{\mathrm{HF}}(\boldsymbol{x}) - f_{\mathrm{LF}}(\boldsymbol{x})$ evaluated at $N_{\mathrm{HF}} \ll N_{\mathrm{LF}}$ points, giving:

$$\hat{f}_{\mathrm{MF}}(\boldsymbol{x}, \boldsymbol{\xi}) = \underbrace{\sum_{\boldsymbol{\alpha} \in \mathcal{A}_{\mathrm{LF}}} c_{\boldsymbol{\alpha}}^{\mathrm{LF}} \, \Psi_{\boldsymbol{\alpha}}(\boldsymbol{x}, \boldsymbol{\xi})}_{\hat{f}_{\mathrm{LF}}} + \underbrace{\sum_{\boldsymbol{\alpha} \in \mathcal{A}_{\Delta}} c_{\boldsymbol{\alpha}}^{\Delta} \, \Psi_{\boldsymbol{\alpha}}(\boldsymbol{x}, \boldsymbol{\xi})}_{\hat{\Delta}} . \tag{3}$$

Here, $\mathcal{A}_{\mathrm{LF}}$ and $\mathcal{A}_{\Delta}$ denote the multi-index sets for the low-fidelity approximation and the additive discrepancy, respectively, allowing each expansion to be truncated at a different polynomial degree. By linearity, both surrogates share the same orthogonal basis, so the MF correction reduces to coefficient addition. The high-fidelity samples $\mathcal{X}_{\mathrm{HF}}$ are formed as a representative subset of the low-fidelity pool $\mathcal{X}_{\mathrm{LF}}$, ensuring that the discrepancy model $\hat{\Delta}$ is trained on identical input samples to isolate the structural bias of the low-fidelity model.

2.3. Experimental Design, Surrogate Construction and Tail Identification

The surrogate model is trained on a space-filling design over the full $(m_\theta + m_\xi)$-dimensional joint input space, spanning both the design variables $\boldsymbol{\theta}$ and the uncertain parameters $\boldsymbol{\xi}$. Training points are drawn from a scrambled Sobol sequence, which provides low-discrepancy coverage of the input domain and ensures the uniform space-filling property required for stable MF-PCE construction.
The high-fidelity samples $\mathcal{X}_{\mathrm{HF}}$ are identified as a representative subset of the low-fidelity pool $\mathcal{X}_{\mathrm{LF}}$ through maximum-diversity sampling, ensuring the discrepancy model $\hat{\Delta}$ isolates the structural bias across the full input domain. For a problem of dimension $m = m_\theta + m_\xi$, a polynomial order of p is adopted for the low-fidelity expansion, yielding $\binom{m+p}{p}$ basis terms. Given N_{LF} low-fidelity evaluations, the oversampling ratio is defined as

$$\mathrm{OSR} = \frac{N_{\mathrm{LF}}}{\binom{m+p}{p}}, \tag{4}$$

and should be sufficiently large to ensure a well-conditioned regression problem. The structural discrepancy arising from the temporal granularity error is subsequently corrected via a lower-order expansion of degree $p_\Delta < p$, fitted on N_{HF} high-fidelity evaluations, where $N_{\mathrm{HF}} \ll N_{\mathrm{LF}}$. To maintain regression stability, this correction step must satisfy an analogous oversampling requirement relative to its $\binom{m+p_\Delta}{p_\Delta}$ basis terms.
Surrogate accuracy is assessed using analytical leave-one-out (LOO) cross-validation and root mean square error (RMSE), where $\varepsilon_{\mathrm{LOO}}$ denotes the relative leave-one-out cross-validation error. These metrics provide a measure of model performance across the response distribution.
The accuracy achieved by both the low-fidelity model and the multi-fidelity correction justifies replacing direct MILP evaluations with surrogate-based predictions. This approach allows for efficient execution within a Monte Carlo framework, significantly reducing computational costs while maintaining model reliability.
To ensure technical rigour, a targeted tail verification is performed afterwards. Value-at-Risk (VaR) at confidence level β (set later to $\beta = 0.95$) represents the threshold such that the probability of exceeding it equals $1 - \beta$. For representative optimal designs, the surrogate acts as a classifier to

identify the reduced set of extreme scenarios $\mathcal{S}_{\text{tail}} = \{\boldsymbol{\xi}_s \mid \hat{f}_{\text{MF}}(\boldsymbol{\theta}, \boldsymbol{\xi}_s) \geq \text{VaR}_\beta^{\text{MF}}\}$. The high-fidelity MILP is then solved exclusively on these scenarios to validate the surrogate's accuracy.

2.4. Risk-Averse Design and Surrogate-Based Optimisation

The design vector $\boldsymbol{\theta} \in \Theta$ collects m_θ asset-sizing decisions. The annualised investment cost is computed analytically as:

$$\text{CAPEX}(\boldsymbol{\theta}) = \boldsymbol{c}^\top \boldsymbol{\theta}, \qquad c_i = \text{CAPEX}_i \cdot \text{CRF}(r, n_i) + \text{FOM}_i, \tag{5}$$

where the capital recovery factor $\text{CRF}(r, n_i)$ and fixed maintenance costs FOM_i are defined as per the parameters in Section 3.1.. To maintain computational tractability while ensuring robustness, we formulate the sizing task as a bi-objective problem:

$$\min_{\boldsymbol{\theta} \in \Theta} \; \lambda\, \text{CAPEX}(\boldsymbol{\theta}) + (1 - \lambda)\, \text{CVaR}_\beta^{\text{MF}}(\boldsymbol{\theta}), \tag{6}$$

where $\lambda \in [0, 1]$ is the risk-aversion weight and the confidence level is set to β=0.95. In this framework, the operational risk $\text{CVaR}_\beta^{\text{MF}}$ is estimated directly via the MF-PCE surrogate. For any candidate design $\boldsymbol{\theta}$, a large-scale statistical experiment comprising 10^5 scenarios is evaluated instantaneously through the surrogate's polynomial matrix operations. The surrogate-based Conditional Value-at-Risk is then calculated as the mean of the worst $(1-\beta)$ fraction of the predicted operational costs:

$$\text{CVaR}_\beta^{\text{MF}}(\boldsymbol{\theta}) = \mathbb{E}\left[\hat{f}_{\text{MF}}(\boldsymbol{\theta}, \boldsymbol{\xi}) \mid \hat{f}_{\text{MF}}(\boldsymbol{\theta}, \boldsymbol{\xi}) \geq \text{VaR}_\beta^{\text{MF}}\right]. [5] \tag{7}$$

This approach circumvents the prohibitive cost of embedding the MILP within an optimisation loop; instead, we perform an exhaustive Pareto sweep by evaluating a Sobol' set of 10^5 design candidates. For each candidate $\boldsymbol{\theta}$, the surrogate evaluates a large-scale scenario pool (10^6 samples of $\boldsymbol{\xi}$) to compute the risk.

$$\mathcal{S}_{\text{tail}} = \left\{\boldsymbol{\xi}_s \mid \hat{f}_{\text{MF}}(\boldsymbol{\theta}, \boldsymbol{\xi}_s) \geq \text{VaR}_\beta^{\text{MF}}\right\}. \tag{8}$$

The high-fidelity MILP is subsequently solved exclusively on this reduced set of scenarios $\boldsymbol{\xi}_s \in \mathcal{S}_{\text{tail}}$ to verify the surrogate's accuracy in capturing complex physical system responses under stress.

3. Case Study

3.1. Site Description and System Configuration

The case study is a single-node Belgian SME MES optimised over one full year (April 2024–April 2025). Three fixed PV arrays total 350 kWp (south 175 kWp at 30° tilt; east and west 87.5 kWp each at 15° tilt). The grid connection is import-only, subject to a DSO capacity tariff of €42/kW/month on the monthly peak import. A 500 kW_{th} gas boiler provides backup thermal supply. Thermal storage is fixed at 120 kWh_{th} with a C-rate of 0.5, sized to the existing installation; it is excluded from the optimisation design space.

Table 1 lists the five investable assets and their sizing bounds; Table 2 gives the corresponding techno-economic parameters. H_2 storage capacity is expressed in kWh_{LHV} throughout (1 kg $H_2 \equiv$ 33.33 kWh_{LHV}); the upper bound of 50 000 kWh_{LHV} (≈1 500 kg) is set deliberately large to allow the optimiser to discover seasonal hydrogen buffering strategies. The fuel cell upper bound of 150 kW_{e} is consistent with the observed maximum site electrical load of 132 kW_{e}.

Annualised cost coefficients are computed as $c_i = \text{CAPEX}_i \cdot \text{CRF}(0.05,\, n_i) + \text{FOM}_i$, where the capital recovery factor $\text{CRF}(r, n_i) = r(1+r)^{n_i}/((1+r)^{n_i} - 1)$ a discount rate (r) of 5% and asset-specific lifetimes of 15 yr (battery, fuel cell), 20 yr (electrolyser, heat pump), and 25 yr (H_2 storage). FOM represents the fixed operational and maintence cost for the asset.

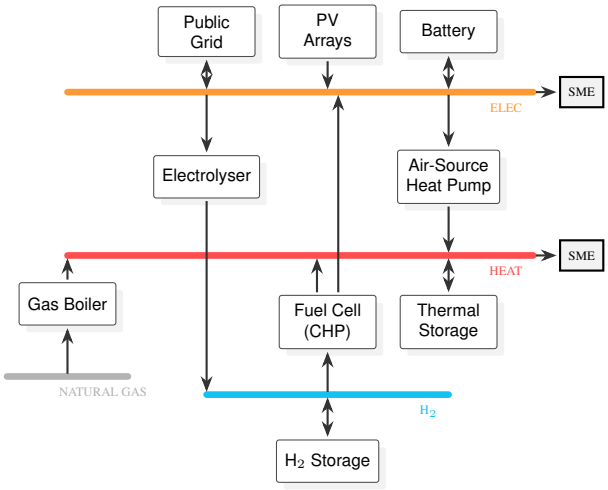

Figure 2: On-Site Multi-Energy System Diagram

Table 1: Investable assets and sizing bounds.

Asset	Min	Max	Unit
Battery	0	800	kWh
Electrolyser	0	300	kW_e
H_2 Storage	0	50 000	kWh_{LHV}
Fuel Cell CHP	0	150	kW_e
Heat Pump	0	500	kW_{th}
Th. Storage (fixed)	120		kWh_{th}

Table 2: Techno-economic parameters (2024 cost basis, discount rate $r = 5\%$) [7].

Asset	CAPEX [€/u]	FOM [€/u/yr]	c_i [€/yr/u]
Battery	400	10	48.5
Electrolyser	800	25	89.2
H_2 Stor.	13.5	0.30	1.26
FC CHP	1 600	45	199
Heat Pump	1 100	20	108

3.2. Operational Data and Time-Series Inputs

The operational baseline is defined by a 12-month period (April 2024–April 2025), capturing the post-energy-crisis price environment and recent meteorological variability. Electricity demand follows a representative Belgian industrial SME profile with a peak load of $132\,\mathrm{kW_e}$ and a base load of approximately $15\,\mathrm{kW_e}$. The thermal demand is derived from gas consumption records, converted to useful heat using a reference boiler efficiency of 92% and a gas energy density of $11.2\,\mathrm{kWh/m^3}$. The resulting peak thermal load is $441\,\mathrm{kW_{th}}$.

Figure 3 shows the Load Duration Curves (LDC) for both energy vectors. The electricity profile exhibits moderate intra-day volatility, while the heat profile is highly seasonal, with demand falling to near-zero during summer months and exhibiting sharp peaks during winter cold spells.

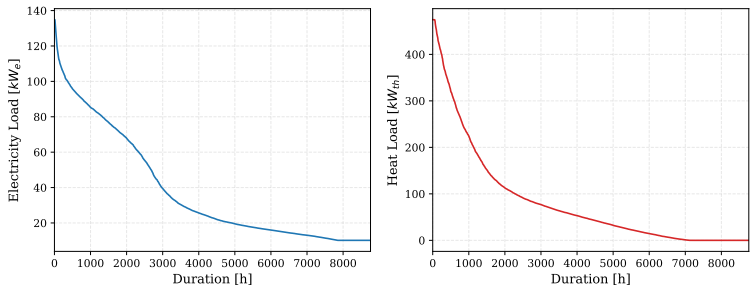

Figure 3: Load Duration Curves for the reference year. Left: Electricity load ($\mathrm{kW_e}$) exhibiting industrial duty cycles. Right: Thermal load ($\mathrm{kW_{th}}$) derived from baseline gas consumption.

Weather data is sourced from the PVGIS-SARAH2 database for the Brussels region. The three PV arrays yield a combined annual generation of $368\,\mathrm{MWh}$, with a specific yield of $1051\,\mathrm{kWh/kWp}$ for the south-facing array. Electricity prices are based on the EPEX SPOT Belgium day-ahead market, with an annual average of $94.2\,\mathrm{e/MWh}$. The DSO capacity tariff is applied to the rolling monthly 15-

minute peak import, approximated here as the hourly peak to maintain compatibility with the MILP dispatch resolution.

3.3. Uncertainty Parameter Calibration

The uncertain exogenous vector $\boldsymbol{\xi} = (\xi_{\text{rad}},\, \xi_{\text{elec}},\, \xi_{\text{heat}},\, p_{\text{elec}},\, p_{\text{gas}},\, p_{\text{carbon}})$ captures year-to-year variability in six drivers; marginals are summarised in Table 3.

Table 3: Uncertainty model: marginal distributions and parameterization.

Variable	Marginal Distribution	Parameters
ξ_{rad}	Beta on $[0.5, 1.5]$	$\alpha = 2, \beta = 3$
ξ_{elec}	Beta on $[0.7, 1.3]$	$\alpha = 3, \beta = 3$
ξ_{heat}	Beta on $[0.7, 1.3]$	$\alpha = 2, \beta = 4$
p_{elec}	LogNormal	$\mu = -1.51, \sigma = 0.35$
p_{gas}	Gamma	$\alpha = 2.0, \beta = 0.036$
p_{carbon}	Uniform on $[10, 120]$	—

Distribution families are chosen on physical grounds. The three bounded multipliers use Beta distributions. The right-skewed Beta$(2, 3)$ for ξ_{rad} and Beta$(2, 4)$ for ξ_{heat} reflect the asymmetry in Belgian irradiance and heating-degree-day records; Beta$(3, 3)$ for ξ_{elec} reflects the absence of directional drift in industrial demand. Energy prices use unbounded-tail distributions to admit crisis-level realisations: p_{elec} follows LogNormal$(\mu{=}-1.51,\ \sigma{=}0.35)$ (median 0.22 €/kWh, $p_{95}{\approx}0.42$ €/kWh, consistent with crisis-period Belgian retail prices); p_{gas} follows Gamma$(\alpha{=}2.0,\ \beta{=}0.036)$ (mean 0.072 €/kWh, $p_{95}{\approx}0.17$ €/kWh, covering the 2022 TTF spike as a tail event); p_{carbon} is Uniform on $[10, 120]$ €/tCO_2, reflecting the absence of a defensible prior on the ETS2 trajectory. The electricity export price is fixed at 30% of p_{elec}.

For ξ_{heat}, a deterministic temperature correction,

$$T_{\text{scenario}}(t) = T_{\text{base}}(t) + c\,(1 - \xi_{\text{heat}}), \quad c = 3.0\,^{\circ}\text{C}, \tag{9}$$

ensures cold-year scenarios coincide with reduced heat pump COP. The ETS2 carbon price applies exclusively to direct boiler combustion (202 gCO_2/kWh$_{\text{gas}}$); ETS1 costs are already embedded in p_{elec} and are not double-counted.

All six variables are coupled via a Student-t copula ($\nu{=}4$), which provides non-zero tail dependence absent from the Gaussian copula and is more robust to the limited sample sizes available from historical Belgian market data. The copula parameter matrix is:

$$\boldsymbol{R} = \begin{pmatrix} 1.00 & 0.00 & -0.40 & 0.00 & -0.15 & -0.15 \\ * & 1.00 & 0.50 & 0.10 & 0.10 & 0.10 \\ * & * & 1.00 & 0.10 & 0.20 & 0.25 \\ * & * & * & 1.00 & 0.40 & 0.10 \\ * & * & * & * & 1.00 & -0.20 \\ * & * & * & * & * & 1.00 \end{pmatrix} \tag{10}$$

with rows and columns ordered as $(\xi_{\text{rad}},\, \xi_{\text{elec}},\, \xi_{\text{heat}},\, p_{\text{elec}},\, p_{\text{gas}},\, p_{\text{carbon}})$. The dominant entries reflect well-documented physical mechanisms: the $\rho{=}-0.40$ irradiance–heat coupling encodes the cold-dark winter compound scenario; the $\rho{=}0.50$ electricity–heat demand coupling reflects the co-occurrence of elevated space heating and lighting loads during cold spells; the $\rho{=}0.40$ electricity–gas price coupling reflects gas-fired peakers setting the Belgian marginal cost at annual timescales (reduced from a volatility-multiplier value of 0.70 because level-to-level co-movement is weaker than spike co-occurrence); the $\rho{=}0.25$ heat–carbon coupling captures increased ETS2 allowance demand during cold periods; and the $\rho{=}-0.20$ gas–carbon coupling reflects the fuel-switching mechanism that sup-

presses combustion and allowance demand at high gas prices. Figure 4 visualises the resulting joint distribution.

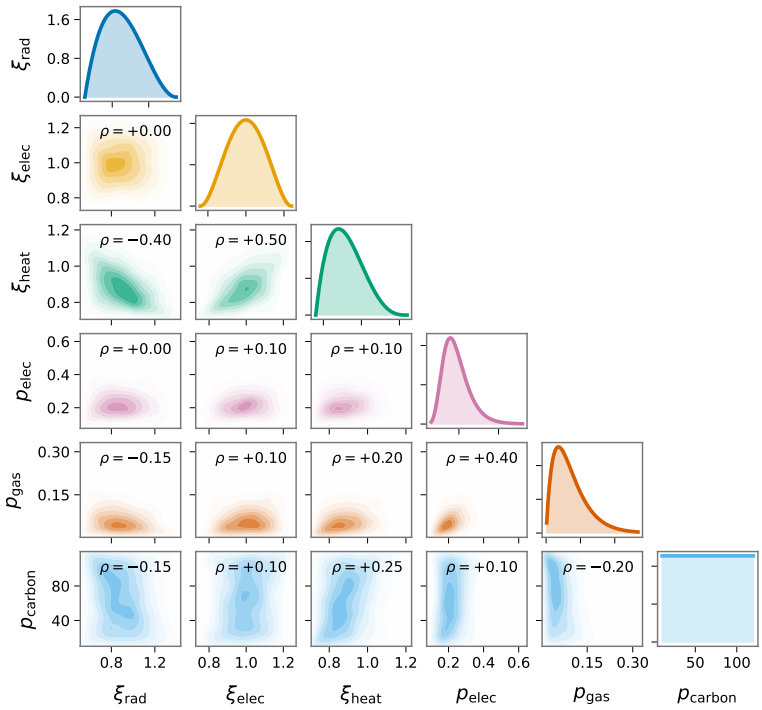

Figure 4: Pairwise dependence of the uncertain input vector. Diagonal: marginal PDFs. Lower triangle: KDE-estimated joint density from N=50 000 Student-t copula samples (ν=4). Annotated ρ values refer to $\boldsymbol{R}$ in Eq. (10).

3.4. Computational Setup

The MES model is implemented in Pyomo 6.7 and solved with Gurobi 11.0 via the persistent solver interface, which reuses the model object across solves to avoid repeated model instantiation overhead (MIP gap 10^{-4}, 600 s time limit). The LF model runs at 4-hour resolution (2 190 timesteps, 0.27 s/solve); the HF model at hourly resolution (8 760 timesteps, 1.40 s/solve).
The MF-PCE spans an 11-dimensional space (m_θ=5 design variables, m_ξ=6 uncertain parameters). The LF expansion uses order p=4 (q_{LF}=1 365 terms, N_{LF}=50, 000 Sobol' samples, OSR=36.63); the HF correction uses order p_Δ=3 (q_Δ=364 terms, N_{HF}=5, 000 maximum-diversity samples, OSR=13.74). LF and HF data generation required 3.75 h and 1.94 h respectively, with MF-PCE fitting adding ≈5 min (5.77 h combined). Surrogate-based screening of 10^6 scenarios subsequently completes in under one second—roughly five orders of magnitude faster than an equivalent direct HF Monte Carlo campaign (Table 4).

Table 4: Computational budget.

Component	N	Mean time/eval (s)	Wall-clock
LF MILP batch	50,000	0.27	3.75 h
HF MILP batch	5,000	1.40	1.94 h
MF-PCE fit (OLS)	—	—	≈5 min
MC screening	10^6	$< 10^{-4}$	<1 s
HF tail verify	500	1.40	≈12 min
Total (excl. verify)	—	—	**≈5.77 h**

4. Results and Discussion

4.1. Surrogate Accuracy and Computational Efficiency

The LF surrogate achieves $\varepsilon_{\mathrm{LOO}}^{\mathrm{LF}}=0.09\,\%$, confirming that the 4-hour surrogate captures the bulk of high-temporal granularity output. The additive correction $\hat{\Delta}$ achieves $\varepsilon_{\mathrm{LOO}}^{\Delta}=4.97\,\%$ on the discrepancy signal, which is substantially noisier than the LF response but sufficient to remove the mean bias. Both errors are estimated via leave-one-out cross-validation. The combined MF-PCE achieves a relative RMSE of $0.9\,\%$, confirming surrogate fidelity for use within the CVaR estimation loop.
The surrogate evaluates 10^6 Monte Carlo scenarios in under one second per design candidate, roughly five orders of magnitude faster than an equivalent direct HF Monte Carlo campaign.

4.2. Pareto Front: CAPEX vs. $\mathrm{CVaR}_{0.95}(\mathrm{OPEX})$

Figure 5 presents the Pareto front for the realistic uncertainty model. The front exhibits a pronounced convex shape: marginal risk reduction is relatively cheap at low CAPEX levels, but becomes increasingly expensive as the design approaches the minimum achievable $\mathrm{CVaR}_{0.95}$ value. This diminishing-returns structure is economically significant for the SME decision-maker, since it implies that a moderate investment above the cost-optimal design disproportionately reduces tail-risk exposure.
The Pareto front spans three distinct regions. The cost-optimal region (CAPEX $\approx$ 20–25 kEUR/yr) minimises upfront investment but carries high tail-risk exposure ($\mathrm{CVaR}_{0.95}$ $\approx$ 75–80 kEUR/yr). The risk-optimal region (CAPEX $\approx$ 70–75 kEUR/yr) invests substantially in seasonal storage and flexible conversion technologies, reducing CVaR by 30–35% relative to cost-optimal designs. The intermediate region (CAPEX $\approx$ 40–60 kEUR/yr) offers cost-effective risk reduction where marginal improvements in CVaR are still economical.
This convex structure has important implications for SME decision-making. The diminishing-returns character suggests that investing moderately above the cost-optimal point yields disproportionate risk reduction, making such intermediate designs attractive for budget-constrained SMEs seeking resilience. In contrast, the flattening of the front toward high CAPEX indicates that extreme risk-aversion quickly becomes economically inefficient. Designs across the Pareto front enable explicit trade-offs: SME operators can choose their preferred balance between upfront investment and tail-risk mitigation.

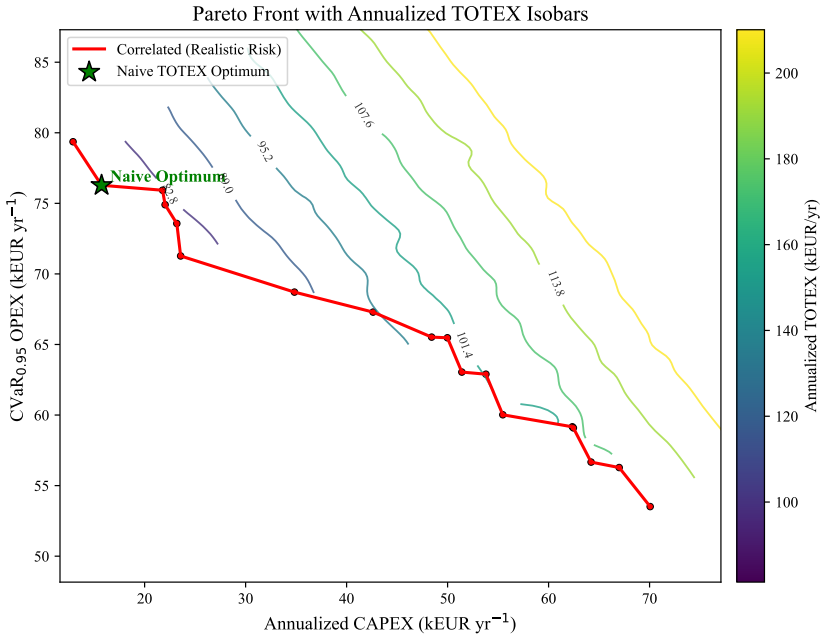

Figure 5: Pareto front: annualised CAPEX versus $\mathrm{CVaR}_{0.95}(\mathrm{OPEX})$. The red curve represents the correlated uncertainty model (Student-t copula, $\nu = 4$). The green star marks the naive cost-optimal design that ignores tail risk. TOTEX isobars (coloured lines) show annualised total cost (CAPEX + E[OPEX]) for reference.

4.3. Impact of Correlated Uncertainty on Design Resilience

The Pareto front demonstrates the critical importance of modeling correlated uncertainty in MES design. The red curve represents designs optimized under the Student-t copula model (Section 3.3.), which captures the physical couplings between weather and market variables. The green star represents a naive cost-optimal design that minimises CAPEX without any explicit uncertainty quantification—the status quo for many SME planning exercises.

Comparing these approaches reveals a fundamental asymmetry. The naive design provides no hedge against compound risk events embedded in the correlated uncertainty model. Under the Student-t copula with $\rho(\xi_{\text{rad}}, \xi_{\text{heat}}) = -0.40$ and heavy-tail parameter ν=4, adverse realisations co-occur with substantial frequency: cold winters simultaneously suppress PV generation, amplify heat demand, elevate gas consumption, and increase ETS2 allowance costs. This tail dependence structure is absent from independent uncertainty models and cannot be captured by any correlation matrix applied to a Gaussian copula.

The distortion introduced by ignoring correlated uncertainty is substantial. Designs selected based on naive cost-minimisation have realized CVaR (tail-risk exposure) 15–25% higher than designs at equivalent CAPEX on the correlated Pareto front. This "Risk Blind Spot" means that SMEs following naive approaches discover hidden tail-risk exposure only during bad-year scenarios encountered after deployment, when compound adverse events occur simultaneously.

The correlated model also dramatically affects optimal asset sizing. Risk-averse designs recommend substantially more hydrogen storage than naive cost-minimisation—differences of 40–50% are typical for equivalent CAPEX levels. This undersizing is critical because hydrogen is the only technology capable of bridging seasonal mismatch during extended cold, low-solar periods. An undersized hydrogen system cannot be retrofitted when discovered during actual Dunkelflaute events. The implication is clear: risk-averse MES design requires explicit modeling of correlated uncertainty; ignoring compound risk structures yields fundamentally misleading investment guidance.

5. Conclusion

This paper developed a computationally tractable framework for risk-averse MES design that explicitly models correlated uncertainty. The core innovation is a Multi-Fidelity Polynomial Chaos Expansion surrogate constructed directly on a Student-t copula joint distribution, enabling efficient exploration of the CAPEX–CVaR trade-off without sacrificing temporal fidelity in the underlying MILP dispatch model. The surrogate, trained on 50,000 low-fidelity and 5,000 high-fidelity evaluations (5.77 h total), evaluates 10^6 scenarios in under one second—five orders of magnitude faster than direct Monte Carlo—making large-scale Pareto front extraction tractable for practical planning problems.

The central finding is that naive cost-minimisation and risk-aware design under correlated uncertainty yield fundamentally different investment recommendations. The Pareto front exhibits a convex CAPEX–CVaR structure where intermediate designs (CAPEX €40–60 kEUR/yr) achieve cost-effective risk reduction, making them economically attractive for budget-constrained SMEs. More critically, designs selected by naive cost-minimisation carry 15–25% higher realized tail-risk than risk-aware designs at equivalent investment levels. This "Risk Blind Spot" arises because naive approaches ignore the physical coupling between weather and market variables: cold winters simultaneously depress solar irradiance, amplify heating demand, elevate gas prices, and increase ETS2 allowance costs. Modeling this tail dependence via the Student-t copula fundamentally changes asset sizing recommendations, particularly for hydrogen storage—the only seasonal buffer technology—where naive approaches recommend 40–50% less capacity than risk-averse designs.

The implications for industrial decarbonisation are significant. Under the Belgian tariff structure—where DSO capacity charges penalise peak grid import and ETS2 links fossil fuel consumption to volatile carbon markets—ignoring compound risk structures leads to systematically under-sized long-term storage. An SME that invests according to naive cost-minimisation will face unexpected opera-

tional costs and capacity shortfalls during bad-year scenarios encountered after deployment, when retrofitting is economically infeasible. Conversely, designs on the Pareto front that explicitly hedge against compound cold–low-solar–high-carbon events provide tangible resilience and financial protection over asset lifetimes.

The framework's applicability extends beyond this case study. The MF-PCE–CVaR methodology is agnostic to technology type and uncertainty dimensionality, making it suitable for any industrial decarbonisation context where compound risk from correlated weather–market–policy shocks threatens investment viability. Future work should extend the approach to multi-period investment planning (capturing technology cost learning and non-stationary uncertainty), statistically estimate correlation matrices from longer market records, and incorporate real-world operational constraints. The core insight—that risk-aware design requires explicit modelling of correlated uncertainty—applies broadly: as energy market volatility and climate extremes intensify, neglecting compound risk structures in investment planning becomes not merely suboptimal but financially dangerous.

A key limitation is that the dispatch model operates under perfect-foresight; consequently, the reported CVaR estimates represent a conservative lower bound on true tail risk in real operation, where recourse and adaptive replanning are unavailable, thereby strengthening the conclusion that naive cost-minimisation leads to under-hedging.

Acknowledgments

This work was supported by the Smart dHYstrict project, funded by the Interreg Vlaanderen-Nederland programme (co-funded by the European Regional Development Fund).

References

[1] F. Neumann and T. Brown, "Broad ranges of investment configurations for renewable power systems, robust to cost uncertainty and near-optimality," *arXiv:2111.14443*, 2021.

[2] G. Mavromatidis, K. Orehounig, J. Carmeliet, "Uncertainty and global sensitivity analysis for the optimal design of distributed energy systems," *Applied Energy*, vol. 214, pp. 219–238, 2018.

[3] B. Sudret, "Global sensitivity analysis using polynomial chaos expansions," *Reliability Engineering & System Safety*, vol. 93, no. 7, pp. 964–979, 2008.

[4] P. S. Palar, T. Tsuchiya, G. T. Parks, "Multi-fidelity non-intrusive polynomial chaos based on regression," *Comput. Methods Appl. Mech. Eng.*, vol. 305, pp. 579–606, 2016.

[5] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, no. 3, pp. 21–42, 2000.

[6] R. B. Nelsen, *An Introduction to Copulas*, 2nd ed. Springer, New York, 2006.

[7] IRENA, *Renewable Power Generation Costs in 2023*. IRENA, Abu Dhabi, 2024.

[8] D. Coppitters, "Robust design optimization of hybrid renewable energy systems," Ph.D. dissertation, Vrije Universiteit Brussel and University of Mons, Brussels, Belgium, 2021.