Journal of Electrotechnology, Electrical Engineering and Management (2026)
Clausius Scientific Press, Canada
DOI: 10.23977/jeeem.2026.090120
ISSN 2560-6697 Vol. 9 Num. 1

Uncertainty-Set Design for Two-Stage Robust Scheduling of Active Distribution Networks with Correlated Forecast Errors

Shijie Chen*, Zhenyu Zhao, Xufei Liu, Yan Jiang, Kai Chen, Binbin Zhou, Congtong Zhang, Junkai Liang

Yunnan Power Grid Dispatching and Control Center, Kunming, Yunnan, China
**Corresponding Author*

Keywords: Active distribution network; Two-stage robust optimization; Uncertainty set; Principal component analysis; Correlated forecast errors

Abstract: Day-ahead scheduling of active distribution networks (ADNs) must hedge photovoltaic (PV), wind and load forecast errors that are correlated across sites and hours. Uncertainty sets built from the principal components of the errors describe this correlation better than independent budget sets, but it is unclear whether they give better two-stage robust schedules. This paper compares such sets on a 69-bus ADN with micro-turbines, batteries, demand response and voltage limits, all solved by the same column-and-constraint generation algorithm with an exact worst-case sub-problem, sized on validation days for a target shedding frequency, and evaluated on 2,000 test days. Under the min–max criterion the independent budget set is better: at shedding targets of 1% and 0.2% of days, its schedules cost 9.5% and 12.2% less. The principal-component sets hedge common-mode PV deficits with day-ahead energy that is mostly unused and curtails 13–19% of the renewable energy. Used only for robust feasibility, with a sample-average cost, the principal-component set gives the best trade-off: shedding on 0.05% of the test days at a mean cost 0.6% above a stochastic program that sheds load on 3.0% of the days.

1. Introduction

Active distribution networks (ADNs) increasingly schedule their own source–grid–load–storage (SGLS) resources: photovoltaics (PV), wind turbines, micro-turbines, batteries, interruptible demand and the upstream grid exchange. Building loads can report per-period flexibility with a disaggregation that stays feasible in real time [1] and secure reserve against temperature uncertainty with distributionally robust chance constraints [2], and ADNs host renewable-based rural heating [3] and ammonia-producing agricultural parks [4]. Day-ahead scheduling must hedge forecast errors that are not independent: PV plants on one feeder see the same clouds, wind turbines the same weather and load zones the same temperature, and the errors persist for hours. SGLS models capture such uncertainty with scenarios [5, 6], multi-scenario park models coordinated with the distribution operator [7], or data-driven flexibility-demand polytopes [8].

Two-stage adaptive robust optimization (ARO) [9] keeps a schedule feasible for every error in an

uncertainty set: the first stage fixes commitment and reserves, and the second re-dispatches after the errors are known. First applied to (security-constrained) unit commitment [10, 11] and solved by column-and-constraint generation (C&CG) [12], it coordinates tap changers, reactive compensation and storage in distribution networks [13, 14]. The uncertainty set decides how conservative the schedule is. The budget set of Bertsimas and Sim [15] bounds each coordinate and the number of simultaneous deviations, sized by a probabilistic bound that assumes independence; data-driven construction can reduce conservativeness at equal robustness [16]. Dynamic uncertainty sets add temporal and spatial wind correlation to robust economic dispatch [17], and storage in the multistage extension [18]. Data-driven sets learn the shape from observations by principal component analysis (PCA) with kernel density estimation [19], Dirichlet-process mixtures [20], hypothesis tests [21], ellipsoidal prediction regions [22], or updated convex hulls of observed daily profiles [23]. Distributionally robust formulations use a ball of distributions [24] or moment-based ambiguity sets, which with linear decision rules give second-order cone programs in generation planning [25] and have been applied to energy–reserve self-scheduling of hydropower-based virtual power plants [26]. In SGLS dispatch, Zhang et al. [8] decorrelated wind, PV and load with PCA, bounded each component by kernel-density intervals and mapped the box back to a flexibility-demand polytope, whose volume the correlation reduced by about 11%. Other work changes how the set is used: unified stochastic and robust unit commitment weights the expected and worst-case costs [27], and further variants connect robust and scenario-based stochastic models [28].

Three questions remain open. First, does a set that describes the error correlation more accurately lead to a better schedule? A correlation-aware set is smaller for the same coverage, but the schedule is chosen by a min–max criterion, and a better description of the errors need not give a better worst-case hedge. Second, how should a set be sized? Budget sets use a bound that assumes independence, and data-driven sets are usually sized to contain a fraction of the historical error vectors (joint coverage), which forces large sets in the hundreds of dimensions of a day-ahead schedule. Third, the reliability of a set is rarely measured out of sample in the operator's terms: load shedding, curtailment and cost. Shedding matters most where critical facilities are supplied, e.g., hospitals, whose resilience depends on interdependent power, water, firefighting and space systems [29].

This paper answers these questions for two-stage robust SGLS scheduling of an ADN by comparing uncertainty sets sized by the same reliability rule, solved by the same algorithm and evaluated on the same out-of-sample days. The contributions are:

(1) A two-stage robust model of a 69-bus ADN with micro-turbine commitment and reserves, demand-response reserve, batteries, renewable curtailment, zonal load shedding, grid-exchange deviations and LinDistFlow voltage limits, with complete recourse through penalized slacks.

(2) A common polyhedral family containing the independent budget set, a principal-component budget set on the leading correlated modes of the errors, a variant with residual coordinates, a joint-coverage set sized by split-conformal calibration [30], and a PCA box as in [8]; a rule that calibrates the budget on validation days for a target shedding frequency; and a C&CG implementation with a sub-problem that is exact for every member, plus a variant that uses the set only for robust feasibility with a sample-average cost.

(3) An out-of-sample study on 2,000 days, including larger errors, showing that under min–max the independent budget set trades off cost and shedding better than the principal-component sets, whose worst cases are common-mode energy deficits hedged with day-ahead energy that is curtailed on most days, and that the principal-component set gives the best schedule of all when used only for robust feasibility.

Section 2 presents the scheduling model, Section 3 the uncertainty sets and their calibration, Section 4 the solution method, and Section 5 the case study; Sections 6 and 7 discuss limitations and conclude.

2. Two-stage robust scheduling model

This section presents the uncertain quantities, the first-stage (day-ahead) and second-stage (re-dispatch) decisions, and the two-stage robust problem.

2.1 System and uncertainty

The ADN is a radial network with buses $n \in \mathcal{N}$ (bus 0 is the substation) and hours $t \in \mathcal{T}$, $|\mathcal{T}| = 24$. It hosts PV plants $i \in \mathcal{I}$ and wind turbines $j \in \mathcal{J}$ with day-ahead availability forecasts $\widehat{P}^{\mathrm{pv}}_{it}$, $\widehat{P}^{\mathrm{w}}_{jt}$, micro-turbines (MTs) $g \in \mathcal{G}$, batteries $b \in \mathcal{B}$, interruptible loads $d \in \mathcal{D}$, and load zones $z \in \mathcal{Z}$ with forecast demand $\widehat{L}_{zt}$. The load of zone z is spread over its buses with fixed shares χ_{nzt}. The forecast errors are collected in

$$\boldsymbol{u} = \left(u^{\mathrm{pv}}_{it},\, u^{\mathrm{w}}_{jt},\, u^{\mathrm{L}}_{zt}\right) \in \mathbb{R}^{N_u}, \qquad N_u = (|\mathcal{I}| + |\mathcal{J}| + |\mathcal{Z}|)\,|\mathcal{T}|, \tag{1}$$

so that the realized availabilities and loads are $\widehat{P}^{\mathrm{pv}}_{it} + u^{\mathrm{pv}}_{it}$, $\widehat{P}^{\mathrm{w}}_{jt} + u^{\mathrm{w}}_{jt}$ and $\widehat{L}_{zt} + u^{\mathrm{L}}_{zt}$.

2.2 First stage: day-ahead decisions

The first stage decides the MT commitment $x^{\mathrm{on}}_{gt} \in \{0,1\}$ and start-ups $x^{\mathrm{su}}_{gt} \geq x^{\mathrm{on}}_{gt} - x^{\mathrm{on}}_{g,t-1}$, the MT set-points p^{0}_{gt} with up and down reserves $r^{+}_{gt}, r^{-}_{gt} \leq R_g$, the day-ahead purchase $P^{\mathrm{DA}}_{t} \in [-\overline{S}, \overline{S}]$, and the DR reserve capacity $r^{\mathrm{dr}}_{dt} \in [0, \overline{D}_d]$:

$$\underline{P}_g\, x^{\mathrm{on}}_{gt} \leq p^{0}_{gt} - r^{-}_{gt}, \qquad p^{0}_{gt} + r^{+}_{gt} \leq \overline{P}_g\, x^{\mathrm{on}}_{gt}. \tag{2}$$

Its cost is $\boldsymbol{c}^{\top}\boldsymbol{x} = \sum_t \lambda^{\mathrm{DA}}_t P^{\mathrm{DA}}_t + \sum_{g,t}\left(c^{\mathrm{nl}}_g x^{\mathrm{on}}_{gt} + c^{\mathrm{su}}_g x^{\mathrm{su}}_{gt} + c^{\mathrm{res}}(r^{+}_{gt} + r^{-}_{gt})\right) + c^{\mathrm{cap}} \sum_{d,t} r^{\mathrm{dr}}_{dt}$.

2.3 Second stage: re-dispatch after the errors are known

Given $\boldsymbol{x}$ and $\boldsymbol{u}$, the second stage chooses MT outputs p_{gt}, battery charge and discharge $p^{\mathrm{c}}_{bt}, p^{\mathrm{d}}_{bt}$ with state of charge e_{bt}, DR deployment δ_{dt}, curtailment $\gamma^{\mathrm{pv}}_{it}$, γ^{w}_{jt}, zonal load shedding ℓ_{zt}, and grid deviations $\Delta^{+}_t, \Delta^{-}_t$:

$$p^{0}_{gt} - r^{-}_{gt} \leq p_{gt} \leq p^{0}_{gt} + r^{+}_{gt}, \qquad |p_{gt} - p_{g,t-1}| \leq R_g, \tag{3}$$

$$0 \leq p^{\mathrm{c}}_{bt}, p^{\mathrm{d}}_{bt} \leq \overline{P}_b, \qquad e_{bt} = e_{b,t-1} + \eta p^{\mathrm{c}}_{bt} - p^{\mathrm{d}}_{bt}/\eta, \qquad \underline{E}_b \leq e_{bt} \leq \overline{E}_b, \qquad e_{b,24} \geq e_{b,0}, \tag{4}$$

$$0 \leq \delta_{dt} \leq r^{\mathrm{dr}}_{dt}, \qquad 0 \leq \gamma^{\mathrm{pv}}_{it} \leq \widehat{P}^{\mathrm{pv}}_{it}, \qquad 0 \leq \gamma^{\mathrm{w}}_{jt} \leq \widehat{P}^{\mathrm{w}}_{jt}, \qquad 0 \leq \ell_{zt} \leq \widehat{L}_{zt},$$
$$0 \leq \Delta^{\pm}_t \leq \overline{\Delta}, \qquad |P^{\mathrm{DA}}_t + \Delta^{+}_t - \Delta^{-}_t| \leq \overline{S}, \tag{5}$$

$$P_t^{\mathrm{DA}} + \Delta_t^{+} - \Delta_t^{-} + \sum_g p_{gt} + \sum_b (p_{bt}^{\mathrm{d}} - p_{bt}^{\mathrm{c}}) + \sum_i (\widehat{P}_{it}^{\mathrm{pv}} + u_{it}^{\mathrm{pv}} - \gamma_{it}^{\mathrm{pv}}) + \sum_j (\widehat{P}_{jt}^{\mathrm{w}} + u_{jt}^{\mathrm{w}} - \gamma_{jt}^{\mathrm{w}}) + \sum_d \delta_{dt} + \sum_z \ell_{zt} + s_t^{+} - s_t^{-} = \sum_z (\hat{L}_{zt} + u_{zt}^{\mathrm{L}}), \tag{6}$$

$$\underline{V}^2 - \sigma_{nt}^{-} \le v_{nt} \le \overline{V}^2 + \sigma_{nt}^{+}, \qquad v_{nt} = v_0^2 - \frac{2}{S_{\mathrm{b}}} \sum_{m \in \mathcal{N}} \left(R_{nm} p_{mt}^{\mathrm{net}} + X_{nm} q_{mt} \right). \tag{7}$$

The net active load p_{mt}^{net} of bus m contains its share of the realized zonal load minus shedding, and subtracts the MT output, battery discharge minus charge, DR deployment and realized PV/wind output minus curtailment located at m; q_{mt} is the forecast reactive load (DERs at unity power factor), and R_{nm}, X_{nm} are the path resistances and reactances of the LinDistFlow model [31]. The balance slacks $s_t^{\pm}$ and voltage slacks $\sigma_{nt}^{\pm}$ make the recourse feasible for every $(\boldsymbol{x}, \boldsymbol{u})$ (complete recourse) and are penalized heavily. The second-stage cost is

$$Q(\boldsymbol{x}, \boldsymbol{u}) = \min \; \sum_{g,t} c_g p_{gt} + c^{\mathrm{bat}} \sum_{b,t} (p_{bt}^{\mathrm{c}} + p_{bt}^{\mathrm{d}}) + c^{\mathrm{dr}} \sum_{d,t} \delta_{dt} + c^{\mathrm{cu}} \sum_t \left(\sum_i \gamma_{it}^{\mathrm{pv}} + \sum_j \gamma_{jt}^{\mathrm{w}} \right) + c^{\mathrm{sh}} \sum_{z,t} \ell_{zt} + \sum_t (\lambda_t^{+} \Delta_t^{+} - \lambda_t^{-} \Delta_t^{-}) + c^{\mathrm{s}} \sum_t (s_t^{+} + s_t^{-}) + c^{\mathrm{v}} \sum_{n,t} (\sigma_{nt}^{+} + \sigma_{nt}^{-}), \tag{8}$$

where upward deviations are bought at $\lambda_t^{+} > \lambda_t^{\mathrm{DA}}$ and downward deviations sold at $\lambda_t^{-} < \lambda_t^{\mathrm{DA}}$. Two modelling choices keep the recourse linear in $\boldsymbol{u}$: curtailment and shedding are bounded by the forecasts, and MT reserve is fixed in the first stage; the out-of-sample evaluation of Section 5 uses the realized bounds.

In compact form, with $\boldsymbol{y} \ge \boldsymbol{0}$ collecting the second-stage variables,

$$Q(\boldsymbol{x}, \boldsymbol{u}) = \min_{\boldsymbol{y} \ge \boldsymbol{0}} \left\{ \boldsymbol{q}^{\top} \boldsymbol{y} : \; \boldsymbol{A}\boldsymbol{y} \le \boldsymbol{b} - \boldsymbol{B}\boldsymbol{x} - \boldsymbol{E}\boldsymbol{u}, \; \boldsymbol{A}^{\mathrm{eq}} \boldsymbol{y} = \boldsymbol{b}^{\mathrm{eq}} - \boldsymbol{B}^{\mathrm{eq}} \boldsymbol{x} - \boldsymbol{E}^{\mathrm{eq}} \boldsymbol{u} \right\}, \tag{9}$$

and the two-stage robust problem for an uncertainty set $\mathcal{U}$ is

$$\min_{\boldsymbol{x} \in \mathcal{X}} \; \boldsymbol{c}^{\top} \boldsymbol{x} + \max_{\boldsymbol{u} \in \mathcal{U}} Q(\boldsymbol{x}, \boldsymbol{u}), \tag{10}$$

where $\mathcal{X}$ collects (2), the start-up logic and the bounds of the first stage. $Q(\boldsymbol{x}, \cdot)$ is convex and piecewise linear, so the inner maximum over a polytope $\mathcal{U}$ is attained at a vertex.

3. Uncertainty sets and their sizing

This section defines the compared sets in a common polyhedral form, the rule that sizes them for a reliability target, and how correlation affects their worst cases.

3.1 A common polyhedral form

All sets considered here have the form

$$\mathcal{U}(\boldsymbol{\Phi}, \rho, \Gamma) = \left\{ \boldsymbol{u} = \boldsymbol{\Phi}\boldsymbol{\zeta} : \; \|\boldsymbol{\zeta}\|_{\infty} \le \rho, \; \|\boldsymbol{\zeta}\|_1 \le \Gamma \right\}, \tag{11}$$

where the columns of $\boldsymbol{\Phi} \in \mathbb{R}^{N_u \times k}$ are the directions in which the errors may deviate, ρ bounds each coordinate and Γ the total deviation ($\Gamma = \infty$ gives a box); the sets differ in $\boldsymbol{\Phi}$ and in how ρ and Γ are chosen. Ellipsoidal sets are excluded: although convex problems with their separable quadratic

constraints are solved efficiently [32], the worst case of the recourse over an ellipsoid is not attained at one of finitely many vertices, which the exact sub-problem of Section 4 exploits. Let $\boldsymbol{U}^{\mathrm{tr}} \in \mathbb{R}^{n \times N_u}$ contain n days of historical forecast errors, $\boldsymbol{\sigma}$ their standard deviations, and $\boldsymbol{\Sigma} = \boldsymbol{V}\,\mathrm{diag}(\boldsymbol{s})^2\boldsymbol{V}^{\top}$ the eigendecomposition of their covariance with eigenvalues in decreasing order.

3.1.1 Independent budget (IB)

$\boldsymbol{\Phi} = \mathrm{diag}(\boldsymbol{\sigma})$ restricted to the k coordinates with non-zero variance (PV at night has none), $\rho = 3$ and $\Gamma = 3t\sqrt{k}$. This is the budget set of [15] with deviations of up to three standard deviations, at most $t\sqrt{k}$ of which can be at their bounds simultaneously. For independent, symmetric coordinates bounded by ρ standard deviations, a constraint protected against this set is violated with probability at most $\exp(-t^2/2)$ [15], so $t = \sqrt{2\ln(1/\varepsilon)}$ targets a violation probability ε; the bound fails for correlated or heavier-tailed coordinates.

3.1.2 Principal-component budget (PK)

$\boldsymbol{\Phi} = \boldsymbol{V}_K\,\mathrm{diag}(\boldsymbol{s}_K)$ with the first K principal components, so that $\boldsymbol{\zeta} = \mathrm{diag}(\boldsymbol{s}_K)^{-1}\boldsymbol{V}_K^{\top}\boldsymbol{u}$ are the standardized, uncorrelated scores of the error vector; $\rho = 3$ and $\Gamma = 3t\sqrt{K}$. The leading components are the correlated error modes: a cloud lowering all PV plants for several hours, a wind front, a temperature deviation of all load zones. The set lets $t\sqrt{K}$ of these modes reach three standard deviations together and excludes deviations outside their span, leaving the residual to the recourse and the sizing rule.

3.1.3 Principal components plus residual (PKR)

$\boldsymbol{\Phi} = [\boldsymbol{V}_K\,\mathrm{diag}(\boldsymbol{s}_K)\,,\ \mathrm{diag}(\boldsymbol{\sigma}^{\mathrm{res}})]$, where $\boldsymbol{\sigma}^{\mathrm{res}}$ are the non-zero standard deviations of the residual $\boldsymbol{u} - \boldsymbol{V}_K\boldsymbol{V}_K^{\top}\boldsymbol{u}$; $\rho = 3$ and $\Gamma = 3t\sqrt{k}$ with k the number of columns of $\boldsymbol{\Phi}$. The K modes and the residual deviations share one budget, so PKR also covers the deviations that PK excludes.

3.1.4 Joint-coverage sizing (JC)

The IB shape with $(\rho, \Gamma) = \theta(\rho_0, \Gamma_0)$, where ρ_0 and Γ_0 are the 99th percentiles of $\|\boldsymbol{\zeta}\|_\infty$ and $\|\boldsymbol{\zeta}\|_1$ on the training days, and θ is chosen by split-conformal calibration [30]: with the gauge $g(\boldsymbol{u}) = \max\{\|\boldsymbol{\zeta}\|_\infty/\rho_0, \|\boldsymbol{\zeta}\|_1/\Gamma_0\}$ evaluated on n_{c} separate calibration days, θ is the $\lceil (n_{\mathrm{c}}+1)(1-\alpha)\rceil$-th smallest value. For exchangeable days, a new day then lies in the set with probability at least $1-\alpha$ [30].

3.1.5 PCA box (PX)

The first K components with a box ($\Gamma = \infty$) sized by the same conformal rule for the projected scores, following the flexibility-demand polytopes of [8].

3.2 Sizing for a reliability target

Joint coverage is demanding in high dimension: a day lies in the set only if all of its N_u errors do. The operator's actual requirement is rather that load shedding occurs on at most a fraction ε of days. For IB and PK we therefore choose t by validation: for a grid of values, the robust schedule is computed on the training data and its recourse simulated on validation days, and the smallest t whose validation shedding frequency does not exceed ε is selected and reported on separate test days. Applying the same rule to both shapes gives each the size it needs to deliver the target.

3.3 Why correlation matters for the worst case

For a single constraint $\boldsymbol{a}^{\top}\boldsymbol{u} \le \beta$, the protection level of a set (11) is its support function $h(\boldsymbol{a}) = \max_{\boldsymbol{u}\in\mathcal{U}} \boldsymbol{a}^{\top}\boldsymbol{u}$, i.e., ρ times the sum of the $\lfloor \Gamma/\rho \rfloor$ largest entries of $|\boldsymbol{\Phi}^{\top}\boldsymbol{a}|$ plus a fractional term; for a Gaussian error the $(1-\varepsilon)$-quantile of $\boldsymbol{a}^{\top}\boldsymbol{u}$ is $z_{1-\varepsilon}\sqrt{\boldsymbol{a}^{\top}\boldsymbol{\Sigma}\boldsymbol{a}}$. If the errors of the six PV plants in one hour are positively correlated, the aggregate PV error has a variance close to $(\sum_i \sigma_i)^2$ rather than $\sum_i \sigma_i^2$; IB reaches it only by spending six units of budget on one hour, whereas PK spends one unit on the component that moves all six plants together. Conversely, IB allows combinations that the data never show (east plants at $+3\sigma$, west plants at -3σ), which PK excludes. Which effect matters depends on the binding constraints and the available first-stage hedges; Section 5 measures it.

4. Solution by column-and-constraint generation

This section describes the master problem, the exact worst-case sub-problem with its local search, and a variant that uses the set only for robust feasibility.

4.1 Master problem

C&CG [12] keeps a list of scenarios $\{\boldsymbol{u}^{1}, \dots, \boldsymbol{u}^{m}\} \subset \mathcal{U}$ (starting with $\boldsymbol{u}^{1} = \boldsymbol{0}$) and solves the mixed-integer master problem

$$\mathrm{LB} = \min_{\boldsymbol{x}\in\mathcal{X},\, \eta,\, \boldsymbol{y}^{1},\dots,\boldsymbol{y}^{m} \ge \boldsymbol{0}} \Big\{ \boldsymbol{c}^{\top}\boldsymbol{x} + \eta :\ \eta \ge \boldsymbol{q}^{\top}\boldsymbol{y}^{l},\ \boldsymbol{y}^{l} \text{ feasible in (9) for } (\boldsymbol{x}, \boldsymbol{u}^{l}),\ l = 1, \dots, m \Big\}, \quad (12)$$

whose optimal value is a lower bound on (10) because only m scenarios are protected.

4.2 Worst-case sub-problem

For the master's first-stage solution $\boldsymbol{x}$, the sub-problem computes $\max_{\boldsymbol{u}\in\mathcal{U}} Q(\boldsymbol{x}, \boldsymbol{u})$. By LP duality of (9),

$$\max_{\boldsymbol{u}\in\mathcal{U}} Q(\boldsymbol{x}, \boldsymbol{u}) = \max_{\boldsymbol{u}\in\mathcal{U},\, (\boldsymbol{\pi},\boldsymbol{\nu})\in\Pi} (\boldsymbol{b} - \boldsymbol{B}\boldsymbol{x})^{\top}\boldsymbol{\pi} + (\boldsymbol{b}^{\mathrm{eq}} - \boldsymbol{B}^{\mathrm{eq}}\boldsymbol{x})^{\top}\boldsymbol{\nu} - \boldsymbol{u}^{\top}\boldsymbol{w}, \qquad \boldsymbol{w} = \boldsymbol{E}^{\top}\boldsymbol{\pi} + \boldsymbol{E}^{\mathrm{eq}\top}\boldsymbol{\nu}, \quad (13)$$

with $\Pi = \{\boldsymbol{A}^{\top}\boldsymbol{\pi} + \boldsymbol{A}^{\mathrm{eq}\top}\boldsymbol{\nu} \le \boldsymbol{q},\ \boldsymbol{\pi} \le \boldsymbol{0}\}$. The objective is bilinear in $\boldsymbol{u}$ and the duals. Because the

maximum over the polytope (11) is attained at a vertex, and the vertices have $n_\rho = \lfloor \Gamma/\rho \rfloor$ coordinates at $\pm\rho$ and at most one at $\pm r$ with $r = \Gamma - n_\rho \rho$, we write

$$\zeta_k = \rho(z_k^+ - z_k^-) + r(f_k^+ - f_k^-), \qquad \sum_k (z_k^+ + z_k^-) \le n_\rho, \qquad \sum_k (f_k^+ + f_k^-) \le 1, \qquad z_k^+ + z_k^- + f_k^+ + f_k^- \le 1, \tag{14}$$

with binary $z_k^\pm, f_k^\pm$. Every point of (14) lies in $\mathcal{U}$ and every vertex of $\mathcal{U}$ is represented, so the formulation is exact. The products of the binaries with $\omega = \boldsymbol{\Phi}^\top \boldsymbol{w}$ are linearized exactly with bounds $[\underline{\omega}_k, \overline{\omega}_k]$, which are finite because every row that $\boldsymbol{u}$ enters (balance and voltage) has a penalized slack that bounds its dual; one pair of LPs over Π per column of $\boldsymbol{\Phi}$ computes them, giving much tighter big-M values than interval arithmetic when $\boldsymbol{\Phi}$ is dense.

4.3 Hybrid sub-problem and termination

The mixed-integer sub-problem is the expensive step. Since $Q(\boldsymbol{x},\cdot)$ is convex, a local search climbs to a vertex quickly: from a start $\boldsymbol{u}$, it solves the recourse LP, takes its dual vector $\boldsymbol{w}$ ($-\boldsymbol{w}$ is a subgradient of Q), moves to the vertex that maximizes $-\boldsymbol{w}^\top \boldsymbol{u}$ over $\mathcal{U}$, found for (11) by sorting $|\boldsymbol{\Phi}^\top \boldsymbol{w}|$, and repeats until the value stops increasing. We run it from 50 random starts and from all scenarios found so far. If the best local value Q^{h} exceeds the master's bound by more than the tolerance, the scenario is added to the master directly. Otherwise the mixed-integer sub-problem is solved (time limit 300 s) from the local solution; its incumbent $\underline{Q}$ is the worst case found and its best bound $\overline{Q}$ certifies $\max_{\boldsymbol{u}\in\mathcal{U}} Q(\boldsymbol{x},\boldsymbol{u}) \le \overline{Q}$. If $\boldsymbol{c}^\top \boldsymbol{x} + \underline{Q}$ exceeds LB by more than 10^{-3} relative, the incumbent is added as a scenario; otherwise, or after six mixed-integer solves, the algorithm returns $\boldsymbol{x}$ with LB, the worst case found and the certified upper bound $\boldsymbol{c}^\top \boldsymbol{x} + \overline{Q}$. Most iterations thus use only LPs, and every reported bound is certified.

4.4 Robust feasibility with a sampled cost

The min–max problem (10) uses the set twice: to require a recourse for every $\boldsymbol{u} \in \mathcal{U}$ and to price the schedule by its worst case. The robust-feasibility variant (RFS) keeps the first use and prices by a sample average. With $\boldsymbol{q}^{\mathrm{p}}$ the part of $\boldsymbol{q}$ that prices load shedding and the balance and voltage slacks, and training days $\boldsymbol{u}^1, \dots, \boldsymbol{u}^N$, RFS solves

$$\min_{\boldsymbol{x}\in\mathcal{X}} \ \boldsymbol{c}^\top \boldsymbol{x} + \tfrac{1}{N}\textstyle\sum_{s=1}^{N} Q(\boldsymbol{x},\boldsymbol{u}^s) \quad \text{s.t.} \quad \min_{\boldsymbol{y}} \left\{ (\boldsymbol{q}^{\mathrm{p}})^\top \boldsymbol{y} : \ \boldsymbol{y} \text{ feasible in (9) for } (\boldsymbol{x},\boldsymbol{u}) \right\} = 0 \quad \forall \boldsymbol{u} \in \mathcal{U}, \tag{15}$$

i.e., it minimizes the sample-average cost among the schedules that avoid shedding and limit violations for every error in the set; in the unified model of [27], it puts all the weight on the expected cost and keeps the set only for feasibility. C&CG applies with two changes: each scenario added to the master carries a copy of the recourse with its penalized variables fixed to zero, and the sub-problem maximizes the penalty cost (the mixed-integer program above with $\boldsymbol{q}^{\mathrm{p}}$ instead of $\boldsymbol{q}$). The local search is of little use because the penalty is zero, with a vanishing subgradient, on most of $\mathcal{U}$,

so every iteration solves the mixed-integer program. The algorithm stops when the certified maximal penalty is below 1 $ or after six mixed-integer solves.

5. Case study

This section describes the test system, data and compared schedules, and then reports the results.

5.1 Network, resources and prices

The test system is the IEEE 69-bus feeder [33] (3.80 MW, 2.69 Mvar base load) with the resources of Table 1. The load is split into a residential zone (main feeder and laterals at buses 66–69), a commercial zone (laterals 28–52) and an industrial zone (lateral 53–65, with the 1.24-MW load of bus 61), each with its own daily shape; the peak total load is 3.45 MW. Six PV plants (3.05 MW, 88% of the peak load), two wind turbines (0.7 MW), two MTs and two batteries are located as in Table 1. The substation voltage is 1.03 p.u., and bus voltages must stay in $[0.95,1.05]$ p.u. The day-ahead price follows a time-of-use pattern, and the deviation limit of 0.4 MW per hour makes the local resources responsible for most of the balancing.

Table 1. Resources of the 69-bus ADN and cost parameters

Resource	Data
Grid exchange	substation 5 MW; day-ahead price 40/70/110 $/MWh; deviations up to 0.4 MW per hour, bought at $1.5\lambda_t^{\mathrm{DA}} + 20$ and sold at $0.5\lambda_t^{\mathrm{DA}}$ $/MWh
PV (bus: MW)	12: 0.40, 21: 0.40, 35: 0.30, 50: 0.65, 61: 0.80, 64: 0.50; forecast 0.8 × clear-sky profile
Wind (bus: MW)	27: 0.40, 46: 0.30; forecast capacity factor 0.42 ± 0.12 (night peak)
Micro-turbines	bus 61: 0.24–0.80 MW, 75 $/MWh; bus 18: 0.15–0.50 MW, 80 $/MWh; ramp 50%/h; no-load 12 / 8 $/h, start-up 30 / 20 $, reserve capacity 4 $/MW-h
Batteries	bus 64: 0.5 MW / 1.5 MWh; bus 27: 0.3 MW / 0.9 MWh; $\eta = 0.95$, SOC 10–100%, 4 $/MWh
Demand response	bus 61: 0.25 MW, bus 49: 0.10 MW, bus 50: 0.10 MW; capacity 8 $/MW-h, deployment 150 $/MWh
Penalties	curtailment 40 $/MWh, shedding 2000 $/MWh, balance slack 5000 $/MWh, voltage slack 2×10^4 $ per p.u.^{2}-h

5.2 Forecast-error data

Forecast errors come from a factor model that mimics real multi-site errors. A cloud factor drives all PV plants, a wind factor both turbines and a temperature factor all load zones; the factors are cross-correlated (cloud and load: -0.35), follow first-order autoregressive processes over the day ($\phi = 0.85$) and have Student-t innovations with five degrees of freedom. Each source adds autoregressive idiosyncratic noise ($\phi = 0.6$), with loadings of 0.9 (PV), 0.8 (wind) and 0.7 (load) on the common factor. PV errors scale with the forecast (relative standard deviation 22%), wind errors with the capacity (14%) and load errors with the forecast load (5%). Realized PV and wind outputs are clipped to [0,capacity], which skews the errors near the bounds. With 6 PV plants, 2 turbines and 3 zones over 24 hours, $N_u = 264$, of which 198 have non-zero variance. In the generated data, same-hour errors are correlated at about 0.8 between PV plants, 0.6 between turbines, 0.5 between load zones and -0.2 between PV and load. The first 10, 20 and 40 principal components explain 64%, 75% and 85% of the variance.

We use four independent data sets: 365 training days (set fitting), 365 calibration days (conformal

sizing), 500 validation days (reliability sizing of t) and 2,000 test days for all reported results. A second test set of 2,000 days with errors scaled by 1.25 represents a forecast degradation absent from the training data.

5.3 Compared schedules

All robust schedules are computed by C&CG on the 365 training days: IB and PK with $K = 20$ for $t \in \{0.05,0.1,0.2,0.35,0.5,1.0\}$ (IB also for $t = 0.7$), and PK with $K = 10$ and $K = 40$ and PKR with $K = 20$ for selected t. JC and PX use $\alpha = 5\%$ and the calibration days. The two non-robust benchmarks are a deterministic schedule for the forecast with MT and DR upward reserve of at least 10% of the forecast load plus renewable output in every hour and half of that as downward reserve (DET), and a two-stage stochastic program over 50 equally weighted training days (SAA). RFS uses 30 training days for the cost and the IB or PK set for feasibility. Every schedule is evaluated by solving the recourse LP with realized bounds on each validation and test day; a shedding day is one with more than 0.1 kWh of shed load or a used balance slack.

5.4 Sets and computation

Table 2 lists the size of every set and the C&CG statistics. Apart from PK with $t = 1.0$ (27.6%), no budget set contains a single test day, whereas JC and PX contain 95.0% and 95.9%, close to their 95% target. Coverage and reliability are thus unrelated here: IB with $t = 1.0$ contains no test day, yet its schedule sheds load on only one test day in 2,000.

C&CG always stopped with a worst case within 0.1% of the master's lower bound. With PK, the sub-problem has only $K = 10$–40 pairs of binaries, and for $t \leq 0.5$ the mixed-integer program also certified the worst case (gap below 0.1%) within 60 s. With IB it has 198 pairs; for $t \geq 0.2$ its bound stayed 25–39% above the incumbent after 300 s, as for PKR with $t = 0.2$, PK with $t = 1.0$, JC and PX. For these sets the reported worst case is the best one found, and the schedules are validated out of sample. The robust objective poorly estimates the realized cost: for the budget sets it lies between 12% below and 14% above the out-of-sample mean, and for JC and PX it is 3.3 and 1.7 times the mean.

Table 2. Uncertainty sets and C&CG statistics

Schedule	k	ρ	Γ	Coverage (%)	Robust obj. ($)	Gap to worst (%)	Certified gap (%)	Scen-arios	Time (s)
IB, $t = 0.05$	198	3.00	2.1	0.0	3185	0.04	0.0	8	26
IB, $t = 0.1$	198	3.00	4.2	0.0	3216	0.09	0.1	8	189
IB, $t = 0.2$	198	3.00	8.4	0.0	3277	0.02	29.8	6	323
IB, $t = 0.35$	198	3.00	14.8	0.0	3354	0.00	31.5	9	339
IB, $t = 0.5$	198	3.00	21.1	0.0	3421	0.06	25.3	12	400
IB, $t = 0.7$	198	3.00	29.5	0.0	3527	0.03	39.0	13	693
IB, $t = 1.0$	198	3.00	42.2	0.0	3635	0.10	30.6	17	509
PK, $t = 0.05$	20	3.00	0.7	0.0	3328	0.00	0.0	2	7
PK, $t = 0.1$	20	3.00	1.3	0.0	3500	0.00	0.0	2	12
PK, $t = 0.2$	20	3.00	2.7	0.0	3861	0.01	0.0	4	20

Schedule	k	ρ	Γ	Coverage (%)	Robust obj. ($)	Gap to worst (%)	Certified gap (%)	Scen-arios	Time (s)
PK, $t = 0.35$	20	3.00	4.7	0.0	4123	0.00	0.0	5	33
PK, $t = 0.5$	20	3.00	6.7	0.0	4326	0.02	0.0	6	59
PK, $t = 1.0$	20	3.00	13.4	27.6	4734	0.05	93.1	12	400
PK_{10}, $t = 0.2$	10	3.00	1.9	0.0	3642	0.04	0.0	3	9
PK_{40}, $t = 0.2$	40	3.00	3.8	0.0	4029	0.07	0.1	5	27
PKR, $t = 0.1$	218	3.00	4.4	0.0	4095	0.08	0.1	5	124
PKR, $t = 0.2$	218	3.00	8.9	0.0	4527	0.01	91.3	7	364
JC (95% coverage)	198	4.86	224.4	95.0	11730	0.00	33.6	34	1261
PX (95% coverage)	20	3.03	∞	95.9	6660	0.00	98.1	19	470
RFS, IB, $t = 0.5$	198	3.00	21.1	0.0	worst penalty 0.05 (bound 1740.12)			4	1620
RFS, IB, $t = 1.0$	198	3.00	42.2	0.0	worst penalty 1289.86 (bound 2334.25)			5	1903
RFS, PK, $t = 0.2$	20	3.00	2.7	0.0	worst penalty 0.00 (bound 0.00)			1	56
RFS, PK, $t = 0.5$	20	3.00	6.7	0.0	worst penalty 0.20 (bound 0.20)			5	167
IB, $t = 0.2$, weak corr.	198	3.00	8.4	0.0	3276	0.00	36.0	6	320
PK, $t = 0.2$, weak corr.	20	3.00	2.7	0.0	3523	0.00	0.0	3	9

Note: Coverage: share of test days inside the set. Gap to worst: relative gap between the master's lower bound and the worst case found; certified gap: between the lower bound and the certified upper bound. RFS rows: worst penalty found and its certified bound ($).

5.5 Out-of-sample cost and reliability

Table 3 and Figure 1 report the out-of-sample results. In Figure 1, the shedding axis is logarithmic beyond 0.1%, labels give the budget parameter t, and the RFS curves connect $t = 0.5$ and 1.0 (IB) and $t = 0.2$ and 0.5 (PK). DET sheds load on 8.1% of the test days; SAA has the lowest mean cost, 3,352 $/day, and sheds on 3.0%.

As t grows, IB and PK trace two cost–reliability curves. Along the IB curve, shedding days fall from 13.5% ($t = 0.05$) to 0.05% ($t = 1.0$), while the mean cost first falls from 3,613 to 3,392 $/day at $t = 0.5$, as less load is shed at 2,000 $/MWh, and then rises to 3,401 and 3,441 $/day at $t = 0.7$ and 1.0. Along the PK curve, shedding days fall from 9.6% to 0% while the mean cost rises from 3,485 to 4,159 $/day. The IB curve lies below the PK curve at every shedding level: at 0.05%, IB with $t = 1.0$ costs 14.7% less than PK with $t = 0.5$, and it dominates PK with $t = 0.2$ (3,441 against 3,759 $/day; 0.05% against 0.70%). PK with $K = 10$ or 40 and PKR lie on or above the PK curve. JC and PX avoid shedding on all test days at a mean cost 3.0% and 10.9% above IB with $t = 1.0$, which costs 2.7% more than SAA and reduces the shedding days from 59 to 1 in 2,000.

With errors scaled by 1.25, every schedule sheds more often and the ranking is unchanged (Figure 1b): IB with $t = 1.0$ sheds on 0.45% of the days at 3,502 $/day, PK with $t = 0.5$ on 0.35% at 4,049 $/day, JC on 0.05% at 3,590 $/day, and SAA on 6.0%, twice as often as before. With halved factor loadings (weaker correlation, last rows of Table 3), IB and PK with $t = 0.2$ shed on 0.80% and

0.50% of the days, and PK costs 5.4% more than IB instead of 8.8%.

Table 3. Out-of-sample results on 2,000 test days (costs in $/day)

Schedule	Test days						Errors × 1.25	
	First stage	Mean cost	95% quantile	Shedding days (%)	Exp. shed (kWh)	Curtail. (MWh)	Mean cost	Shedding days (%)
DET (10% reserve)	2473	3577	4283	8.05	28.1	0.17	3727	17.50
SAA (50 days)	2575	3352	3989	2.95	14.5	0.02	3491	6.00
IB, $t = 0.05$	2972	3613	5860	13.50	140.4	0.16	3982	22.55
IB, $t = 0.1$	2868	3542	5333	10.90	104.7	0.14	3873	19.35
IB, $t = 0.2$	2955	3457	4324	7.35	59.5	0.16	3717	14.75
IB, $t = 0.35$	2980	3398	4010	2.55	26.5	0.14	3588	7.50
IB, $t = 0.5$	3068	3392	3950	1.60	11.1	0.15	3541	4.80
IB, $t = 0.7$	2976	3401	3913	0.35	0.7	0.14	3481	1.45
IB, $t = 1.0$	3076	3441	3873	0.05	0.0	0.24	3502	0.45
PK, $t = 0.05$	3054	3485	4183	9.55	50.5	0.54	3715	17.85
PK, $t = 0.1$	3225	3516	3911	4.35	15.6	1.25	3638	8.45
PK, $t = 0.2$	3288	3759	3917	0.70	1.4	3.19	3802	2.60
PK, $t = 0.35$	3552	3918	4048	0.25	0.4	4.14	3940	0.70
PK, $t = 0.5$	3758	4035	4159	0.05	0.2	4.61	4049	0.35
PK, $t = 1.0$	4008	4159	4333	0.00	0.0	3.61	4176	0.05
PK_{10}, $t = 0.2$	3367	3617	3838	2.65	5.5	2.08	3698	5.85
PK_{40}, $t = 0.2$	3457	3858	3998	0.20	0.4	3.87	3886	0.95
PKR, $t = 0.1$	3524	3900	4030	0.15	0.4	4.06	3924	0.90
PKR, $t = 0.2$	3727	4079	4216	0.05	0.0	4.34	4095	0.25
JC (95% coverage)	2873	3546	3927	0.00	0.0	0.22	3590	0.05
PX (95% coverage)	3356	3817	4070	0.00	0.0	0.66	3854	0.00
RFS, IB, $t = 0.5$	2581	3353	3974	1.10	7.2	0.00	3478	3.70
RFS, IB, $t = 1.0$	2637	3366	3969	0.70	1.8	0.01	3463	2.70
RFS, PK, $t = 0.2$	2541	3362	4050	2.90	18.1	0.01	3516	6.75
RFS, PK, $t = 0.5$	2701	3371	3928	0.05	0.2	0.01	3439	0.70
IB, $t = 0.2$, weak corr.	2901	3273	3658	0.80	1.3	0.01	3328	2.60
PK, $t = 0.2$, weak corr.	3250	3450	3597	0.50	0.4	1.04	3478	1.15

Note: Shedding days: share of days with load shedding; expected shedding and renewable curtailment are daily means; last two columns: errors scaled by 1.25.

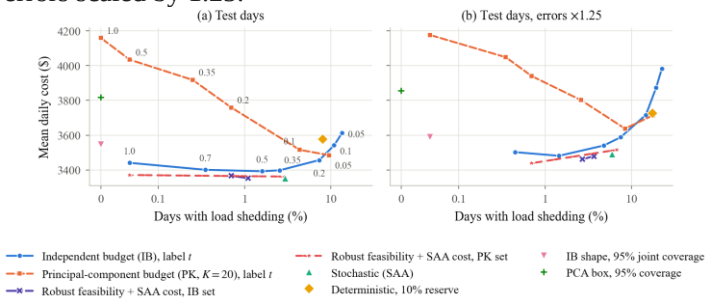


Figure 1. Mean daily cost against the share of days with load shedding: (a) test days; (b) errors scaled by 1.25

5.6 Why the principal-component sets cost more

Table 4 compares the first-stage decisions with the protection of the sets against three aggregate deficits: the daily net-load deficit $\sum_t(\sum_z u^{\mathrm{L}}_{zt} - \sum_i u^{\mathrm{pv}}_{it} - \sum_j u^{\mathrm{w}}_{jt})$, the PV deficit between 09:00 and 16:00, and the largest single-hour net-load deficit. The protection is the support function of the set in the corresponding direction (Section 3); the last row gives empirical quantiles on the test days.

The sets place their protection differently. A large aggregate deficit in IB requires many coordinates to deviate in the same direction, so its daily protection stays at 1.1–4.3 MWh for $t = 0.2$–1.0, below the 99th percentile of 8.9 MWh, while its single-hour protection reaches 2.24 MWh at $t = 1.0$, above the 99.9th percentile. PK spends its budget on correlated modes that move all PV plants or all load zones over several hours: at the shedding frequency of IB with $t = 1.0$, PK with $t = 0.5$ protects 13.9 MWh of daily and 7.5 MWh of midday PV deficit, 3.2 and 1.8 times as much as IB, but less single-hour deficit.

The first stage hedges these worst cases differently. As t grows, the PK schedules buy more day-ahead energy (44.0 MWh at $t = 0.05$ to 53.3 MWh at $t = 1.0$; SAA: 36.4 MWh), 12.8–15.5 MWh of it between 09:00 and 16:00 for $t \geq 0.2$ (IB: 8.3–9.6 MWh; SAA: 4.6 MWh); for $t \geq 0.2$ they commit the MTs for 11–15 h, and no PK schedule holds DR reserve. The IB schedules keep the purchase at 41–43 MWh and add MT upward reserve and DR capacity (up to 1.8 and 3.8 MWh at $t = 1.0$). The min–max criterion values a hedge at the worst case, where the hedged energy is always used. A PV deficit of several megawatt-hours over several hours would need MT commitment and deployment at 75–80 $/MWh plus no-load cost, and the DR capacity (0.45 MW) and deviation limit (0.4 MW) are too small, so day-ahead energy at 70–110 $/MWh is the cheapest cover. The sparse worst cases of IB, with a few coordinates at three standard deviations, are covered more cheaply by reserve at 4–8 $/MW-h.

Out of sample, reserve is paid as capacity and deployed only when needed, whereas the extra day-ahead energy is paid every day but needed only on days with a large PV deficit; otherwise the surplus is sold back at half the day-ahead price within the 0.4-MW deviation limit, and the rest displaces wind and PV output. The PK schedules with $t \geq 0.2$ curtail 3.2–4.6 MWh per day, 13–19% of the forecast renewable energy of 24.3 MWh, against 0.14–0.24 MWh for IB and 0.02 MWh for SAA. Curtailment and the higher first-stage cost account for the gap between the curves in Figure 1.

PK does not protect the wrong deviations: in the six schedules for which the hour of shedding was recorded (DET, SAA, and IB and PK with $t = 0.2$ and 0.5), 68–100% of the shed energy falls between 08:00 and 16:00, i.e., shedding is caused by PV shortfalls, the common mode that PK represents. The difference lies in how the min–max schedule hedges them. JC confirms this: with the IB shape but more daily protection (18.1 MWh) than any PK set, its schedule still hedges with reserve (27 committed MT-hours, 5.2 MWh of MT upward reserve, 5.8 MWh of DR capacity) rather than day-ahead energy (38.0 MWh), and avoids shedding on all test days at a lower mean cost than every PK schedule with $t \geq 0.2$.

Table 4. First-stage decisions and protection levels

	First stage and curtailment						Protection of the set		
Schedule	DA	DA 9–16 h	MT on	MT res.	DR	Curtail.	Daily	Midday PV	One hour
SAA (50 days)	36.4	4.6	15	1.7	0.8	0.02	–	–	–
IB, $t = 0.2$	42.5	8.8	5	0.3	0.5	0.16	1.1	1.1	0.89
IB, $t = 0.5$	43.0	9.6	6	0.8	1.2	0.15	2.4	2.4	1.72
IB, $t = 1.0$	42.0	9.2	11	1.8	3.8	0.24	4.3	4.2	2.24

	First stage and curtailment						Protection of the set		
Schedule	DA	DA 9–16 h	MT on	MT res.	DR	Curtail.	Daily	Midday PV	One hour
PK, $t = 0.2$	46.4	12.8	11	0.8	0.0	3.19	8.5	6.0	1.19
PK, $t = 0.5$	51.7	15.5	11	2.0	0.0	4.61	13.9	7.5	1.86
PK, $t = 1.0$	53.3	13.5	15	3.6	0.0	3.61	17.9	7.9	2.48
JC (95% coverage)	38.0	7.0	27	5.2	5.8	0.22	18.1	14.2	3.62
PX (95% coverage)	42.6	8.8	31	7.8	4.2	0.66	22.2	8.8	3.73
Test days							8.9 (11.5)	5.9 (7.5)	1.49 (1.98)

Note: DA: day-ahead purchase (total and 09:00–16:00); MT on: committed MT-hours; MT res.: MT upward reserve; DR: DR reserve capacity (MWh per day except MT-hours); curtailment: test-day mean. Protection: largest daily net-load, midday PV and single-hour net-load deficits in the set (MWh); last row: 99th (99.9th) percentiles on the test days.

5.7 Sizing for a reliability target

Table 5 applies the sizing rule of Section 3. On the test days, the IB selections meet all three targets, and the PK selections meet 5% and 1% but slightly exceed 0.2% (0.25%, five days instead of four), since 500 validation days resolve a frequency of 0.2% only to a single day. At every target, the IB selection is cheaper, by 3.3%, 9.5% and 12.2%, and curtails an order of magnitude less renewable energy. The a priori bound of [15] would suggest $t = \sqrt{2\ln(1/\varepsilon)} = 2.45$, 3.03 and 3.53, three to seven times the values selected on validation data, although it does not even hold for correlated, heavy-tailed errors.

Table 5. Reliability-targeted sizing

			Validation	Test days			Errors × 1.25	
ε (%)	Set	t	shedding (%)	Mean	Shedding (%)	Curtail.	Mean	Shedding (%)
5.0	IB	0.35	4.20	3398	2.55	0.14	3588	7.50
	PK	0.1	5.00	3516	4.35	1.25	3638	8.45
1.0	IB	0.7	0.60	3401	0.35	0.14	3481	1.45
	PK	0.2	1.00	3759	0.70	3.19	3802	2.60
0.2	IB	1.0	0.00	3441	0.05	0.24	3502	0.45
	PK	0.35	0.20	3918	0.25	4.14	3940	0.70

Note: For each target ε: the smallest t with a validation shedding frequency of at most ε, and its performance on the test days (costs in $/day, curtailment in MWh/day).

5.8 Robust feasibility with a sampled cost

The analysis suggests using a set only to define the deviations a schedule must withstand and pricing the schedule by its expected cost; the RFS rows of Tables 3 and 2 and Figure 1 report this variant, with the cost averaged over 30 training days.

With the PK set and $t = 0.5$, RFS gives the best trade-off of all schedules: it sheds load on one test day in 2,000, like IB with $t = 1.0$ and PK with $t = 0.5$ under min–max, at a mean cost of 3,371 $/day, 2.0% and 16.5% below these and only 0.6% above SAA, which sheds load on 59 days; with errors scaled by 1.25 it sheds on 0.70% of the days. It hedges with capacity rather than energy: 37.1 MWh day-ahead, 5.5 MWh of it between 09:00 and 16:00 (SAA: 36.4 and 4.6 MWh), 3.6 MWh of MT upward reserve, 4.1 MWh of DR capacity and 0.01 MWh/day of curtailment. The set decides which PV shortfalls the schedule must survive; the sample average decides how. C&CG needed five

robust scenarios and three mixed-integer solves (167 s), and the certified worst-case penalty of 0.20 $ keeps shedding and limit violations below that cost (0.1 kWh of shed load) for every error in the set.

With $t = 0.2$, a single robust scenario suffices and the schedule stays close to SAA (2.9% of days with shedding): the smaller set does not contain the PV shortfalls that cause shedding.

With the IB set, RFS converges slowly, as each iteration removes only one of the many sparse worst cases of the 198-dimensional set. With $t = 0.5$, the worst penalty found fell to 0.05 $ after five mixed-integer solves, but its bound stayed at 1,740 $; the schedule costs as much as SAA (3,353 $/day) and sheds load on 1.1% of the days. With $t = 1.0$, the worst penalty was still 1,290 $ after six solves, and the schedule costs 3,366 $/day and sheds on 0.70% of the days. Neither matches RFS with the PK set, whose five robust scenarios, built from the correlated error modes, cover the PV shortfalls that cause shedding.

6. Discussion

This section discusses when the ranking of the sets would change and how the main modelling assumptions affect the results.

6.1 When would the ranking of the sets change?

The disadvantage of PK under min–max comes from the cost of unwinding day-ahead energy: surpluses are sold at half the day-ahead price within a 0.4-MW deviation limit, and the rest is curtailed. Where surplus energy can be resold near the day-ahead price, the energy hedge of PK would cost less and the gap should narrow; where reserve capacity is expensive or scarce, the reserve hedge of IB would lose part of its advantage. RFS removes this dependence, as the sample average, not the worst case, chooses the hedge.

6.2 Two-stage recourse versus real-time operation

The second stage re-dispatches knowing the whole day's errors, as in most two-stage robust scheduling models [10, 11], and the out-of-sample evaluation uses the same recourse. Real-time operation is non-anticipative: a battery discharged at 17:00 cannot be recharged in anticipation of a shortfall revealed at 20:00. Multistage robust models with non-anticipative policies [18, 34] would give more conservative but more realistic first-stage decisions. Using the same recourse for all sets keeps the comparison consistent, but the absolute shedding frequencies are optimistic.

6.3 Linearized network and forecast-bounded recourse

The linearized branch-flow model ignores losses and so overestimates voltages on long, heavily loaded laterals, and the robust model bounds curtailment and shedding by the forecasts. Both keep the recourse linear in the errors, as the vertex argument of Section 4 requires; the out-of-sample evaluation uses the realized bounds.

6.4 Synthetic errors

The errors come from a factor model with heavy-tailed innovations, not from measured forecasts. Real forecast errors have seasonal and weather-dependent structure, calling for conditional (weather-regime) sets, and the grouping of sites into zones could be learned by partitioning the graph of error correlations, for which global semidefinite-relaxation algorithms exist [35]. The sizing rule and RFS

apply unchanged to measured data, and the experiment with weaker correlation indicates that the advantage of IB under min–max shrinks as the correlation weakens.

6.5 Exactness of the sub-problem

The mixed-integer sub-problem is the bottleneck. With PK it has K rather than 198 pairs of binaries and closes within seconds for $t \leq 0.5$, a practical advantage of the principal-component set. With IB, its bound remained loose after 300 s, although no worst case more than 0.1% above the lower bound was found, so the IB schedules are validated out of sample and RFS with IB could not be certified. Branch-and-bound algorithms for very large separable problems with one or a few linear constraints exist for continuous variables with indefinite quadratic costs [36] and for mixed-binary variables [37, 38]; the budget constraint of (11) has this knapsack structure, which the local search exploits for a fixed dual vector, but the recourse duals couple the coordinates, so such algorithms would need adaptation to close the gap of IB.

6.6 Larger networks

The master problem grows with the number of units. With many identical units, e.g., fleets of MTs or batteries, its commitment variables become symmetric; aggregated formulations remove this symmetry [39] and, with perspective reformulation, tighten the continuous relaxation [40], as aggregated models do for pumped-storage plants whose units share one reservoir [41].

7. Conclusion

This paper compared uncertainty sets for two-stage robust SGLS scheduling of an ADN with correlated forecast errors, all sized by the same reliability rule, solved by the same algorithm and evaluated on the same 2,000 out-of-sample days. Four conclusions follow.

(1) Joint coverage does not make a schedule reliable. The independent budget sets contain none of the test days, yet the schedule of the largest sheds load on one test day in 2,000; sets sized for 95% joint coverage avoid shedding, but at 3–11% higher cost.

(2) Under min–max, the independent budget set trades off cost and shedding better than the principal-component sets, although these describe the errors better and make the worst-case sub-problem easy to certify: at shedding targets of 1% and 0.2%, it is 9.5% and 12.2% cheaper.

(3) The reason is the hedge that the worst case selects: the common-mode PV deficit of a principal-component set is covered with day-ahead energy that is paid every day, needed on few days, and displaces 13–19% of the renewable energy on the others, whereas the sparse worst cases of the independent set lead to reserve hedges.

(4) Used only for robust feasibility, with a sample-average cost, the principal-component set gives the best schedule: shedding on 0.05% of the test days at a mean cost 0.6% above the stochastic program, which sheds on 3.0% of the days, with certified robust feasibility.

For practice, this suggests using correlation-aware, data-driven sets to define what a schedule must withstand, calibrating their size on validation data for the reliability target, and pricing the schedule with scenarios rather than its worst case. Measured forecast errors, non-anticipative multistage recourse and AC power flow are left for future work.

References

[1] Liu J, Ai X, Chen C, Li S, Cui S, Fang J, et al. Reachable-Set Decomposition for Real-Time Aggregation of Multi-Zone HVAC Fleets. arXiv preprint arXiv:2605.13836. 2026. doi:10.48550/arXiv.2605.13836.

[2] Liu J, Ai X, Fang J, Cui S, Wang S, Yao W, et al. Continuous-Time Aggregation of Massive Flexible HVAC Loads Considering Uncertainty for Reserve Provision in Power System Dispatch. IEEE Transactions on Smart Grid. 2024; 15(5):4835-4849.
[3] He J, Wu Y, Wu J, Li S, Liu F, Zhou J, et al. Towards cleaner heating production in rural areas: Identifying optimal regional renewable systems with a case in Ningxia, China. Sustainable Cities and Society. 2021; 75:103288.
[4] Wu S, Zhang L, Li S, Pan D, Li S. From green electricity to green fertilizer: Economic and environmental benefits of ammonia diversified utilization in agricultural parks hybrid renewable energy systems. Energy Economics. 2026; 161:109537.
[5] Wang J, Zeng P, Yin L, Dong Y, Liu X, Yang R. A two-stage distributed stochastic planning method for source-grid-load-storage flexibility resources considering flexible ramp capacity. International Journal of Electrical Power & Energy Systems. 2024; 160:110134.
[6] Jia C, Wang H, Cao Y, Jia H, Mu Y, Yu X. A risk-averse multi-timescale stochastic integration framework for a community multi-energy system under correlated uncertainties. Applied Energy. 2026; 408:127364.
[7] Zhang N, Chen Y, Deng H, Ma T, Zheng S, Zhang Z. Flexibility-driven distributed coordination for low-carbon operation of active distribution networks with multi-park integrated energy systems. Applied Energy. 2026; 415:127861.
[8] Zhang Y, Zhao H, Qi Z, Li B. A two-stage low-carbon economic coordinated dispatching model for generation-load-storage resources considering flexible supply-demand balance. Applied Energy. 2024; 373:123981.
[9] Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A. Adjustable robust solutions of uncertain linear programs. Mathematical Programming. 2004; 99(2):351-376.
[10] Bertsimas D, Litvinov E, Sun XA, Zhao J, Zheng T. Adaptive Robust Optimization for the Security Constrained Unit Commitment Problem. IEEE Transactions on Power Systems. 2013; 28(1):52-63.
[11] Jiang R, Wang J, Guan Y. Robust Unit Commitment With Wind Power and Pumped Storage Hydro. IEEE Transactions on Power Systems. 2012; 27(2):800-810.
[12] Zeng B, Zhao L. Solving two-stage robust optimization problems using a column-and-constraint generation method. Operations Research Letters. 2013; 41(5):457-461.
[13] Ding T, Liu S, Yuan W, Bie Z, Zeng B. A Two-Stage Robust Reactive Power Optimization Considering Uncertain Wind Power Integration in Active Distribution Networks. IEEE Transactions on Sustainable Energy. 2016; 7(1):301-311.
[14] Gao H, Liu J, Wang L. Robust Coordinated Optimization of Active and Reactive Power in Active Distribution Systems. IEEE Transactions on Smart Grid. 2018; 9(5):4436-4447.
[15] Bertsimas D, Sim M. The Price of Robustness. Operations Research. 2004; 52(1):35-53.
[16] Guan Y, Wang J. Uncertainty Sets for Robust Unit Commitment. IEEE Transactions on Power Systems. 2014; 29(3):1439-1440.
[17] Lorca A, Sun XA. Adaptive Robust Optimization With Dynamic Uncertainty Sets for Multi-Period Economic Dispatch Under Significant Wind. IEEE Transactions on Power Systems. 2015; 30(4):1702-1713.
[18] Lorca A, Sun XA. Multistage Robust Unit Commitment With Dynamic Uncertainty Sets and Energy Storage. IEEE Transactions on Power Systems. 2017; 32(3):1678-1688.
[19] Ning C, You F. Data-driven decision making under uncertainty integrating robust optimization with principal component analysis and kernel smoothing methods. Computers & Chemical Engineering. 2018; 112:190-210.
[20] Ning C, You F. Data-Driven Adaptive Robust Unit Commitment Under Wind Power Uncertainty: A Bayesian Nonparametric Approach. IEEE Transactions on Power Systems. 2019; 34(3):2409-2418.
[21] Bertsimas D, Gupta V, Kallus N. Data-driven robust optimization. Mathematical Programming. 2018; 167(2):235-292.
[22] Golestaneh F, Pinson P, Azizipanah-Abarghooee R, Gooi HB. Ellipsoidal Prediction Regions for Multivariate Uncertainty Characterization. IEEE Transactions on Power Systems. 2018; 33(4):4519-4530.
[23] Velloso A, Street A, Pozo D, Arroyo JM, Cobos NG. Two-Stage Robust Unit Commitment for Co-Optimized Electricity Markets: An Adaptive Data-Driven Approach for Scenario-Based Uncertainty Sets. IEEE Transactions on Sustainable Energy. 2020; 11(2):958-969.
[24] Mohajerin Esfahani P, Kuhn D. Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. Mathematical Programming. 2018; 171(1-2):115-166.
[25] Lai K, Nie Y, Zhou H, Xu D, Li S. Generation Construction Planning Considering Renewable Energy Sources. In: 2025 6th International Conference on Clean Energy and Electric Power Engineering (ICCEPE). 2025; 565-572.
[26] Yan R, Leng S, Fan Z, Li S. Distributionally Robust Chance-Constrained Hydropower-Based Virtual Power Plants. In: 2025 5th IEEE International Conference on Energy Engineering and Power Systems (EEPS). 2025; 223-228.
[27] Zhao C, Guan Y. Unified Stochastic and Robust Unit Commitment. IEEE Transactions on Power Systems. 2013; 28(3):3353-3361.
[28] An Y, Zeng B. Exploring the Modeling Capacity of Two-Stage Robust Optimization: Variants of Robust Unit Commitment Model. IEEE Transactions on Power Systems. 2015; 30(1):109-122.
[29] Wang R, Li L, Lin W, Wang E, Yuan J. Mastering the complexity: An enhanced cellular automata-based framework

for simulating resilience of hospital Power-Water-Firefighting-Space nexus system. Simulation Modelling Practice and Theory. 2025; 144:103177.
[30] Angelopoulos AN, Bates S. Conformal Prediction: A Gentle Introduction. Foundations and Trends in Machine Learning. 2023; 16(4):494-591.
[31] Baran ME, Wu FF. Network reconfiguration in distribution systems for loss reduction and load balancing. IEEE Transactions on Power Delivery. 1989; 4(2):1401-1407.
[32] Li S, Wu J, Lu C, Deng Z, Fang SC. An Efficient Solution Method for Solving Convex Separable Quadratic Optimization Problems. Pacific Journal of Optimization. 2025; 21(4):677-689.
[33] Baran ME, Wu FF. Optimal capacitor placement on radial distribution systems. IEEE Transactions on Power Delivery. 1989; 4(1):725-734.
[34] Lorca Á, Sun XA, Litvinov E, Zheng T. Multistage Adaptive Robust Optimization for the Unit Commitment Problem. Operations Research. 2016; 64(1):32-51.
[35] Wu J, Lu C, Li S, Deng Z. A semidefinite relaxation based global algorithm for two-level graph partition problem. Journal of Industrial and Management Optimization. 2023; 19(9):7036-7053.
[36] Li S, Deng Z, Lu C, Wu J, Dai J, Wang Q. An efficient global algorithm for indefinite separable quadratic knapsack problems with box constraints. Computational Optimization and Applications. 2023; 86(1):241-273.
[37] Lu C, Wu J, Deng Z, Li S. A fast global algorithm for singly linearly constrained separable binary quadratic program with partially identical parameters. Optimization Letters. 2023; 17(3):613-628.
[38] Qu G, Li S, Deng Z, Lu C. A fast global algorithm for multi-linearly constrained separable binary quadratic program. Journal of Industrial and Management Optimization. 2025; 21(2):1456-1473.
[39] Qu G, Lu C, Wu J, Li S, Deng Z, Fang SC. Reformulations of Separable Quadratic Optimization Problems with Symmetric Structures via Variable Aggregation. Journal of Optimization Theory and Applications. 2026; 209(3):71.
[40] Wu J, Li S, Lu C, Deng Z, Fang SC. Variable Aggregation-based Perspective Reformulation for Mixed-Integer Convex Optimization with Symmetry. arXiv preprint arXiv:2602.04123. 2026. doi:10.48550/arXiv.2602.04123.
[41] Li S, Wu J, Deng Z. Variable aggregation-based formulations for pumped storage hydro model in the day-ahead unit commitment problem. arXiv preprint arXiv:2510.07751. 2025. doi:10.48550/arXiv.2510.07751.