

Cost Attribution And Risk-Averse Unit Commitment In Power Grids Using Integrated Gradient

René Carmona^{a,*}, Ronnie Sircar^a, Xinshuo Yang^l^a

^aPrinceton University, Princeton, NJ 08540, USA

Abstract

This paper introduces a novel approach to addressing uncertainty and associated risks in power system management, focusing on the discrepancies between forecasted and actual values of load demand and renewable power generation. By employing Economic Dispatch (ED) with both day-ahead forecasts and actual values, we derive two distinct system costs, revealing the financial risks stemming from uncertainty. We present a numerical algorithm inspired by the Integrated Gradients (IG) method to attribute the contribution of stochastic components to the difference in system costs. This method, originally developed for machine learning, facilitates the understanding of individual input features' impact on the model's output prediction. By assigning numeric values to represent the influence of variability on operational costs, our method provides actionable insights for grid management. As an application, we propose a risk-averse unit commitment framework, leveraging our cost attribution algorithm to adjust the capacity of renewable generators, thus mitigating system risk. Simulation results on the RTS-GMLC grid demonstrate the efficacy of our approach in improving grid reliability and reducing operational costs.

Keywords: risk allocation, unit commitment, economic dispatch, integrated gradient, risk aversion.

1. Introduction

In the pursuit of a more sustainable and environmentally responsible energy landscape, the world is experiencing a remarkable shift towards high levels of renewable energy penetration. The relentless expansion of wind, solar, and other clean energy technologies has brought us closer to reducing our carbon footprint and combating climate change. However, the integration of renewable sources introduces an unprecedented level of complexity into many fundamental aspects of power system management. Unlike their conventional counterparts, renewable resources are inherently variable, contingent upon factors like weather patterns, sunlight availability, and wind speed. Consequently,

new sources of randomness as well as inherent risk factors should be introduced into Unit Commitment (UC) and Economic Dispatch (ED), two crucial components of modern power system management. In the short term operational scheduling of electric power systems, UC involves the strategic selection of power generation units, determining which sources, whether they are conventional or renewable, should be dispatched to meet electricity demand over a specific time horizon. ED, on the other hand, optimizes the allocation of generation output among these chosen units in order to minimize costs while ensuring the reliable provision of electricity.

Traditionally, UC models take a deterministic approach in which the parameters and variables are assumed to be known exactly [11]. Due to their deterministic nature, these models are not capable of capturing variability and uncertainty. In recent years, there has been significant research in modeling and addressing the reliability of UC under uncertainty. Among those works, stochastic programming and robust optimization are two common methodologies frequently used.

*Corresponding author

Email addresses: rcarmona@princeton.edu (René Carmona), sircar@princeton.edu (Ronnie Sircar), xinshuoyang@gmail.com (Xinshuo Yang)

In stochastic programming models, the expected operating cost is minimized under a set of representative scenarios to meet the demand while satisfying the operational constraints [15, 17, 20, 13]. These representative scenarios may sometimes encompass highly rare events, leading to solutions that come with substantial costs. To tackle this issue, risk-averse UC models were developed with additional constraints to manage the risk exposures through the use of risk measures, conditional value at risk (CVaR) being the most widely used [7, 6, 19, 1]. Despite this, stochastic UC models can present computational difficulties, primarily arising from their inherent complexity, particularly when a substantial number of scenarios are integrated into the model.

Unlike models based on stochastic programming, robust UC models aim to address uncertainty by considering only the range of uncertainty, rather than relying on information about underlying probability distributions. Instead of minimizing the total expected cost, as is done in stochastic UC, robust UC focuses on minimizing the worst-case cost across all potential outcomes of uncertain parameters [2, 8, 16]. While robust UC solutions tend to be quite conservative, they have the advantage of avoiding the computational complexity associated with incorporating a large number of Monte Carlo scenarios.

In the present paper, we address the aforementioned uncertainty and its associated risks from a different perspective. The primary source of uncertainty stems from the deviations between forecasted and actual values, whether we are examining load demand or renewable power generation. By running ED with two sets of parameters—the day-ahead forecasts and the actual values—we obtain two distinct system costs. The difference between these costs represents the financial risks associated to these uncertainties. We present a numerical algorithm inspired by a widely recognized method known as the Integrated Gradients (IG) [14] to assess the individual contributions to each stochastic component to the difference of system costs. Originally introduced as an attribution mechanism in machine learning, the IG is typically used to attribute the significance or contribution of individual input features (e.g., pixels in an image or words in text) on a model’s output prediction, especially in deep neural network learning. The IG chooses a baseline input to ideally represent the absence or neutrality of input features.

It then calculates the contribution of each feature to the model’s output by integrating the gradients of the model’s prediction with respect to the input along a straight path from a baseline to the target input. In the proposed algorithm, the power production cost (ED) model is treated as a function that take input consisting of the random quantities such as load demand and renewable generations, and yields the system cost. We designate the day-ahead forecast as the baseline input, while the actual values serve as the target input. The algorithm will assign a numeric monetary value to represent the influence of this variability on the grid operational costs.

Certainly, these values can be used as mere informational tools in various contexts. As an example of application of the cost attribution algorithm, we propose a risk averse unit commitment framework to adjust the capacity of renewable generators using a large number of Monte Carlo scenarios as the target inputs in the attribution method. This framework considers the distribution of the cost attribution of individual generator through the Monte Carlo scenarios. It results in a reduction of their generation capacity (maximum generation level) if a certain portion of that capacity is deemed unlikely to be delivered and could potentially lead to elevated operational costs. By integrating this adjustment mechanism into the Unit Commitment (UC) optimization process, we can avoid relying on highly uncertain renewable productions without sacrificing the benefits of inexpensive and eco-friendly renewable energy, thereby reducing the overall risk in grid operations.

This work is part of research from the Princeton team ORFEUS¹, funded by ARPA-E under its program PERFORM². A related paper that uses Shapley values for renewables reliability risk allocation is [12].

The rest of the paper is organized as follows: in Section 2 we provide a comprehensive description of our problem setup and model configuration and introduce a numerical algorithm designed for attributing operational costs to individual components within this setup. We outline certain desirable properties of the attribution algorithm along with implementation techniques. Section 3 is devoted to a case study for the illustrative RTS-GMLC grid. We report simulation results

¹ORFEUS: Operational Risk Financialization of Electricity Under Stochasticity, orfeus.princeton.edu

²PERFORM—Performance-based Energy Resource Feedback, Optimization, and Risk Management

and performance of our cost attribution algorithm. In Section 4, we describe a framework on how to adjust the generation capacity of renewable generators using Monte Carlo scenarios and present yearly simulation results on the RTS-GMLC grid, highlighting its efficacy in mitigating system risk concerning loss of load. We conclude with a short recall of the main contributions of the paper in Section 5.

2. Cost Attribution

2.1. Integrated Gradient

Integrated Gradients (IG) is a technique widely used to attribute the predictions of a machine learning model to its input features. It is often used to interpret and explain the predictions of complex models, especially deep neural networks. The main goal of IG is to understand the contribution of each feature in the input to the final prediction made by the model. The idea is to compute the integral of the partial derivatives of the model’s output with respect to the input features while gradually transitioning from a baseline input (usually a neutral or zero input) to the actual input for which one wants to explain the prediction. This process helps in attributing the prediction incrementally to each feature while considering their contributions in the context of the overall prediction. This attribution method is reminiscent of the way sensitivities of derivative prices and portfolio values are understood and computed in financial engineering.

Mathematically speaking, the premise is the first order Taylor expansion, also known as the fundamental theorem of calculus. Suppose we have a function $F : \mathbb{R}^n \rightarrow [0, 1]$ that represents a deep neural network, an arbitrary input vector $x \in \mathbb{R}^n$ and a baseline input vector $x' \in \mathbb{R}^n$. The IG provide a quantitative measure of the contribution of each input feature x_i to the model’s prediction $F(x)$. Higher attributions indicate greater influence on the prediction, while lower attributions suggest lesser impact. The attribution for the i -th feature is obtained by integrating the gradients along the straight path from the baseline input x' to the actual input x :

$$I_i(x) := (x_i - x'_i) \int_{\lambda=0}^1 \frac{\partial F(x' + \lambda(x - x'))}{\partial x_i} d\lambda. \quad (1)$$

This integration is typically performed using numerical methods like the trapezoidal rule or Simpson’s

rule. One important property of Integrated Gradients is that the sum of the attribution scores across all features equals the difference in the model’s output between the actual input and the baseline, i.e.

$$\sum_{i=1}^n I_i(x) = F(x) - F(x'). \quad (2)$$

2.2. Production Cost Model

We adapt the technique of IG to the problem of cost attribution in power production cost models (PCM) using a single scenario of electricity demand and renewable production. Production cost modeling is a technique to simulate and analyze the economic and operational characteristics of electricity generation and distribution. It involves creating mathematical models that represent the various components of an electrical power system, including power plants, transmission lines, distribution networks, and demand profiles. Modern production cost modeling involves simulations of electricity markets based on forms of Unit Commitment (UC) and Economic Dispatch (ED) for generation and load scheduling. While the mathematical formulations in UC and ED can vary widely depending on the software package and its options [9], the cost attribution method proposed here is agnostic to the choice of their formulation. Power cost modeling typically requires alternating day-ahead UC and real-time ED market optimizations over a fixed horizon, whether it is a day, month or year. Our method can be used for most of the common simulation frameworks, however, for the purpose of illustration, we consider a day-ahead and real-time optimization cycle framework executed on a daily basis. The day-ahead simulation consists of a single UC optimization step that determines 48 hours of system operations subject to load and renewable production forecasts. Subsequently the real-time simulation consists of 24 ED on the hourly timescale, which re-optimize system operations subject to actual load and renewable power productions at each hour with a h -hour look-ahead horizon. A common choice of h can vary from 1 to 4 hours. Because many baseload generators are inflexible at short timescales, the real-time simulation treats day-ahead baseload UC profiles as fixed constraints.

Given the actual and day-ahead forecasts of the load demand and renewable production, our approach to the operational cost attribution is top-down. We

first solve the day-ahead UC using the load and renewable generation forecasts. Following the operational constraints in the UC profile, we then assess the system-level costs by solving 24 hourly EDs twice, once with the actual data and the second time with the day-ahead forecasts as the actual data. Finally, we attribute the system-wide cost difference between the two ED to individual load and production assets.

To be more specific, we consider a system with M dispatchable (conventional) power plants and denote their power generation at time t by

$$\mathbf{p}_t = (p_t^1, p_t^2, \dots, p_t^M)^T \in \mathbb{R}_+^M.$$

For each $\tau = 1, 2, \dots, 24$, the ED model with look-ahead horizon h can be written as,

$$\begin{aligned} \min \quad & \sum_{t=\tau}^{\tau+h} \left(\sum_{g=1}^M C_g^t(p_t^g) + C_p^t(\mathbf{p}_t) \right) \\ \text{s.t.} \quad & (\mathbf{p}_\tau^T, \mathbf{p}_{\tau+1}^T, \dots, \mathbf{p}_{\tau+h}^T) \in \Omega \end{aligned} \quad (3)$$

where C_g^t is the operational cost associated with generation unit g at time t , C_p^t is a penalty term accounting for loss of load, unmet reserve requirement, etc. at time t , and Ω is the feasible set representing the system constraints (load balance, generation limits, ramping capacities etc). The costs C_g^t and C_p^t as well as the feasible set Ω are determined by the grid configuration, and so are the load demand profile, renewable generation data and the initial grid operating condition. In other words, given the load demand profile, renewable generation data and initial state of the dispatchable units \mathbf{p}_0 , one can solve (3) and obtain an optimal operational cost (assuming that the model is feasible).

Without loss of generality, we set the look-ahead horizon to $h = 1$, and we assume that the system has L load buses, and N renewable power plants with load demand profile and power generation capacity at time t given by $\mathbf{d}_t = (d_t^1, d_t^2, \dots, d_t^L)^T \in \mathbb{R}_+^L$ and $\mathbf{q}_t = (q_t^1, q_t^2, \dots, q_t^N)^T \in \mathbb{R}_+^N$ respectively. Now the ED model at time τ can be viewed as a mapping F_τ :

$$\mathbb{R}_+^M \times \mathbb{R}_+^L \times \mathbb{R}_+^N \ni (\mathbf{p}_{\tau-1}, \mathbf{d}_\tau, \mathbf{q}_\tau) \rightarrow F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1} \\ \mathbf{d}_\tau \\ \mathbf{q}_\tau \end{pmatrix} \in \mathbb{R}_+, \quad (4)$$

providing the overall cost to the system. The goal of the proposed method is to quantify accurately the contributions of each individual load and each generation

asset to the system-wide cost difference computed with the day-ahead forecasts and the actual data, namely, for $\tau = 1, 2, \dots, 24$.

$$C_\tau^{act} - C_\tau^{fcst} := F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1}^{act} \\ \mathbf{d}_\tau^{act} \\ \mathbf{q}_\tau^{act} \end{pmatrix} - F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1}^{fcst} \\ \mathbf{d}_\tau^{fcst} \\ \mathbf{q}_\tau^{fcst} \end{pmatrix}. \quad (5)$$

2.3. Cost Attribution of PCM

We assume that F_τ is differentiable almost everywhere, and let \mathcal{C} be a straight line given by

$$\vec{r}_\tau(\lambda) = \begin{pmatrix} \mathbf{p}_{\tau-1}^{fcst} + \lambda(\mathbf{p}_{\tau-1}^{act} - \mathbf{p}_{\tau-1}^{fcst}) \\ \mathbf{d}_\tau^{fcst} + \lambda(\mathbf{d}_\tau^{act} - \mathbf{d}_\tau^{fcst}) \\ \mathbf{q}_\tau^{fcst} + \lambda(\mathbf{q}_\tau^{act} - \mathbf{q}_\tau^{fcst}) \end{pmatrix}, \quad 0 \leq \lambda \leq 1 \quad (6)$$

then by the fundamental theorem of calculus for line integral, we can write

$$\begin{aligned} F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1}^{act} \\ \mathbf{d}_\tau^{act} \\ \mathbf{q}_\tau^{act} \end{pmatrix} - F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1}^{fcst} \\ \mathbf{d}_\tau^{fcst} \\ \mathbf{q}_\tau^{fcst} \end{pmatrix} &= \int_{\mathcal{C}} \nabla F_\tau \begin{pmatrix} \mathbf{p}_{\tau-1} \\ \mathbf{d}_\tau \\ \mathbf{q}_\tau \end{pmatrix} \cdot d\vec{r} \\ &= \begin{pmatrix} \mathbf{p}_{\tau-1}^{act} - \mathbf{p}_{\tau-1}^{fcst} \\ \mathbf{d}_\tau^{act} - \mathbf{d}_\tau^{fcst} \\ \mathbf{q}_\tau^{act} - \mathbf{q}_\tau^{fcst} \end{pmatrix} \cdot \int_0^1 \nabla F_\tau(\vec{r}_\tau(\lambda)) d\lambda. \end{aligned}$$

In particular, the IG assigns to the initial state of dispatchable generator m the attribution:

$$C_{init}^{m,\tau} = (p_{\tau-1}^{m,act} - p_{\tau-1}^{m,fcst}) \int_0^1 \frac{\partial F_\tau}{\partial p_{\tau-1}^m}(\vec{r}_\tau(\lambda)) d\lambda \quad (7)$$

to the load asset l ,

$$C_{load}^{l,\tau} = (d_\tau^{l,act} - d_\tau^{l,fcst}) \int_0^1 \frac{\partial F_\tau}{\partial d_\tau^l}(\vec{r}_\tau(\lambda)) d\lambda \quad (8)$$

and to the renewable generator n

$$C_{renew}^{n,\tau} = (q_\tau^{n,act} - q_\tau^{n,fcst}) \int_0^1 \frac{\partial F_\tau}{\partial q_\tau^n}(\vec{r}_\tau(\lambda)) d\lambda. \quad (9)$$

We also have

$$C_{act}^\tau - C_{fcst}^\tau = \sum_{m=1}^M C_{init}^{m,\tau} + \sum_{l=1}^L C_{load}^{l,\tau} + \sum_{n=1}^N C_{renew}^{n,\tau}. \quad (10)$$

Equation (10) is called the *Completeness Axiom*. Basically, it says that the attributions to the initial states of the dispatchable generators, the load demands

and the renewable generations add up to the difference between the cost of operating the grid under the day-ahead forecasts and the actual data. The assumption of differentiability of F_τ is obviously a strong condition and can be difficult to check in practice, especially as it depends upon the grid configuration and the ED model formulation. Nevertheless, we observe empirically that indeed the Completeness Axiom holds in our experiments (see Table 3).

Another important property which is desirable for attribution methods is *Symmetry Preserving*. In the present context of production cost modeling this property can be understood in the following way. A function is said to be symmetric with respect to two input variables if the value of the function remains unchanged when the values of the variables are interchanged. Accordingly, an attribution method is said to be *Symmetry Preserving*, if whenever the function is symmetric with respect to two input variables, these input variables receive identical attributions. For a power grid system, *Symmetry Preserving* suggests that if two assets have the same impact in terms of the operational cost under all circumstances, they should receive identical attributions. For instance, if two wind farms are connected to the same bus and always produce the same amounts of power, their cost attributions by the proposed method should be identical. At an intuitive level, such a requirement is desirable for a practical attribution to be fair and equitable, and it holds in the attribution method introduced in this paper.

The original introductions and implementations of the IG also mentioned a few other desirable properties, such as *Sensitivity*, *Linearity*, *Implementation Invariance*, etc, however, they are not relevant to models of grid management, and we chose to ignore them in our context.

2.4. Computation of the Gradient

The integrals in (7)-(9) can be efficiently approximated by adaptive quadrature, e.g. adaptive trapezoidal rule. Adaptivity usually requires less than 100 function (integrand) evaluations to reach a desire error threshold (within 5%).

For the gradient, noting that the ED model we consider is a linear program, we compute the partial derivative w.r.t. a quantity of interest by the dual value of the constraint corresponding to that quantity. For example, let us assume that at time τ the generation capacity of the n -th renewable generator is \bar{q}_τ^n .

To compute the partial derivative $\frac{\partial F_\tau}{\partial \bar{q}_\tau^n}$, we first add a constraint

$$q_\tau^n \leq \bar{q}_\tau^n \quad (11)$$

to the ED model (3). Once the ED model is optimized, we obtain the desired partial derivative as the dual value of the constraint (11). This is what is usually called the *shadow price*

3. Case Study: the RTS-GMLC Grid Model

In this section, we present numerical results of the cost attribution method for the RTS-GMLC grid model whose characteristics can be found at <https://github.com/GridMod/RTS-GMLC>. The RTS-GMLC grid contains 73 buses, 157 generators (including 4 wind and 57 solar stations) and 120 transmission lines. We use the production cost modeling tool Vatic v0.4.1-a1 (<https://github.com/PrincetonUniversity/Vatic/releases/tag/v0.4.1-a1>) and GUROBI solver v10.0.1 [5] to run UC and ED optimization. Vatic applies mixed-integer linear programming optimization to power grid formulations created using the software package EGRET [10]. Our Python implementation of the cost attribution algorithm with detailed examples is available at <https://github.com/PrincetonUniversity/PGrisk>[18]. For all examples presented in this section, we choose a maximum of $2^{12} = 4096$ for the adaptive trapezoidal quadrature nodes with 5% error threshold. Below, we summarize the formulations in our experiments in Table 1 and refer to [10, 9] for more details.

In our initial case study, we simulate the RTS-GMLC grid by executing UC with day-ahead forecasts and ED with both forecasted and actual data. Across each day of 2020, we calculate cost attributions comparing forecasted (baseline) and actual data over 24-hour periods, then average these attributions hourly. Table ?? presents the top 10 renewable generators with the highest average cost attributions. Notably, wind generators rank highest due to their substantial capacity, followed by solar generators, which operate predominantly during daylight hours. Specifically, the capacities for 317-WIND-1, 309-WIND-1, and 319-PV-1 are 799.1MW, 148.3MW, and 188.2MW, respectively. These renewable generators are distributed across five sub-areas within the RTS-GMLC grid. Additionally, we calculated the sum of generation deficiencies (actual - forecast), the sum of cost attributions, and the average shadow price within each sub-area for 2020-06-18. They are provided in Figure 1.

	UC	ED
Status	garver_3bin_vars	garver_3bin_vars
Power	garver_power_vars	garver_power_vars
Reserve	garver_power_avail_vars	MLR_reserve_vars
Generation	pan_guan_gentile_KOW_generation_limits	MLR_generation_limits
Ramping	damcikurt_ramping	damcikurt_ramping
Production	KOW_production_costs_tightened	CA_production_costs
Up/Down	rajan_takriti_UT_DT	rajan_takriti_UT_DT
StartUp	KOW_startup_costs	MLR_startup_costs
Network	ptdf_power_flow	ptdf_power_flow

Table 1: Formulation used in UC and ED models.

Generator	Bus	Attribution	Shadow Price
317-WIND-1	317	54625.9	-159.2
122-WIND-1	122	44742.8	-151.1
303-WIND-1	303	32993.6	-148.8
309-WIND-1	309	6318.4	-168.7
319-PV-1	319	768.9	-161.6
215-PV-1	215	481.5	-154.2
113-PV-1	113	380.0	-153.1
313-PV-2	313	257.2	-164.6
312-PV-1	312	237.7	-165.2
310-PV-2	310	220.2	-164.3

Table 2: Top 10 renewable generators with the largest average cost attributions.




Figure 1: Generation deficiency, cost attributions and shadow prices for sub-areas in RTS-GMLC.

	number of nodes	run time (s)	relative efficiency gap
max	124	1307.9	5.3%
median	23	239.0	0.2%
mean	26.2	275.8	0.6%
std	18.0	195.8	0.8%

Table 3: Cost attribution performance.

Next, we examine the performance of the algorithm using Monte Carlo scenarios. The load demand and renewable production scenarios used for computing the cost attributions are generated by the open source Python package PGscen [4]. For each day in 2020, we first generated $K = 1000$ Monte Carlo scenarios for load demand and renewable (wind and solar) production. With the risk level $\alpha = 0.05$, the algorithm is then applied to obtain cost attributions for the worst $\alpha K = 50$ scenarios.

All the experiments are conducted on a 80-core cluster equipped with two Intel(R) Xeon(R) Platinum 8380 CPU @ 2.30GHz and 512GB memory. For each scenario, the cost attribution procedure is run on a single core with no parallelism. To examine the performance of the cost attribution method, we consider the following relative efficiency gap:

relative efficiency gap =

$$\frac{\max_{1 \leq \tau \leq 24} |C_{act}^{\tau} C_{fcst}^{\tau} \sum_{m=1}^M C_{init}^{m,\tau} \sum_{l=1}^L C_{load}^{l,\tau} \sum_{n=1}^N C_{renew}^{n,\tau}|}{\max_{1 \leq \tau \leq 24} |C_{act}^{\tau}|}$$

We report in Table 3 the average number of actual quadrature nodes, run time (in seconds) and the relative efficiency gap over $50 \times 364 = 18200$ scenarios from 2020-01-02 to 2020-12-30.

For the purpose of illustration, we present numerical results for the test day 2020-07-08. We performed the cost attribution computations for 1000 Monte Carlo

scenarios and we discuss the results for the load bus Abel, the wind farm 317-WIND-1 and the solar plant 324-PV-1. In each case, we show in Figures 2, 3 and 4 the histograms of the values of the scenarios (left) and their cost attributions (right) computed scenario by scenario. Figure 2 shows that on that day, the load forecast significantly over estimated the actual load, and the Monte Carlo generation engine (which does not have a crystal ball and could not use the value of the actual load to generate the scenarios) used a heavy tail distribution to produce reasonable values despite the forecast overshoot. The two histograms follow similar patterns. This should be expected due to the positive correlation between load demand and system costs. The two histograms follow similar patterns. This should be expected due to the positive correlation between load demand and system costs.




Figure 2: Scenarios and cost attributions for Abel at 2020-07-08 04:00.

Figure 3 shows that on that day, the wind forecast at this location was pretty accurate. Still, the patterns of the histograms seem to be mirror images of each other. This is consistent with the negative correlation between power production and system costs. The same remarks apply to Figure 4. On the other hand, it's worth noting that when a loss of load occurs during a specific hour within a scenario, the attribution of the associated cost is likely to be significantly amplified. This is due to the substantial penalty incurred for failing to meet the load demand. As illustrated in Figure 5, we can observe an example of this phenomenon where certain scenarios with low production are attributed with notably high costs. This emphasizes the critical importance of addressing load shedding and its cost implications in our analysis.




Figure 3: Scenarios and cost attributions for 317-WIND-1 at 2020-07-08 17:00.




Figure 4: Scenarios and cost attributions for 324-PV-1 at 2020-07-08 15:00.




Figure 5: Scenarios and cost attributions for 317-WIND-1 at 2020-07-08 19:00.

4. Risk-Averse Unit Commitment

4.1. Scenario-Based Risk Allocation

In this section, we demonstrate an application of the cost attribution algorithm within a simulation framework tailored for a risk-averse Unit Commitment (UC) approach. First, we introduce a risk allocation algorithm built upon the aforementioned cost attribution method. The primary objective of this algorithm is to assess the impact of uncertainties on daily operational costs of the system. This is achieved by incorporating the variability in both load and renewable generation

across numerous Monte Carlo scenarios. Initially, we solve the UC and ED problems using day-ahead forecasts. Subsequently, we use these scenarios as proxies for actual load and renewable production data and solve ED once again. The scenarios employed in this study are generated utilizing the model described in [3], as implemented in [4]. This model captures a high-dimensional spatial-temporal correlation structure and encompasses extreme and risky events.

To be more specific, let us assume that K Monte Carlo scenarios have been generated and a risk level $\alpha \in (0, 1)$ has been chosen. The risk allocation algorithm comprises the following steps:

1. The UC and ED are solved based on the day-ahead forecasted data.
2. For each scenario, its values are substituted for the actual data and we solve the 24 hourly ED models to obtain hourly costs.
3. We select the worst αK scenarios in terms of their total cost. We denote the set of the worst scenarios by \mathcal{S} .
4. For each scenario in \mathcal{S} , we calculate the cost attribution between the baseline model, which uses forecasted data, and the target model, which uses scenario data. This enables us to assign costs to all initial states of dispatchable generators, load demands, and renewable generators for every hour.
5. For a given asset (load demand or renewable generator) at any hour, we compute its risk score R as the average of its cost attributions across the scenarios in \mathcal{S} . For example, the risk score for the renewable generator n at hour τ is $R = \frac{1}{\alpha K} \sum_{j \in \mathcal{S}} C_{renew}^{n,\tau,j}$, where $C_{renew}^{n,\tau,j}$ is the cost attribution in the j -th scenarios.

4.2. Reliability Adjustment

As previously explained, our proposed method assigns a reliability risk score to each stochastic input during every hour of the day. While these stochastic inputs include various factors such as loads and outputs from renewable generation assets, our attention in this study will be concentrated on the renewable generation assets. To mitigate the risk associated with

the uncertainty of these renewable generators, our proposal is to modify their generation capacities whenever elevated risk levels are identified.

Specifically, given a risk level $\alpha \in (0, 1)$, we recall that the risk score $C_{renew}^{n,\tau}$ for the renewable generator n at hour τ is computed as the average costs attributions over a set of worst case scenarios. These costs are frequently linked to the under-production of the renewable generators relative to the forecasts. Consequently, our initial step involves the computation of a "per MWh" risk score R_τ^n through

$$R_\tau^n = \frac{C_{renew}^{n,\tau}}{q_\tau^{n,fcst} - \frac{1}{K} \sum_{j \in \mathcal{S}} q_\tau^{n,j}}$$

where $q_\tau^{n,j}$ is the generation capacity of the renewable generator n at hour τ in the j -th scenario. Following that, we introduce two positive parameters, denoted as \underline{R} and \overline{R} . Broadly speaking, these parameters are meant to represent the per MWh cost of the grid under usual conditions and in a situation where the generation needs to be completely disregarded. Subsequently, a generation capacity adjustment percentage can be calculated as follows:

$$r_\tau^n = \max \left(0, \min \left(1, \frac{R_\tau^n - \underline{R}}{\overline{R} - \underline{R}} \right) \right).$$

where min and max functions guarantees that $r_\tau^n \in [0, 1]$. Finally, the adjusted renewable generation capacity is determined by the equation:

$$q_\tau^{n,adj} = q_\tau^{n,fcst} - r_\tau^n (q_\tau^{n,fcst} - q_\tau^{n,min}).$$

Here $q_\tau^{n,min}$ represents the minimum (*guaranteed*) generation capacity of renewable generator n at hour τ across all scenarios. By design, when $r_\tau^n = 0$, indicating no need for adjustment, the equation simplifies to $q_\tau^{n,adj} = q_\tau^{n,fcst}$. Conversely, if $r_\tau^n = 1$, implying a necessity for maximum adjustment, the generation capacity is reduced to the minimum (*guaranteed*) level: $q_\tau^{n,adj} = q_\tau^{n,min}$. As an illustrative example, we demonstrate the generation capacity adjustments for 317-WIND-1 at 2020-07-08 during the hours 19:00 and 20:00. For this example, we choose $\underline{R} = 20$ and $\overline{R} = 200, 300, \text{ and } 500$. The histograms of generation scenarios are presented in Figure 6, and the relevant quantities for computing the adjustments are detailed in Table 4.

Hour	$q_{\tau}^{n,fcst}$	$C_{renew}^{n,\tau}$	R_{τ}^n	$\bar{R} = 200$		$\bar{R} = 300$		$\bar{R} = 500$	
				r_{τ}^n	$q_{\tau}^{n,adj}$	r_{τ}^n	$q_{\tau}^{n,adj}$	r_{τ}^n	$q_{\tau}^{n,adj}$
19	107.0	15630.2	210.8	95.4%	4.9	63.6%	38.9	38.2%	66.2
20	262.0	6163.0	35.0	7.5%	242.3	5.0%	248.9	3.0%	254.1

Table 4: Generation capacity adjustments for 317-WIND-1 at hours 19 and 20 of 2020-07-08.




Figure 6: Scenarios and the worst case scenarios for two wind farms on 2020-04-26.

4.3. Simulation results

To examine the impact of the proposed risk-averse UC approach, we run simulations for the RTS-GMLC grid from 2020-01-02 to 2020-12-30. For the purpose of comparison, in addition to our approach, we include three benchmark approaches, in which the standard alternating UC and ED are performed with a system spinning reserve equal to 10%, 20% and 30% of the system forecasted load demand. In the benchmark approaches, the load mismatch (over-generation and loss of load) and reserve shortfall penalties in the objective function of the UC model are selected to be 10000\$/MWh and 1000\$/MWh. In contrast, our approach adopted a reserve factor of 5% and used the same load mismatch and reserve shortfall penalties as the benchmark methods. For generation capacity adjustments, \underline{R} was set to 20\$, and the value of \bar{R} was varied by 200\$, 300\$, and 500\$. In Figure 7, we show the daily system operational costs and the load shedding patterns for the entire year using different UC approaches. Additionally, we report the average daily operational costs and total yearly load shedding in the same plot. We observe that the proposed risk-averse UC approach is capable of largely reducing load shedding (which is the largest risk to the grid operation) under all choices of \bar{R} . Specifically, selecting the least aggressive parameter $\bar{R} = 500$ results in a 1525 MWh

loss of load, roughly a quarter of the load shedding experienced with a 30% reserve factor. The production costs using the risk-averse UC approach are approximately \$1.47 million, which is much lower than the \$1.53 million with a 30% reserve factor and only slightly higher than with a 20% reserve factor.

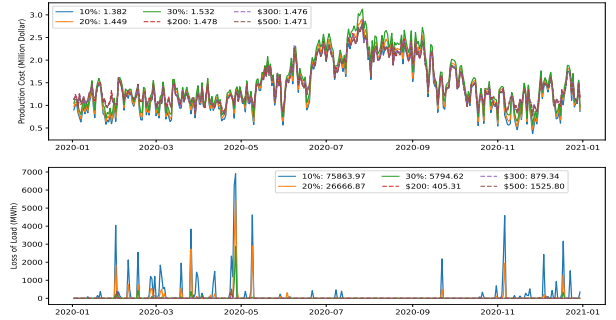


Figure 7: Production costs and loss of load in RTS-GMLC system using different UC from 2020-01-02 to 2020-12-30.

In Figure 8, we create the same plot from 2020-01-01 to 2020-04-30, a shorter period in which activities like load shedding happen more frequently. For 2020-03-26, 2020-04-26 and 2020-04-27, it is clear that the risk-averse UC hedges the risk against the worst case scenarios by depending less on the renewable generations (and thus higher production cost) to avoid possible load shedding due to uncertainties in the day-ahead forecasts. The loss of load for all choices of \bar{R} is less than 1000 MWh, whereas reserve factors of 10%, 20%, and 30% result in unmet load demands ranging from 4000 to 34000 MWh. The production costs using the risk-averse UC approach, however, are comparable to those incurred with a 20% reserve factor.


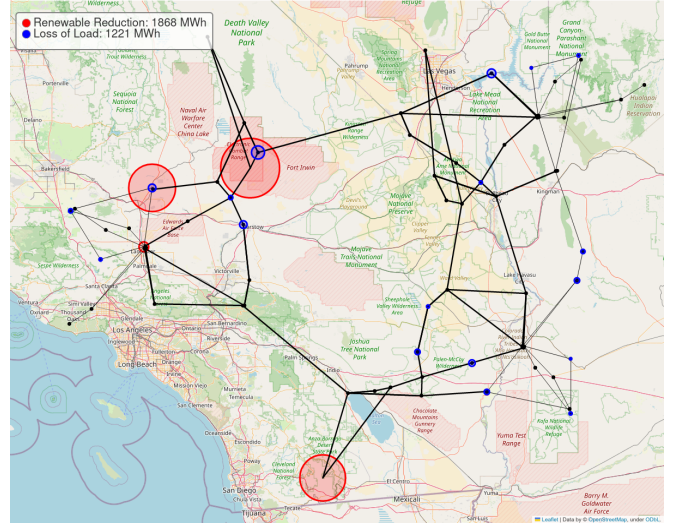


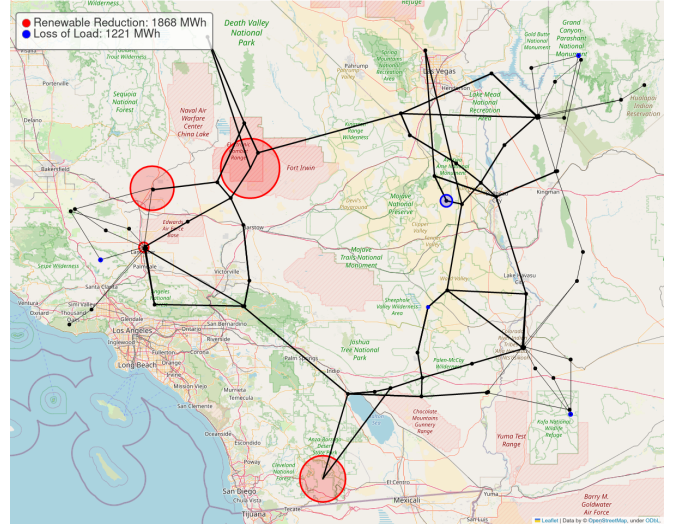
Figure 8: Production costs and loss of load in RTS-GMLC system using different UC for 2020-03-01 to 2020-04-30.

4.4. Load Loss Abatement

To further explore how the renewable generation adjustment assists in mitigating risk, we present a more comprehensive comparison with the use of a fixed percentage of the load demand spinning reserves factor (referred to as the risk-neutral UC) on 2020-04-26. During the 24-hour period, when operating the ED with the risk-averse UC, there were no instances of load loss. However, running it with the risk-neutral UC using 5% and 30% reserve factors resulted in a total load loss of approximately 860 and 7903 MWh, respectively. This loss primarily occurred due to an overestimation of wind power generation in the day-ahead forecasts. The day-ahead predictions estimated wind power plants to generate as much as 37,046 MWh throughout the day, while the actual real-time production was only 12590 MWh. Notably, at 21:00, there was an over-forecast of wind power production by about 2019 MWh. In Figure 9, we illustrate the RTS-GMLC grid on a geographical map, indicating buses with load shedding (depicted as blue circles) in simulations using the risk-neutral UC. Additionally, we calculate the difference in renewable power dispatch in simulations using the risk-averse UC, revealing a significant reduction in renewable power dispatch due to the associated capacity adjustments, shown as differences represented by red circles on the map.



(a)



(b)

Figure 9: Comparison of risk-neutral UC using 5% (A) and 30% (B) reserve requirements with risk-averse at 2020-04-26 21:00:00.

5. Conclusion

In conclusion, this paper offers a new perspective on addressing uncertainty and associated risks in power system management, driven by the increasing integration of renewable energy sources. We propose a novel algorithm that leverages the discrepancies between forecasted and actual values to quantify the financial risks associated with uncertainty. Inspired by Integrated Gradients (IG), a renowned method in machine learn-

ing, our algorithm attributes the contributions of stochastic components to differences in system costs. By treating the power production cost model as a function that takes inputs of random quantities like load demand and renewable generation, we provide a practical approach for understanding the impact of variability on grid operational costs.

Moreover, we demonstrate the utility of our approach in a risk-averse unit commitment framework, where adjustments to renewable generator capacities are made based on Monte Carlo scenarios. This framework mitigates system risk by reducing reliance on highly uncertain renewable productions, thereby enhancing grid reliability. On the other hand, the overall operational costs will not increase significantly under this framework (see Figures 7 and 8) as it identifies critical times and locations where highly uncertain renewable productions occur.

The results of our simulations on the RTS-GMLC system serve as a compelling validation of the effectiveness of our algorithms. Furthermore, they highlight the superior performance of the risk-averse UC framework, particularly in its ability to significantly reduce the occurrence of load shedding. As part of our future work, we plan to further refine and extend our approach, conducting tests on larger and more realistic grid systems to verify its performance in more complex and practical scenarios.

6. Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

7. Acknowledgments

The authors were partially supported by ARPA-E grants DE-AR0001289 and DE-AR0001390 under the PERFORM program of the US Department of Energy. We would like to thank Mike Ludkovski (University of California at Santa Barbara) and Glen Swindle (Scoville Risk Partners) for enlightening conversations on the content of the paper; and Michal Grzadkowski (Princeton University) for software development and implementation.

References

- [1] Miguel Asensio and Javier Contreras. Stochastic unit commitment in isolated systems with renewable penetration under cvar assessment. *IEEE Transactions on Smart Grid*, 7(3):1356–1367, 2015.
- [2] Dimitris Bertsimas, Eugene Litvinov, Xu Andy Sun, Jinye Zhao, and Tongxin Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE transactions on power systems*, 28(1):52–63, 2012.
- [3] René Carmona and Xinshuo Yang. Joint granular model for load, solar and wind power scenario generation. *IEEE Transactions on Sustainable Energy*, pages 1–13, 2023.
- [4] Michal Grzadkowski, Carmona Rene, and Xinshuo Yang. Pgsцен. <https://github.com/PrincetonUniversity/Pgsцен/>, 2023.
- [5] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023.
- [6] Yuping Huang, Qipeng P Zheng, and Jianhui Wang. Two-stage stochastic unit commitment model including non-generation resources with conditional value-at-risk constraints. *Electric Power Systems Research*, 116:427–438, 2014.
- [7] Rabih A Jabr. Robust self-scheduling under price uncertainty using conditional value-at-risk. *IEEE Transactions on Power Systems*, 20(4):1852–1858, 2005.
- [8] Ruiwei Jiang, Jianhui Wang, and Yongpei Guan. Robust unit commitment with wind power and pumped storage hydro. *IEEE Transactions on Power Systems*, 27(2):800–810, 2011.
- [9] Bernard Knueven, James Ostrowski, and Jean-Paul Watson. On mixed-integer programming formulations for the unit commitment problem. *INFORMS Journal on Computing*, 32(4):857–876, 2020.
- [10] Bernard Knueven et al. Egret. <https://github.com/grid-parity-exchange/Egret>, 2023.
- [11] Narayana Prasad Padhy. Unit commitment-a bibliographical survey. *IEEE Transactions on power systems*, 19(2):1196–1205, 2004.
- [12] Arvind Shrivats, Ronnie Sircar, and Xinshuo Yang. Quantifying renewables reliability risk in modern and future electricity grids. *Submitted*, 2024.
- [13] Elina Spyrou, Ben Hobbs, Deb Chattopadhyay, and Neha Mukhi. How to assess uncertainty-aware frameworks for power system planning? *IEEE Transactions on Energy Markets, Policy and Regulation*, 2024.
- [14] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pages 3319–3328. PMLR, 2017.
- [15] Aidan Tuohy, Peter Meibom, Eleanor Denny, and Mark O’Malley. Unit commitment for systems with significant wind penetration. *IEEE Transactions on power systems*, 24(2):592–601, 2009.
- [16] Qianfan Wang, Jean-Paul Watson, and Yongpei Guan. Two-stage robust optimization for nk contingency-constrained unit commitment. *IEEE Transactions on Power Systems*, 28(3):2366–2375, 2013.
- [17] Lei Wu, Mohammad Shahidehpour, and Tao Li. Stochastic security-constrained unit commitment. *IEEE Transactions on power systems*, 22(2):800–811, 2007.

- [18] Xinshuo Yang. PGrisk. <https://github.com/PrincetonUniversity/PGrisk.git>, 2023.
- [19] Yao Zhang, Jianxue Wang, Tao Ding, and Xifan Wang. Conditional value at risk-based stochastic unit commitment considering the uncertainty of wind power generation. *IEE Generation, Transmission & Distribution*, 12(2):482–489, 2018.
- [20] Qipeng P Zheng, Jianhui Wang, and Andrew L Liu. Stochastic optimization for unit commitment—a review. *IEEE Transactions on Power Systems*, 30(4):1913–1924, 2014.