

Stochastic Optimization for Resource Adequacy in Capacity Markets with Storage and Renewables

Baptiste Rabecq, Andy Sun
Operations Research Center, MIT
Cambridge, MA, USA
(brabecq, sunx)@mit.edu

Feng Zhao, Tongxin Zheng, Xiaochu Wang, Yufan Zhang
ISO-New England
Holyoke, MA, USA
(fzhao, tzheng, xiwang, yuzhang)@iso-ne.com

Abstract—The integration of storage and renewable resources fundamentally alters resource-adequacy analysis. Because storage couples decisions across time, it invalidates the traditional reliability models that are based on time-independent capacity demand curves. Moreover, renewables introduce temporally correlated intermittency. To address this, we formulate the capacity procurement problem as a two-stage stochastic program, where the capacity decision is made in the first stage, while the expected unserved energy is evaluated by a second-stage dispatch problem that considers uncertainties such as generator failures via Markov chains, temporally correlated renewable output, and stochastic load. We implement the resulting stochastic capacity procurement (SCP) model on a New England system with 305 generators, including conventional, renewable, and storage units. Using the stochastic decomposition (SD) algorithm, we solve the SCP with up to 20,000 Monte Carlo samples, each representing a six-month trajectory of more than 4,300 hours of uncertainty across all units. We analyze the convergence behavior of SD and show that convergence for the stochastic program happens faster than reliable estimation of the reliability metrics, which require more samples than are used in typical stochastic programs. These results show that chronologically detailed Monte Carlo sampling can be integrated into capacity procurement optimization in a computationally tractable manner, enabling reliability evaluation with controlled statistical accuracy at realistic system scales.

Index Terms—Resource Adequacy, Capacity Market, Stochastic Optimization, Energy Storage, Stochastic Decomposition

I. INTRODUCTION

In deregulated electricity markets, capacity markets were introduced to achieve long-term resource adequacy [1]. They treat resource adequacy as a trade-off between the cost of supplying capacity and the value of avoided unserved energy. Generators incur investment and operating costs to access the grid, while consumers suffer an unobserved disutility when load is shed, typically represented by administratively defined downward-sloping demand curves [2]. These demand curves are often calibrated to probabilistic reliability criteria such as Loss of Load Expectation (LOLE), or Expected Unserved Energy (EUE) [3], using the “1-in-10 years” standard.

With increasing renewable penetration and more rapid variations in demand, assessing the contribution of individual generation units to system reliability becomes increasingly difficult. Recent work from ESIG and DOE argues that, in high-renewable systems, loss-of-load risk is no longer confined to a few peak hours, and that adequacy assessments must evaluate risk across many hours and seasons [4], [5]. In response,

TABLE I
NOTATION TABLE FOR SETS, VARIABLES, AND PARAMETERS

Notation	Description
Sets	
$\mathcal{T}$ G, R, S	Set of hourly time steps, of length T . Set of conventional generators, renewable generators and energy storage systems
$\mathcal{D}, \mathcal{W}, \mathcal{M}$	Set of days, weeks, months of length D, W , respectively
$\mathcal{T}_d, \mathcal{T}_w, \mathcal{T}_m \subset \mathcal{T}$	Set of hours in month m , week w and day d , respectively.
$\mathcal{X}$	Set of Capacity Decisions, polyhedral
Optimization Variables	
x	Capacity decision
LS_t	Load shed in time period t
z_t	Total power output in time period t
$p_{i,t}$	Power output of generator $i \in R \cup G$ in time period t
$e_{s,t}$	Battery state of charge of storage system s in time period t
$p_{s,t}^+, p_{s,t}^-, \eta_s^+, \eta_s^-$	Charging and discharging power, efficiency of storage system s in time period t
Economic parameters	
c	Capacity bids of the generators (\$/kW-month)
$VOLL$	Value of Lost Load (\$/kWh)
Physical Parameters	
$K_g^{\text{day}}, K_g^{\text{week}}, K_g^{\text{month}} \in [0, 1]$	Proportion of energy that can be generated by generator g in a day, week or month, respectively.
H_s	Duration of storage s , in hours.

demand curves are now typically built using Monte Carlo tools that simulate outage, load, and renewable uncertainty [6]. However, in systems with significant storage, reliability is inherently intertemporal: a resource’s ability to serve load at a given hour depends on the feasibility of its entire prior state-of-charge trajectory. Scalar capacity metrics and hour-by-hour availability models fail to enforce these long-horizon feasibility constraints and can therefore misrepresent the contribution of storage to system reliability.

This paper formulates resource adequacy (RA) as a stochastic capacity procurement problem (SCP): a two-stage stochastic optimization problem that explicitly accounts for storage dynamics and chronologically detailed uncertainty. First-stage decisions determine generation and storage capacity. In the

second stage, adequacy is evaluated through Monte Carlo sampling of long-horizon uncertainty trajectories, where each sampled realization is assessed by solving a load-shedding minimization problem subject to dispatch and state-of-charge constraints. This formulation guarantees that reliability is measured using feasible operating trajectories rather than approximations.

While closely related to the stochastic capacity expansion literature [7]–[9], our setting differs in two critical respects. First, resource adequacy outcomes are driven by temporally correlated rare events that are poorly captured by pre-sampled representative days or small scenario sets, leading to underrepresentation of loss-of-load conditions. Second, convergence of a finite-scenario stochastic program does not guarantee accurate estimation of reliability metrics such as LOLE or EUE. Indeed, any finite-sample approximation produces a downward-biased estimate of expected unserved energy, and statistical error in reliability criteria decreases only with increasing sample size [10]. The resulting capacity mix may still produce inaccurate reliability estimates because of insufficient sampling of rare events.

We address this challenge by explicitly separating optimization error from statistical error in our SCP model. We adopt a stochastic decomposition (SD) algorithm [11] that dynamically samples uncertainty trajectories during the optimization process, allowing capacity decisions to be optimized while assessing the statistical accuracy of reliability metrics.

The paper makes three contributions. First, we propose a storage-aware formulation of SCP that enforces long-horizon intertemporal feasibility. Second, we demonstrate that this formulation can be solved at realistic scales using SD, even with tens of thousands of Monte Carlo scenarios. Third, we provide a statistical assessment of reliability metrics at the optimal solution, showing the distinction between optimization convergence and reliability estimation accuracy. We present numerical results on a stylized ISO New England system with hourly resolution on a six-month horizon and tens of thousands of Monte Carlo scenarios.

Demand curves are an adequate tool to generate market price signals. By contrast, our centralized optimization approach does not directly address this, which is a potential limitation. A thorough treatment of market design and pricing implications lies beyond the scope of this paper.

II. MODEL

In this work, we take the central planner's point of view, solving a risk-neutral surplus maximization problem [2], [12]. We introduce a compact form for the SCP model:

$$\min_{x \in \mathcal{X}} c^\top x + V(x) \quad (1)$$

where x is the capacity decision, $c^\top x$ is the total capacity payment, and $V(x)$ is the expected cost of unserved energy. Once the capacity decision x is fixed, we evaluate the reliability it delivers to the grid, dependent on uncertainties in system operation. We model three sources of uncertainty:

- 1) **Generator failures.** Following [3], each unit i 's availability is a two-state, time-homogeneous, discrete-time Markov chain on $\{0, 1\}$. When $\xi_{i,t}^A = 0$, unit i is unavailable at hour t , and when $\xi_{i,t}^A = 1$, it can generate electricity to its full capacity. Reliability is parameterized by the Forced Outage Rate (FOR_i) and Mean Time To Repair (MTTR_i). Define the per-period repair probability $\lambda_i := 1/\text{MTTR}_i$ and per-period failure probability $\mu_i := \frac{\text{FOR}_i}{(1-\text{FOR}_i)\text{MTTR}_i}$. The availability process $(\xi_{i,t}^A)$ has transition matrix

$$P_i = \begin{pmatrix} 1 - \lambda_i & \lambda_i \\ \mu_i & 1 - \mu_i \end{pmatrix}. \quad (2)$$

For each unit i , a scenario is a sample path $(\xi_{i,1}^A, \dots, \xi_{i,T}^A)$ of this Markov chain over the horizon $1, \dots, T$.

- 2) **Renewable generation.** For a renewable unit r and time period t , we model generation uncertainty via a capacity factor $\xi_{r,t}^{\text{CF}} \in [0, 1]$, which scales the plant's generation capacity proportionally. To capture temporal correlation within a day, we work with daily capacity-factor profiles. We construct from historical data K representative daily profiles $\{(\xi_{r,t,k}^{\text{CF}})_{t \in 1, \dots, |\mathcal{T}_d|}, k \in \{1, \dots, K\}\}$, and define a discrete distribution $(\pi_{r,1}, \dots, \pi_{r,K})$ over these profiles. For each day d and unit r , we draw a profile index $\kappa_{r,d} \in \{1, \dots, K\}$ independently with $\mathbb{P}(\kappa_{r,d} = k) = \pi_{r,k}$, and set $\xi_{r,t}^{\text{CF}} := \xi_{r,\kappa_{r,d},t}^{\text{CF}}$ for all t in day d .
- 3) **Load uncertainty.** We model load uncertainty via a multiplicative factor. Let $(L_t)_{t=1}^T$ be a fixed historical load profile and let $\xi_L \sim \text{Unif}[0.8, 1.2]$. The stochastic load process is $\xi_t^{\text{Load}} := \xi_L L_t$, $t = 1, \dots, T$.

We combine all randomness into a vector $\xi := (\xi^A, \xi^{\text{CF}}, \xi^{\text{Load}})$, and define $(\Xi, \mathcal{F}, \mathbb{P})$ a probability space supporting ξ . We assume that generator failures, renewable profiles, and load are mutually independent, so that the joint law $\mathbb{P}$ factorizes as a product measure of the marginals.

Following [2], we assume that the expected cost of unserved load can be written as:

$$V(x) = \text{VOLL} \cdot \mathbb{E}[U(x, \xi)], \quad (3)$$

where $U(x, \xi)$ is the unserved energy in scenario ξ with capacity decision x . For every x, ξ , $U(x, \xi)$ can be found as the value function of the following second-stage dispatch problem:

$$U(x, \xi) = \min \sum_{t \in \mathcal{T}} LS_t \quad (4a)$$

$$\text{s.t. } LS_t = \xi_t^{\text{LOAD}} - z_t, \quad \forall t \in \mathcal{T} \quad (4b)$$

$$z_t = \sum_{g \in G} p_{g,t} + \sum_{r \in R} p_{r,t} \quad \forall t \in \mathcal{T}, \quad (4c)$$

$$+ \sum_{s \in S} (p_{s,t}^- - p_{s,t}^+),$$

$$0 \leq p_{g,t} \leq \xi_{g,t}^A x_g, \quad \forall g \in G, t \in \mathcal{T} \quad (4d)$$

$$0 \leq p_{r,t} \leq \xi_{r,t}^A \xi_{r,t}^{\text{CF},k} x_r, \quad \forall r \in R, t \in \mathcal{T} \quad (4e)$$

$$0 \leq p_{s,t}^+ \leq \xi_{s,t}^A x_s, \quad \forall s \in S, t \in \mathcal{T} \quad (4f)$$

$$0 \leq p_{s,t}^- \leq \xi_{s,t}^A x_s, \quad \forall s \in S, t \in \mathcal{T} \quad (4g)$$

$$0 \leq e_{s,t} \leq H_s x_s, \quad \forall s \in S, t \in \mathcal{T} \quad (4h)$$

$$e_{s,1} = 0, \quad \forall s \in S, \quad (4i)$$

$$e_{s,t+1} = e_{s,t} + \eta_s^+ p_{s,t}^+ - \frac{p_{s,t}^-}{\eta_s^-}, \quad \forall s \in S, t \in \mathcal{T} \quad (4j)$$

$$\sum_{t \in \mathcal{T}_d} p_{g,t} \leq K_g^{\text{day}} x_g | \mathcal{T}|_d, \quad \forall g \in G, d \in \mathcal{D} \quad (4k)$$

$$\sum_{t \in \mathcal{T}_w} p_{g,t} \leq K_g^{\text{week}} x_g | \mathcal{T}|_w, \quad \forall g \in G, w \in \mathcal{W} \quad (4l)$$

$$\sum_{t \in \mathcal{T}_m} p_{g,t} \leq K_g^{\text{month}} x_g | \mathcal{T}|_m, \quad \forall g \in G, m \in \mathcal{M} \quad (4m)$$

$$L S_t \geq 0, \quad \forall t \in \mathcal{T}. \quad (4n)$$

The objective (4a) minimizes total load shed. Constraints (4b)–(4c) enforce power balance, (4d)–(4g) impose unit and storage power limits, (4h)–(4j) describe storage state-of-charge dynamics, and (4k)–(4m) impose daily, weekly, and monthly energy limits as a proxy for fuel replenishment.

III. METHODS

Problem (4) is a linear program with fixed recourse and the capacity decision x only appears in the right hand side. As a result $U(x, \xi)$ is a convex piecewise linear function of x , for almost every ξ . Moreover, for any $x \in \mathcal{X}$ and almost every ξ , setting $LS = \xi^{\text{LOAD}}$ and setting all other variables to 0 yields a feasible solution. Since all variables are bounded above and below, we conclude that for every x and almost every ξ , $U(x, \xi) < \infty$. Therefore, problem (4) has relatively complete recourse. That is, there exists a (single) dual feasible set Π , T and h measurable maps such that:

$$U(x, \xi) = \max_{\theta \in \Pi} \theta^\top (h(\xi) - T(\xi)x), \quad (5)$$

and there exists a feasible solution for every x and almost every ξ .

Problem (1) is therefore a two-stage stochastic convex program, where the second-stage function $V(x)$ is the expectation of a linear program with relatively complete recourse. While $V(x)$ is not available in closed form, we approximate it by Monte Carlo sampling: given N_k scenarios $(\xi^i)_{i=1, \dots, N_k}$, we form the sample-average estimator

$$\hat{V}_{N_k}(x) = \frac{1}{N_k} \sum_{i=1}^{N_k} U(x, \xi^i), \quad (6)$$

which is an unbiased estimator of $V(x)$. We use the stabilized SD algorithm [11], which can be considered a sampling version of a proximally stabilized Stochastic Benders algorithm. Contrary to applying Benders decomposition to a fixed scenario set, SD draws new samples over the course of the algorithm and is designed to approximate the true expectation $V(x)$ rather than a particular finite-sample problem.

At every iteration k , the algorithm constructs a piecewise linear lower approximation $\hat{f}_k$ of the sample average $\hat{V}_{N_k}$. Let $\mathcal{U}_k := \{\alpha_j^k, \beta_j^k\}_{j=1, \dots, k}$ such that $\hat{f}_k(x) = \max_{1 \leq j \leq k} \alpha_j^k + \beta_j^k^\top x$. Given a proximal penalty ρ , and a proximal center

$\bar{x}_k := x_{i_k}$ for some $i_k \leq k$, the iterate x_{k+1} is obtained by solving the following convex Quadratic Program:

$$x_{k+1} = \arg \min_{x \in \mathcal{X}} c^\top x + \hat{f}_k(x) + \frac{\rho}{2} \|x - \bar{x}_k\|_2^2. \quad (7)$$

The details of the cutting plane construction in the SD algorithm are presented in 1. We refer the interested reader to [11] for a more precise description.

Algorithm 1 Stabilized Stochastic Decomposition

```

1: Initialize:  $E_0 \leftarrow \emptyset$ ,  $x_0 \in \mathcal{X}$ ,  $r \in (0, 1)$ ,  $\rho \in \mathbb{R}^{++}$ ,  $i_0 = 0$ ,  $\mathcal{U}_0 = \emptyset$ 
2: for  $k \in \{1, \dots, n\}$  do,
3:   Sample  $\xi^{N_{k-1}+1}, \dots, \xi^{N_k}$  from  $\mathbb{P}$ .
4:   Solve dispatch problems: Solve (4) at  $x_k$  and  $x_{i_{k-1}}$  for every sample, record optimal duals  $\Theta_k := \{\theta_{N_k+1}^k, \dots, \theta_{N_{k+1}}^k, \bar{\theta}_{N_k+1}^k, \dots, \bar{\theta}_{N_{k+1}}^k\}$ 
5:    $\Pi_k \leftarrow \Pi_{k-1} \cup \Theta_k$ .
6:   Update Models:
7:   for  $j \in \{1, \dots, k\}$  do
8:     if  $j \in \{i_k, k\}$  then
9:        $\theta_t^j := \arg \max_{\Pi_k} \theta^\top (h(\xi^t) - T(\xi^t)x_j)$ ,  $\forall t \in [N_{k+1}]$ 
10:       $\alpha_j^k \leftarrow \frac{1}{N_k} \sum_{t=1}^{N_k} \theta_t^j^\top h(\xi^t)$ ,  $\beta_j^k \leftarrow \frac{1}{N_k} \sum_{t=1}^{N_k} \theta_t^j^\top T(\xi^t)$ 
11:     else
12:        $\alpha_j^k \leftarrow \frac{k-1}{k} \alpha_j^{k-1}$ ,  $\beta_j^k \leftarrow \frac{k-1}{k} \beta_j^{k-1}$ 
13:     end if
14:   end for
15:    $\mathcal{U}_k \leftarrow (\alpha_j^k, \beta_j^k)_{j \leq k}$ 
16:   if  $\hat{f}_k(x_k) - \hat{f}_k(\bar{x}_{k-1}) \leq r(\hat{f}_{k-1}(x_k) - \hat{f}_{k-1}(\bar{x}_{k-1}))$  then
17:      $i_k \leftarrow k$ 
18:   else  $i_k \leftarrow i_{k-1}$ 
19:   end if
19:   Get  $x_{k+1}$  from (7)
20: end for

```

We emphasize that the cuts are explicitly recomputed only at the current proximal center and the current iterate; at all other points, the corresponding cuts are simply decayed. Because problem (4) has a nonnegative optimal value, this decay preserves the property that the model is a statistical lower approximation of the true objective $V(x)$. The described dual caching procedure allows for increasing the accuracy of the model compared to batch first-order methods (all samples are used to generate cuts), while keeping the iteration cost low compared to Benders decomposition.

The quantity $\Delta_k := \hat{f}_k(\bar{x}_k) - \hat{f}_k(x_{k+1})$ is the model gap, representing the gap between the current minimizer of the model and the proximal center. A small gap is a good indicator of good in-sample convergence. Evaluating the expected cost of lost load $V(x)$ is a tail estimation problem. Under acceptable reliability standards, $U(x, \xi) = 0$, in almost all scenarios, so only a small fraction of realizations are informative about V . Coupled with the high-dimensional uncertainty, this makes finite-sample estimators prone to overfitting in low to medium sample regimes, and achieving a small generalization error typically requires large scenario sets.

IV. EXPERIMENTS

We solve the stochastic capacity procurement model using data from the ISO New England (ISO-NE) system. Feasible




Fig. 1. Convergence diagnostics of the SD algorithm: (a) model gap Δ_k , (b) incumbent objective value, and (c) expected cost of unserved energy, all as functions of the number of iterations for the summer and winter capacity auction, at aggregation level 10 (5 hours) MW and 100 MW (2 hours).

capacity decisions are based on cleared capacity from the ISO-NE 18th Forward Capacity Auction [13], with bidding prices derived from EIA cost data [14] and assumed capacity market revenue shares. Forced-outage uses ISO-NE’s Equivalent Forced Outage Rate demand statistics by technology [15], and mean time to repair values from NREL’s Regional Energy Deployment System 2.0 documentation [16].

Renewable generation profiles are obtained from one year of hourly solar and wind data from the NREL National Solar Radiation Data Base [17] and the NREL Bias-Corrected High-Resolution Rapid Refresh dataset [18], respectively. We convert these series into capacity factor profiles using Direct Normal Irradiance for solar and the `windpowerlib` Python package for wind [19]. We capture intra-day temporal correlation using k -medoids clustering to obtain representative daily profiles and probabilities, sampling five profiles per season.

We clear capacity in two distinct auctions: one summer auction, from May 1 to October 30, and one winter auction, from November 1 to April 30. We solve the problems with hourly dispatch decisions on the corresponding time horizons of approximately 4300 hours. To limit the size of the problem, we aggregate units of the same technology type when their capacity is below a prescribed threshold, taken as 100 MW and 10 MW. With a 100 MW threshold, the system comprises 107 units in total, including 10 energy storage units, 2 solar units, 4 wind units, and 91 conventional units. With a 10 MW threshold, the system comprises 305 units, including 17 energy storage units, 10 solar units, 9 wind units, and 269 conventional units. We use a Value of Lost Load of \$100,000/MWh. According to market practice, we set a market cap based on the net Cost of New Entry, set to \$14.5/kW-month.

All experiments are conducted on the Engaging MIT computing cluster, using AMD EPYC 9654 96-core processors. We set the SD batch size to $b = 32$ scenarios per iteration, solved in parallel at every iteration of the algorithm. The choice of b balances per-iteration computational speed with efficiency: larger batches slow each iteration, but smaller batches require more iterations to reach comparable solution quality.

V. RESULTS

Throughout the runs, we monitor the in-sample standard deviation of our *EUE* estimator and stop when it is on the order of 25% of the estimated *EUE*. On our system,

this number is reached when approximately 20,000 samples are used. We therefore use this total number of samples as a stopping criterion for our experiments. For the 100MW aggregated system, this happens after 2 hours, and 5 hours for the 10MW aggregated system.

A. Convergence of the algorithm

Figure 1 reports convergence diagnostics for the SD algorithm at aggregation level 10 and 100 MW over 5 and 2 hours of wall-clock time, respectively. Panel 1a shows the model gap Δ_k as a function of the iteration index. For both aggregation levels, the gap decreases rapidly and falls below 10^{-4} after a few hundred iterations, indicating that the master problem is well solved relative to the current sample information. Panel 1b displays the incumbent objective value $\hat{f}_k(\hat{x}_k)$; it decreases substantially during the first ~ 100 iterations, after which improvements are marginal and subsequent iterations produce small fluctuations around a stable level. The difference in objective value between aggregation levels is due to a different per-unit scaling. Panel 1c shows the corresponding estimated expected cost of unserved energy, which also stabilizes after the first 300 iterations, and is the same for both aggregation levels. We observe that the lower-dimensional runs (100 MW aggregation) converge faster than the higher-dimensional runs.

A small model gap and visually stable objective values indicate satisfactory convergence of the SD procedure, but they do not certify optimality for Problem (1). Because the estimates $\hat{f}_k(\hat{x}_k)$ are computed from finite samples, the incumbent decisions x_k continue to move slightly even when Δ_k is very small, reflecting residual sampling noise. In this work, we assess the final solution x^* via OOS statistical validation.

B. Statistical validation of the solution

We evaluate the final solution x^* returned by the algorithm both in-sample (IS) (i.e. during the algorithm execution) and out-of-sample (OOS), using the 10MW-thresholded system. Table II reports the estimated reliability metrics under the IS scenario set used during optimization, together with OOS estimates obtained from an independent Monte Carlo simulation of 20,000 samples. Across all metrics, the IS and OOS confidence intervals overlap, and the IS–OOS differences are of the same order as these intervals, with no consistent

TABLE II
IS VS OOS PERFORMANCE OF x^* . LOLE, EUE WITH 95% CI

Season/Metric	IS mean	OOS mean	Rel. OOS CI width
Summer EUE	91.70 $\pm$ 11.14	82.24 $\pm$ 9.11	0.22
Summer LOLE	0.045 $\pm$ 0.003	0.057 $\pm$ 0.004	0.15
Winter EUE	36.82 $\pm$ 5.48	44.34 $\pm$ 6.78	0.30
Winter LOLE	0.046 $\pm$ 0.004	0.049 $\pm$ 0.004	0.16




Fig. 2. Comparison between summer and winter capacity mix of x^* (MW). The blue and pink bars represent the cleared capacity and total available capacity, respectively.

sign. This provides no statistically significant evidence of systematic optimization-induced bias at the current Monte Carlo resolution. The relative CI widths in Table II range between 0.15 and 0.30, with larger values for EUE than for LOLE. Combined with the small EUE levels (below 100 MWh/year), this confirms that expected costs are driven by rare, high-impact events rather than by typical operating conditions.

Finally, Table II shows that at the optimal solution, LOLE levels are relatively conservative for this stylized system. Indeed, the LOLE is around 45% of the standard value of 0.1 days per year.

C. Solution and seasonal mix

Figure 2 displays the capacity mix x^* in Summer and Winter. The generators are divided into three categories: “Renewable”, denoting wind and solar generators, “Energy Storage”, and “conventional”, denoting the remaining ones. We observe that more capacity is contracted in the summer auction than in the winter auction. This is expected, as the capacity is primarily driven by the peak load, which is higher in the summer. In the winter auction, there is little renewable generation and no storage. In the summer, the renewable capacity allocation increases, and some storage capacity is allocated. In Figure 3, we observe that Loss of Load is concentrated on a relatively small set of peak hours, as expected. In some scenarios, the generator’s fuel constraints are binding, showing dependence of the EUE on long-term coupling constraints.

VI. CONCLUSION

In this work, we propose a two-stage stochastic program for achieving resource adequacy in capacity markets with significant storage and renewables. We solve the resulting large-scale stochastic program using stochastic decomposition with chronological Monte Carlo sampling. Numerical experiments on an ISO-NE system show that the proposed




Fig. 3. Load Shedding Distribution for the summer and winter auction (MWh)

method is computationally efficient with statistically reliable performance.

REFERENCES

- [1] P. L. Joskow, “Capacity payments in imperfect electricity markets: Need and design,” *Utilities Policy*, vol. 16, no. 3, pp. 159–170, 2008.
- [2] F. Zhao, T. Zheng, and E. Litvinov, “Constructing Demand Curves in Forward Capacity Market,” *IEEE Transactions on Power Systems*, vol. 33, no. 1, pp. 525–535, 2018.
- [3] R. Billinton and R. N. Allan, *Reliability Evaluation of Power Systems*. Springer US, 1996.
- [4] DOE, “The future of resource adequacy,” U.S. DOE, Tech. Rep., 2024. [Online]. Available: <https://www.energy.gov/sites/default/files/2024-04/2024%20The%20Future%20of%20Resource%20Adequacy%20Report.pdf>
- [5] D. Stencilik, A. Bloom, W. Cole, A. Figueroa Acevedo, G. Stephen, and A. Tuohy, “Redefining resource adequacy for modern power systems,” ESIG, Tech. Rep., 2021.
- [6] GE, “Resource Adequacy Software | GE Vernova.”
- [7] R. S. Go, F. D. Munoz, and J.-P. Watson, “Assessing the economic value of co-optimized grid-scale energy storage investments in supporting high renewable portfolio standards,” *Applied Energy*, vol. 183, pp. 902–913, 2016.
- [8] C. J. Newlun, J. D. McCalley, R. Amitava, A. J. Ardakani, A. Venktraman, and A. L. Figueroa Acevedo, “Adaptive Expansion Planning Framework for MISO Transmission Planning Process,” in *2021 IEEE Kansas Power and Energy Conference (KPEC)*, 2021, pp. 1–6.
- [9] T. Valencia Zuluaga, A. Musselman, J.-P. Watson, and S. S. Oren, “Parallel computing for power system climate resiliency: Solving a large-scale stochastic capacity expansion problem with mpi-sppy,” *Electric Power Systems Research*, vol. 235, p. 110720, 2024.
- [10] A. Shapiro, “Chapter 5: Statistical inference,” in *Lectures on Stochastic Programming: Modeling and Theory*, Third Edition, pp. 151–221.
- [11] J. L. Hight and S. Sen, “Finite master programs in regularized stochastic decomposition,” *Mathematical Programming*, vol. 67, no. 1, pp. 143–168, 1994.
- [12] C. Guo, C. Kroer, Y. Dvorkin, and D. Bienstock, “Incentivizing Investment and Reliability: A Study on Electricity Capacity Markets,” 2023.
- [13] ISO New England, “Forward capacity market reports — iso express,” <https://www.iso-ne.com/isoexpress/web/reports/auctions/-/tree/forward-capacity-mkt>.
- [14] U.S. EIA, “Cost and performance characteristics of new generating technologies, Annual Energy Outlook 2023,” U.S. EIA, Tech. Rep., 2023.
- [15] ISO New England, “Nerc GADS EFORD class averages as used by ISO new england,” ISONE, Technical report, 2024.
- [16] N. R. E. Laboratory, “Model documentation — ReEDS 2.0,” https://nrel.github.io/ReEDS-2.0/model_documentation.html, 2025.
- [17] M. Sengupta, Y. Xie, A. Lopez, A. Habte, G. Maclaurin, and J. Shelby, “The National Solar Radiation Data Base (NSRDB),” *Renewable and Sustainable Energy Reviews*, vol. 89, pp. 51–60, 2018.
- [18] G. Buster, P. Pinchuk, L. Lavin, B. Benton, and N. Bodini, “Bias Correcting NOAA’s High-Resolution Rapid Refresh (HRRR) Wind Resource Data for Grid Integration Applications,” 2024.
- [19] S. Haas, U. Krien, B. Schachler, S. Bot, V. Zeli, F. Maurer, K. Shivam, F. Witte, S. J. Rasti, Seth, and S. Bosch, “wind-python/windpowerlib: Update release (v0.2.2),” 2024.