

Quantifying Renewables Reliability Risk in Modern and Future Electricity Grids

Arvind Shrivats^{*†} Ronnie Sircar^{‡†} Xinshuo Yang^{*†}

First version: May 2022. Final version September 2025.

Abstract

We propose and implement a methodology to quantify, allocate and account for the risk introduced to electricity production from the unpredictable intermittency of renewable resources such as wind and solar. Incorporating this stochasticity into grid risk management is viewed by the industry (which has remained almost entirely tethered to a deterministic viewpoint, and in particular to weather forecasts) as increasingly crucial, as we aim for greater renewables penetration to reduce dependence on carbon-emitting fuels. Our methodology involves feeding Monte Carlo simulations of solar, wind and demand into a grid optimization software that emulates the performance and costs of the Texas electricity grid. This outputs a distribution of running costs, from which we can numerically extract a measure of system (grid) risk. The more challenging part is to allocate this risk back (top down) to the individual renewable assets to assign them a reliability cost. This adapts existing approaches for the risk allocation problem related to Shapley values, but is computationally intensive. We show results, project to potential future grids, and propose a way to incorporate the reliability costs back into the day ahead bid curve and thereby re-optimize unit commitment and economic dispatch of assets taking into account the probabilistic nature of supply from renewables.

Keywords: Renewable energy, electricity grid, risk allocation, reliability costs, CVaR, Shapley values

1 Introduction

As is by now well-documented, the unpredictability of electricity production from renewable generators such as solar or wind arising from their sources' intermittency is of increasing significance as we transition to electric grids that are much cleaner in terms of carbon emissions. We refer, for instance, to [16], and in particular to the sections entitled *Public Policies*

^{*}Formerly Operations Research & Financial Engineering Department, Princeton University.

[†]Work supported by U.S. Department of Energy award DE-AR0001289.

[‡]Operations Research & Financial Engineering Department and Andlinger Center for Energy and the Environment, Princeton University.

and Clean Energy Technologies Are Reshaping the Grid and Planning Challenges with the Growing Reliance on Intermittent Generation. It therefore becomes essential to address and quantify the stochastic (uncertain, unpredictable or unreliable) nature of renewable assets, beyond just crude forecasts (expectations), and their impact on a grid’s ability to meet power demand.

From the vantage of a system operator charged with reliable electricity production for, say, the day ahead, there are a number of obstacles to incorporating probabilistic information into their forward planning procedures. First, the “black box” optimization software used to solve the AC optimal power flow (OPF) model and to compute unit commitments (UC) and economic dispatch (ED) is inherently dealing with a deterministic system, where forecasts (of production and demand) are treated as actuals; that is, random variables are effectively replaced by their expectations (or some function thereof). More precisely, the forecasts are purchased from firms analyzing weather patterns, and translated into expected electricity production from each renewable generator. They are not computed from a stochastic model, such as we will discuss using here, to provide more refined statistical information than just the mean.

However, the black boxes cannot be materially altered (at least in the US) without major revisions to regulations governing electric grids, which may need congressional involvement. As described in a 2012 Federal Energy Regulatory Commission (FERC) article [5]: “The ACOPF is at the heart of Independent System Operator (ISO) power markets, and is solved in some form every year for system planning, every day for day-ahead markets, every hour, and even every 5 minutes. It was first formulated in 1962, and formulations have changed little over the years.”

As such, it would be ideal if the ISO could solve a *stochastic optimization* problem of the form

$$\min_{a \in \mathcal{A}} \mathbb{E}\{f_a(X)\},$$

where X is random, describing intermittent production (and demand uncertainty), and a are decision variables (unit commitment and dispatch among others) from some permissible action set \mathcal{A} , which includes power flow and other constraints. A decision a results in system cost $f_a(X)$ which depends on the real time outcome X of renewables production the day ahead. See, for instance, [17] for such an approach to cope with wind production uncertainty in California, demonstrating unit commitment policies that outperform the approach actually employed in practice.

In terms of the simplified model notation above, what ISOs currently do is akin to solving the *deterministic optimization* problem

$$\min_{a \in \mathcal{A}} f_a(\mathbb{E}\{X\}),$$

where random production (and demand) X are reduced to their forecasted values $\mathbb{E}\{X\}$. Given that we are practically tied to using ISOs’ grid software that solves the deterministic problem, we can only experimentally vary the inputs into the black box, which is what we will discuss with a Monte Carlo simulation approach.

A second obstacle is that renewables bid-in production costs are zero, compared with more reliable thermal generators which have fuel and operating costs. As such, when commitment

decisions are made, renewables, being at the bottom of the bid stack, are typically committed at their forecasted capacity. But if actual production falls short of expectations, thermal generators (“peakers”) may have to be called up in a pinch, which can be costly and even more environmentally detrimental than if they had been planned for earlier. Additionally, renewable resources are not dispatchable, that is they may not deliver on forecasts due to weather unpredictability, and they cannot reliably be called upon when demand exceeds expectations. The uncertainty that we have with respect to how much wind or solar power can be produced in some future time span creates issues for grid management, as a failure to provide enough supply to match the demand for electricity can result in sizeable costs and, potentially, human loss.

This motivates modernizing the risk management processes and strategies of electricity grid operators to account for the forthcoming change in the production mix. To that end, we will quantify the impact of uncertainty on daily system operation, ultimately producing a distribution of system-wide daily operational cost which incorporates the stochasticity present in load, renewable generation, and other uncertainties. We then apply a risk measure to this distribution to obtain a single metric which provides an indication of overall system risk. This is then allocated and distributed to the assets based on an assessment of their contribution to it.

We develop methods to allocate the overall operational cost to individual units or grid sectors. This is a risk allocation problem which has been studied in the financial and actuarial mathematics literature, and we develop algorithms for its efficient computation from Monte Carlo simulations of load and supply uncertainty fed through a grid optimization software. We will utilize the theory of capital allocation (see, for instance [13, 21]) to decompose system-level risk to individual assets or cohorts of assets.

The framework of capital/risk allocation incorporates the concepts of “fairness” (for example ensuring a consistent definition of risk diversification across any subset of assets, allowing reliability grouping), “stability” (collective risk is always less than the sum of individual costs) and *monetary* cost (so that all risk is expressed in dollar amounts—essential for engaging financial participants and accounting-driven regulations). The rules are designed to be resistant to market manipulation or strategic gaming and have been extensively validated across a range of US and global regulatory frameworks. As such, they will offer all stakeholders a clear assurance of validity, an obvious prerequisite for piloting and market adoption.

The starting point of the advocated top-down risk allocation is that the overall grid reliability index (expressed as a risk measure) is less than the sum of individual asset risk scores if they were to be measured independently. Risk allocation thus directly captures the definition of reliability cost as a conditional quantity ascribed not just based on marginal asset performance, but relative to the overall grid. Therefore, assets that might be unreliable on their own could in fact be highly beneficial to the grid if they are negatively correlated to other sources of variability. This is precisely the hedging concept that underlies our approach. An important challenge is for developers and planners to be able to *project* likely reliability costs of a new asset that is being constructed.

At the heart of our proposed system risk allocation procedure is computation of Shapley values, named after (Nobel laureate) Lloyd Shapley, who introduced it in the seminal work [20]. In a grid context, it has been theorized and applied to a small-scale synthetic example in

[15], and the related Aumann-Shapley approximation has been used to allocate transmission costs in [12]. There are also connections between risk allocation and game-theory [8]. We give the details of our approach in Section 2.4. Section 3 contains our computational results, the risk allocations first for a small synthetic grid, and then the 2020 and projected 2030 Texas grids. In Section 4, we detail how we turn the allocations into reliability cost curves and re-run the unit commitment on this basis, showing improved outcomes. We conclude and discuss future directions in Section 5. This work is part of research from the Princeton team ORFEUS¹, funded by ARPA-E under its program PERFORM².

2 ORFEUS Methodology

In this section, we describe our methodology for quantifying reliability risk from renewables. The statistical model behind the previously-developed scenario generator that we employ is summarized in Section 2.1; our OPF software and its measurement of system costs are described in Section 2.2; in Section 2.3, we detail how we quantify system costs by a risk measure; and in Section 2.4, we present our risk allocation methodology for assigning the system risk back to the individual assets.

We have for our use a stochastic model for load, solar, and wind generation, such as the one detailed in [6]. In our study, we use the synthetic Texas power grid constructed in [4]. The original version, referred to as the 2020 grid, is a 6717-bus network covers the geographic footprint of the electric reliability council of Texas (ERCOT) and serves 74 GW of peak load across the majority of the U.S. state of Texas. Additionally, we consider a future projection of the grid network, referred to as 2030 grid, and also produced by the PERFORM team at Texas A&M. This version comprises 7132 buses and assumes a complete phase-out of coal plants in favor of renewable alternatives. The 2030 grid’s micro-structure is designed to handle about 89 GW of peak load across Texas.

We will also show results from a much smaller (synthetic) test grid RTS-GMLC, which covers parts of Southern California and Nevada. It has 60 renewable assets: 25 solar, 31 rooftop solar and 4 wind.

2.1 Load and Renewable Power Scenarios

In our simulations, once the UC optimization is obtained, we solve ED problems with Monte Carlo scenarios generated with the open source Python package PGScen [10] based on [6]. The output Monte Carlo scenarios include zone-level load demands and generator-level solar and wind productions at hourly frequency. As described in [6], the scenario generation algorithms capture the intricate dependencies and interactions between temporal and geographical components. The main features of the simulation engine include:

1. The use of generalized Pareto distributions when heavy-tailed distributions are detected.

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

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

2. Transformation of marginal distributions into Gaussian in the spirit of Gaussian copulas.
3. Use of separable covariance structures to disentangle temporal and spatial contributions to the correlations.
4. Development of a marginal statistical model for solar power production based on Principal Component Analysis.

Due to the training data used in `PGScen`, the output solar and wind power production scenarios are restricted to the locations in the input data [19]. This limitation prevents us from directly applying these scenarios to the Texas 2020 and 2030 grids because the geographical locations and generator capacities between [19] and the Texas grids do not match one-to-one. To overcome this hurdle, we map the solar and wind generators in the scenarios to the nearest solar and wind generators in the grid data, using a proper scaling factor to match the capacity.

2.2 OPF Software `Vatic` and System Costs

We have adapted the grid modeling software of `Egret` and `Prescient` developed by Lawrence Livermore National Laboratories (LLNL) [14] for use in a high performance computing environment, which we call `Vatic`³. `Vatic` is a Python package for running simulations of a power grid using the PJM framework consisting of alternating day-ahead unit commitment (UC) and real-time economic dispatch (ED) steps. `Vatic` was originally designed as a lightweight adaptation of `Prescient`; it likewise applies mixed-integer linear programming optimization as implemented in `Pyomo` to power grid formulations created using `Egret`.

We next run `Vatic` for a specific day. The day-ahead unit-commitment (UC) step is run once using load and renewable production forecasts. Then, for each of hundreds or thousands of Monte Carlo simulations of probability-weighted scenarios of load and renewables outcomes (actuals) produced by `PGScen` (described in Section 2.1), we run the real-time economic dispatch (ED). This outputs, among other quantities, the total system costs of operating the grid during each hour of the day. Specifically, this is

$$\text{Total System Costs} = \text{Fixed Costs} + \text{Variable Costs} + \lambda \times \text{Load Shedding},$$

where Fixed and Variable Costs represent the fixed and variable production costs from generators, Load Shedding is the amount, in MWh, from the commitment and dispatch which could not be serviced without out-of-market corrections, and λ is a monetary penalty per unit MWh of load shedding (\$/MWh). In our experiments, we will take $\lambda = \$10\text{K}/\text{MWh}$. This is consistent in magnitude with the price cap in ERCOT, which was \$9,000/MWh in 2020.

In scenarios where renewables underperform forecasts, these costs will be high due to the need to pull in alternative generators, while in other scenarios where actuals turn out to be close to forecasts, the system costs may be lower. Thereby, we obtain a Monte Carlo distribution of system costs for each hour. In our context, we are looking to allocate the risk of excess grid operational costs due to stochasticity in renewables production and load,

³<https://github.com/PrincetonUniversity/Vatic>




Figure 1: Total system costs in the BAU (gray) and PW (red) settings on RTS for June 30 2020 (left) and January 20 2020 (right). Dashed vertical lines represent 95th percentiles (VaR at 5%).

over a **perfect world (PW)** situation in which there is no uncertainty from renewables intermittency and actuals exactly match up with forecasts.

The former **business as usual (BAU)** system costs are described by the gray histograms in Figure 1, while the PW system costs are generated by re-running the Monte Carlo simulations through *Vatic*, but with actual renewable production set to meet their forecasted values. These are shown in the red histograms in the figure, and reflect the only remaining uncertainty from load, where we also use the term “idealized” to refer to the removal of renewables uncertainty.

The dashed vertical lines represent the 95th percentiles of each histogram. From this, we can see that the BAU grid has higher variance and tail risk. This can be thought of as representing the worst case scenarios that occur due to the stochasticity of renewable generators. In contrast, the PW grid has significantly lower variance, and significantly lower tail risk, as measured by the 95th percentile (or the VaR at the 5% significance level). This shows the risks created by the unreliability of renewable assets.

The difference between these system costs is the excess grid operational cost due to renewables stochasticity. We suppose that this excess system cost, which depends on the random outcomes of load and renewable supply, is described by a bounded random variable X with cumulative distribution function $F_X(x) = \mathbb{P}\{X \leq x\}$:

$$X = \text{Total System Costs}^{BAU} - \text{Total System Costs}^{PW}. \quad (2.1)$$

The empirical probability density function of X is known from the simulated histograms such as in Figure 1.

2.3 Risk Measures

Our goal is to address the extreme worst-case scenarios, when renewables under-performance leads to large excess system costs, that is, the right-tail of X . We map an empirical distribution of excess system costs X in a given time period into a single metric which quantifies

the system risk in that time period. In the literature of finance and actuarial science, such a mapping is referred to as a *risk measure*. In modern risk management, much emphasis is placed on *coherent* risk measures, in particular, which were introduced in [1].

A function ρ is said to be a *coherent risk measure* if it has the following four properties:

1. **Monotonicity:** if X, Y are two random variables and $X \geq Y$ almost surely, then $\rho(X) \geq \rho(Y)$. Intuitively, this states that if the costs X always exceed the costs Y , the excess system risk of X must be greater than the excess system risk of Y .
2. **Risk measure is monetary:** For any $m \in \mathbb{R}$, $\rho(X - m) = \rho(X) - m$. Lowering excess costs X by a constant monetary amount $m > 0$ reduces the excess system risk by the same amount. (Similarly for $m < 0$, increasing costs).
3. **Sub-additivity/Convexity:** if X, Y are two random variables, then

$$\rho(\lambda X + (1 - \lambda)Y) \leq \lambda\rho(X) + (1 - \lambda)\rho(Y), \quad \text{for any } \lambda \in [0, 1].$$

Intuitively, this states that the excess system risk of a convex combination of two sub-grids cannot be more than the same convex combination of excess system risk of the excess system costs of the sub-grids.

4. **Homogeneity:** For any $\alpha > 0$, $\rho(\alpha X) = \alpha\rho(X)$. This states that the excess system risk scales linearly .

We note that as we are quantifying the risk of extreme excess costs, the “bad” cases are in the upper tail. In the context of financial profits and losses, the “bad” cases (losses) are in the left tail, where X would be negative, so formulas in properties 1 and 2 above are slightly different in that literature.

The features discussed above for coherence are all logical and desirable, and the unit of system risk is dollars. There are many such risk measures. However, one of the most common of these measures is Conditional Value at Risk (CVaR). CVaR also goes by many other names, such as Expected Shortfall, or Expected Tail Loss. It is defined as follows. For $\alpha \in (0, 1)$:

$$\text{CVaR}_\alpha(X) = \frac{1}{\alpha} \int_{1-\alpha}^1 \text{VaR}_q(X) dq, \quad (2.2)$$

where value at risk $\text{VaR}_q(X) = \inf\{x : F_X(x) \geq q\}$. In essence, CVaR captures the expected downside in the worst $\alpha \times 100\%$ of scenarios, and an equivalent formulation to (2.2) is

$$\text{CVaR}_\alpha(X) = \mathbb{E}\{X \mid X \geq \text{VaR}_\alpha(X)\}.$$

CVaR is of course, not the only possible valid choice of risk measure, and is also not perfect. All risk measures have strengths and weaknesses, and as discussed in [18], the appropriate course of action to decide the optimal risk measure will depend on the particular situation and risks being measured (and in this case, allocated). However, what makes CVaR particularly attractive in this situation is that it explicitly targets the extreme upside system costs that exist in a distribution.

As it relates to system costs, it seems plausible that the distribution of costs would be skewed heavily, with large tails representing potentially very costly days which can occur when there is extreme undersupply of electricity, and grids are forced either to turn to costly thermal generators to make up the difference at the last minute, or in a worst case scenario, shed load entirely. These extreme events corresponding to extremely large excess costs are the most relevant part of the distribution of grid costs, as they represent the most critical risk to the system as a whole, and therefore justify the choice of CVaR as the risk measure of choice. The choice of α may vary, but $\alpha = 5\%$ is a common initial choice that we will employ throughout.

2.4 Risk allocation methodology

Our goal is to allocate excess system risk which arises from stochasticity in renewables production, and allocate it to the renewable generators who are most contributory to this risk. As previously mentioned, this excess system risk is connected to the idea of counterfactuals between the grid as it is, and a world where renewables have perfect production reliability. In particular, the distribution across Monte Carlo scenarios of the cost differences between these worlds will be the genesis of excess system risk, and therefore the risk allocations and, later, reliability costs at the individual asset level. We primarily consider risks in an hourly context, noting that the ideas we discuss generalize to any time span.

The key idea behind the risk calculation and allocation methodology will be to idealize groups of renewable generators and observe how grid costs change, ultimately using the results across different combinations of idealized renewable generators to identify the assets which contribute to grid costs most significantly due to the stochasticity in their power output. In essence, we will adjust the simulations to modify specific groups of renewable assets into being idealized versions of themselves, which do not have the hour-to-hour randomness typically associated with renewables. Specifically, this means that they will produce exactly their day-ahead forecast on an hour-by-hour basis.

Armed with the excess system risk discussed in the previous subsection, our next goal is to allocate this risk in an additive manner to cohorts (groups) of generators, as well as power generators individually, using Shapley values. The cohorting may be by type of asset (*e.g.* solar, wind) or by location, and is used to reduce the dimension of the computation. The output of this process should be costs or proportions of the CVaR $\rho(X)$ that are associated with each renewable generator within the grid system whose system risk is being allocated. Suppose there are N renewable assets (or cohorts of renewables). We seek allocations $a_i, i = 1, 2, \dots, N$ to each such that

$$\rho(X) = a_1 + a_2 + \dots + a_N.$$

As with risk measures, which are chosen to satisfy some desired axiomatic properties, as described in Section 2.3, there is a literature on desirable axiomatic properties for risk allocation methods. These properties have names such as full domain, core compatibility, diversification, strong monotonicity, incentive compatibility, efficiency, equal treatment property, riskless portfolio, covariance, decomposition invariance, and we refer to [2] for details. There are also 8 risk allocation methods discussed there which each satisfy different subsets

Method / Property	FD	CC	Div	SM	IC	Eff	ETP	RP	Cov	DI
Activity based	∅	∅	∅	∅	∅	✓	✓	∅	∅	∅
Beta	∅	∅	∅	∅	∅	✓	∅	∅	∅	✓
Incremental	∅	∅	∅	∅	∅	✓	✓	∅	∅	∅
Cost gap	✓	∅	✓	∅	∅	✓	✓	✓	✓	∅
MRC (directional)	✓	∅	∅	∅	∅	∅	∅	✓	✓	✓
MRC (partial)	∅	✓	✓	∅	∅	✓	∅	✓	✓	✓
Shapley	✓	∅	✓	✓	✓	✓	✓	✓	✓	∅
Nucleolus	✓	✓	✓	∅	∅	✓	✓	✓	✓	∅

Figure 2: Properties satisfied by different risk allocation methods. Source [2].

of these 10 properties: see Figure 2. We choose to work with the Shapley values allocation, which satisfies 8 of the 10 properties.

Let G denote the set of renewable assets in the grid. The Shapley value of asset $i \in G$ is defined as follows. We pick a subset of renewable assets $H \subseteq G \setminus \{i\}$ not containing asset i and run the Monte Carlo simulations through the grid software with the assets in H idealized (their actuals adjusted to meet their forecasts). From this, we obtain the distribution of system costs $\nu(H)$. The Shapley value attributable to asset $i \in G$ is:

$$\phi_i = \frac{1}{N} \sum_{H \subseteq G \setminus \{i\}} \binom{N-1}{|H|}^{-1} (\nu(H) - \nu(H \cup \{i\})), \quad (2.3)$$

where $\nu(H \cup \{i\})$ is the system cost when the the asset i is also idealized, and $|H|$ is the size of the coalition H .

This formula may seem esoteric, but it has a natural interpretation. The excess cost attributable to generator i is the marginal cost added by that generator to a coalition of a given size, averaged across all the possible permutations in which the coalition of that size could be constructed, summed across all possible coalition sizes. We note the combinatorial difficulty that there are $2^N - 1$ terms in the sum, and we observe that

$$\sum_i \phi_i = \nu(\emptyset) - \nu(G) = X,$$

as in (2.1). We also note that there has been renewed interest in combining Shapley values with counterfactuals for interpretability of machine learning, explainable AI, and reinforcement learning. As we are not experts in this area, we defer providing references found from a search, but we mention the application to feature importance assessment in the prediction of electricity price spike prediction in [9].

2.4.1 Two Asset Example

Consider a hypothetical grid where there are only two renewable assets (or cohorts), denoted A and B , so $G = \{A, B\}$, and $\nu(\{A, B\})$ denotes BAU total system costs in a given hour.

Let $\nu(\mathbf{A}, \mathbf{B})$ represent PW total system costs in the same hour, where boldfaced asset name means that asset is idealized. It is clear from our previous definition (2.1) that as $X = \nu(\{A, B\}) - \nu(\{\mathbf{A}, \mathbf{B}\})$. We re-write X as:

$$\begin{aligned} X &= \nu(\{\cdot\}) - \nu(\{A, B\}) \\ &= \frac{1}{2} \left(\nu(\{\cdot\}) - \nu(\{A\}) \right) + \frac{1}{2} \left(\nu(\{B\}) - \nu(\{A, B\}) \right) \end{aligned} \quad (2.4)$$

$$+ \frac{1}{2} \left(\nu(\{\cdot\}) - \nu(\{B\}) \right) + \frac{1}{2} \left(\nu(\{A\}) - \nu(\{A, B\}) \right), \quad (2.5)$$

where $\nu(\{B\})$ is system cost when A is kept as is, but B is idealized, and similarly $\nu(\{A\})$. Note that (2.4) is the average across the two states of B (idealized or not) of the marginal value of idealizing A , and this is the Shapley value of A . Similarly, (2.5) is the average across both states of A of the marginal value of idealizing B , and this is the Shapley value of B . This motivates the Shapley decomposition $X = \phi_A + \phi_B$, with the definitions

$$\begin{aligned} \phi_A &:= \frac{1}{2} \left(\nu(\{\cdot\}) - \nu(\{A\}) \right) + \frac{1}{2} \left(\nu(\{B\}) - \nu(\{A, B\}) \right) \\ \phi_B &:= \frac{1}{2} \left(\nu(\{\cdot\}) - \nu(\{B\}) \right) + \frac{1}{2} \left(\nu(\{A\}) - \nu(\{A, B\}) \right), \end{aligned}$$

which corresponds to (2.3) with $N = 2$.

2.4.2 Risk Allocation of Decomposed Excess Costs

Given the Shapley decomposition of excess costs (2.3), we define the fair risk allocation for renewable asset i as

$$a_i = \mathbb{E} \{ \phi_i | X > \text{VaR}_\alpha(X) \}, \quad (2.6)$$

as discussed in [3]. This corresponds to the well-known Euler risk allocations (see [21]). The formula (2.6) effectively takes the average of the contributions of asset i to excess costs in the scenarios where excess costs are particularly large. It is clear from (2.6) that

$$\sum_{i=1}^N a_i = \mathbb{E} \left\{ \sum_{i=1}^N \phi_i | X > \text{VaR}_\alpha(X) \right\} = \mathbb{E} \{ X | X > \text{VaR}_\alpha(X) \} = \rho(X),$$

as desired. Therefore, in obtaining and calculating (2.6), we have established a sound methodology for allocating the overall risk of excess system costs due to stochastic renewable generation to the individual generators responsible.

3 Computational Results

The decomposition of X via Shapley value requires running `Vatic` for each combination of assets being idealized. This results in 2^N combinations, each of which must be run with hundreds of Monte Carlo scenarios. To reduce the computational burden, we will use Shapley decompositions for cohorts of assets, rather than each individual asset. We can then allocate risk to this smaller number of cohorts, and use a first-order Shapley approximation to

subdivide cohort level risk allocations to the individual members of the cohort. Cohorts can be chosen in a myriad of ways, such as separating renewable assets by geographic area, or by type. Ideally, we should choose a small number of cohorts, as there are 2^C grid specifications, where C is the number of cohorts.

Once the cohorts are chosen, we follow the methodology described in the previous section, treating assets within the same cohorts similarly. In Figure 3, we show the different idealization groupings for a hypothetical example where renewables are broken up into three groups, based on asset type (RTPV, PV, Wind), where PV is solar and RTPV is rooftop solar. As




Figure 3: Schematic of cohort idealization and coalitions of sizes 0,1,2,3.

Figure 3 suggests, this means that entire cohorts are idealized together. The computational costs are on the order of a few hours with this reduction.

3.1 RTS Grid

We first discuss the methodology applied to the RTS-GMLC test grid, which covers southern California and Nevada with 73 load buses, 3 geographic zones, and 60 renewable assets (25 PV, 31 RTPV, 4 Wind). Applying the methodology to RTS allowed for greater experimentation and testing as RUC / SCED instances on RTS are significantly faster than those run on Texas7K (roughly 1 minute for a RUC/SCED instance per scenario). This is due to the small size of RTS in comparison to Texas7K.

3.1.1 Cohort Risk Allocations

We provide here the results for the cohort allocations on the RTS grid for two days within 2020 (June 30 and January 20). The former is pictured on the left of Figure 4 and the latter is on the right of the same figure. The top panel in each column shows the hourly risk allocations (reliability costs to the grid) by generator type. The panels below that show the forecasts and PGScen scenarios for load, wind, PV, and RTPV assets, with the dots indicating the particular scenarios which correspond to the most extreme values of X (excess costs).

Focusing on the left panel first (June 30), we see that PV and RTPV contribute to stochasticity risk only during sunlight hours, as expected. PVs and RTPVs also appear to have roughly the same overall profile and effect on this day. Meanwhile, we see that wind generators are allocated a large risk in the later stages of the day, due to the likelihood of

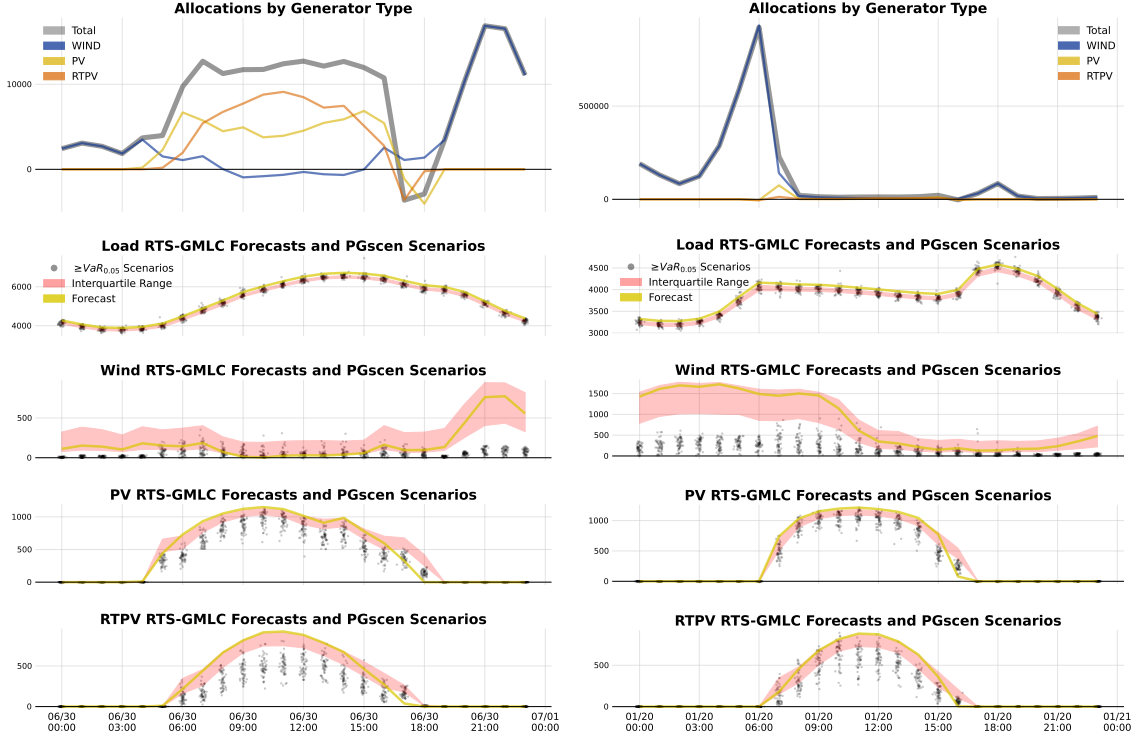


Figure 4: Cohort Allocation for June 30 2020 (left) and January 20 2020 (right) on RTS grid with load and renewable generation scenarios

generating far lower than the forecast amount (which can be seen in the middle plot of the left panel).

Turning our attention to the right panel (January 20) of Figure 4, we see the risks wind assets pose due to stochasticity in the morning dominates the day, to the point of other asset types appearing relatively inconsequential. Once again, the wind allocation is due to the number of scenarios in which wind assets produce significantly less than forecasted, and are compounded by the fact that morning hours see increasing load, without a PV or RTPV buffer to help account for that. As a result, situations where wind assets fail to deliver as expected are hugely costly for the grid in those hours, as they force the ISO to call upon fast-starting thermal generators, which are expensive.

3.1.2 Asset Risk Allocations: first order Shapley approximation

Next, we must distribute these cohort-wide allocations within the cohort, to its individual renewable assets. In principle, there are many ways to perform such a distribution, such as an equal distribution, distributing proportional to size, or to some long-run estimate of reliability (such as the average percentage of their forecast each asset delivers on a rolling basis). However, we will use a first-order approximation to the Shapley decomposition.

True Shapley decompositions require calculating differences in system costs for every idealizing grouping in the power set of generators. Instead, we will only calculate these differences by comparing the grid in the BAU sense to the grid with just one asset idealized. In other words, we compute (2.3) using only subsets H of size one. This results in a first-

order estimate of each asset’s Shapley decomposition. For all assets in a given cohort, we then sum these estimates, and find the relative weight of each asset. Finally, we apply this to the cohort-wide risk allocation in order to get an individual allocation. Once again, the sum of the individual allocations will be exactly the CVaR of X .

To conclude, we show the results of this methodology for June 30 and January 20, 2020 in Figure 5 and Figure 6 respectively. From each of these figures, we can see the same general shapes as the cohort allocations. However, it is important to note that we are able to pick up individual differences in assets, despite their closeness in proximity and type. As an example, we can see marked differences in wind assets in Figure 5 across different hours. This suggests our methodology has the fidelity to appropriately allocate system risks that arise due to stochasticity to the generators who most contribute to it, and to distinguish between superficially similar assets, both of which are important traits of the methodology.

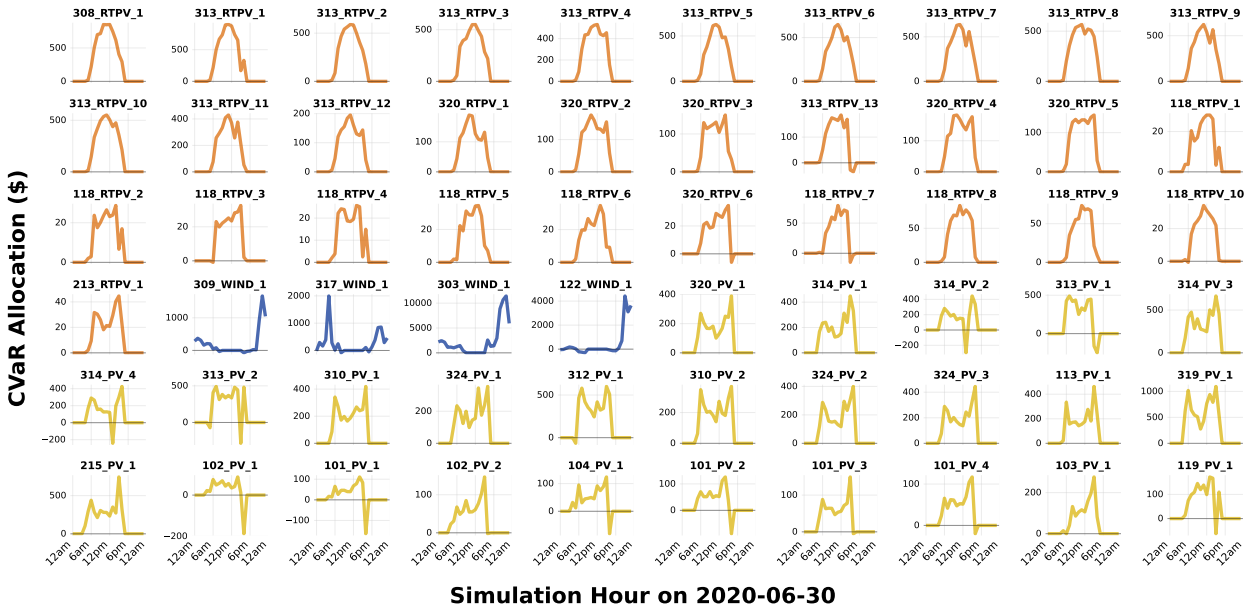


Figure 5: Hourly asset allocations for June 30 2020 on RTS grid

3.2 Texas Grids

We work with two grid models of Texas constructed by the group at Texas A&M. The Texas 2020 (current) grid has 673 total generators, of which 36 are solar and 149 wind. The 2030 (projection) grid has 977 total generators, of which there are 178 solar and 293 wind. Details are in Figure 7. For feasibility of the Shapley computation, we first show results from cohorting the renewables in their two types, wind and solar. Then we compute individual asset allocations following the procedure described in Section 3.1.2.

3.2.1 Texas 2020 & 2030 grid cohort allocations

In Figure 8, we show wind and solar (PV) allocations per MWh of production for a summer and a winter day in the Texas 2020 (current) grid. In both cases, solar risk outweighs wind




Figure 6: Hourly asset allocations for January 20 2020 on RTS grid

Category	Texas7K (Current)	Texas7K (2030 Projection)
Overall Generator Capacity (MW)	102,887 MW	166,912 MW
Solar Capacity (MW & %)	2,335 MW (2.27%)	26,835 MW (16.1%)
Wind Capacity (MW & %)	25,131 MW (24.4%)	55,702 MW (33.4%)

Figure 7: Texas grids 2020 & 2030

risk, peaking around sunset, with a smaller peak around sunrise. The lower graphs show load, wind and solar forecasts and Monte Carlo scenarios. In Figure 9, we show the allocations for

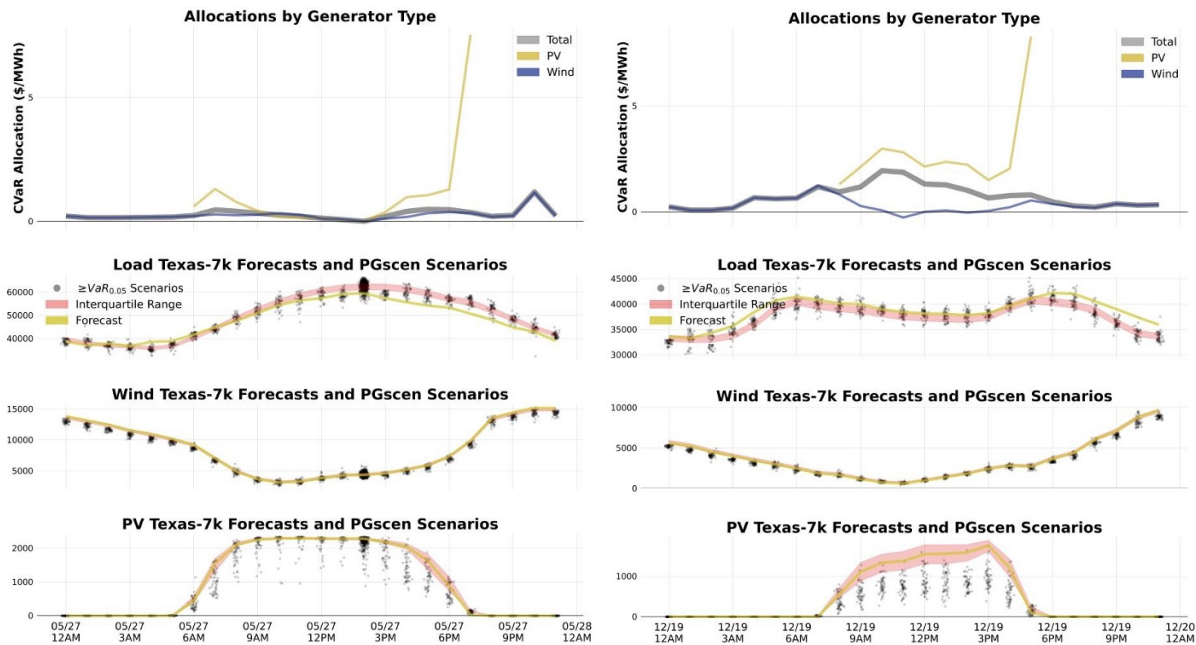


Figure 8: Cohort Allocations for May 27 (left) and December 18 (right) on Texas 2020 grid

the Texas 2030 futuristic grid on a day in September. Again, solar dominates, but increased renewable penetration (up by a factor of 12 in capacity from 2020) means the risk allocation is large for a good portion of the daylight hours.

3.2.2 Individual asset allocations for Texas 2030 grid

In Figure 10, we show individual asset risk allocations on a heat map of Texas for six hours on September 4, for the 2030 grid. Red denotes high risk allocations, while blue are negative (beneficial allocations) meaning those assets at that time are positively helping the grid’s reliability (for instance by exceeding forecasts and making up for other assets that are below forecasts). Wind assets are diamond-shaped and solar assets are circles. The size is proportional to their production at each hour, and allocations are scaled to dollars per MWh of production.

We see that:

- 12am At night all risk is from wind, obviously, and production is relatively low, with greater risk from the assets in the west, and less so from the south of the state.
- 7am By morning, solar risk emerges in far west Texas, and there is still some wind risk.
- 3pm By the afternoon, solar risk dominates. The method identifies less reliable solar in the west, and some more reliable solar in the east.
- 5pm As the sun drops, we see a mix again of some wind risk and the last of the solar risk.
- 9pm At night, wind risk in the west dominates that from the south.

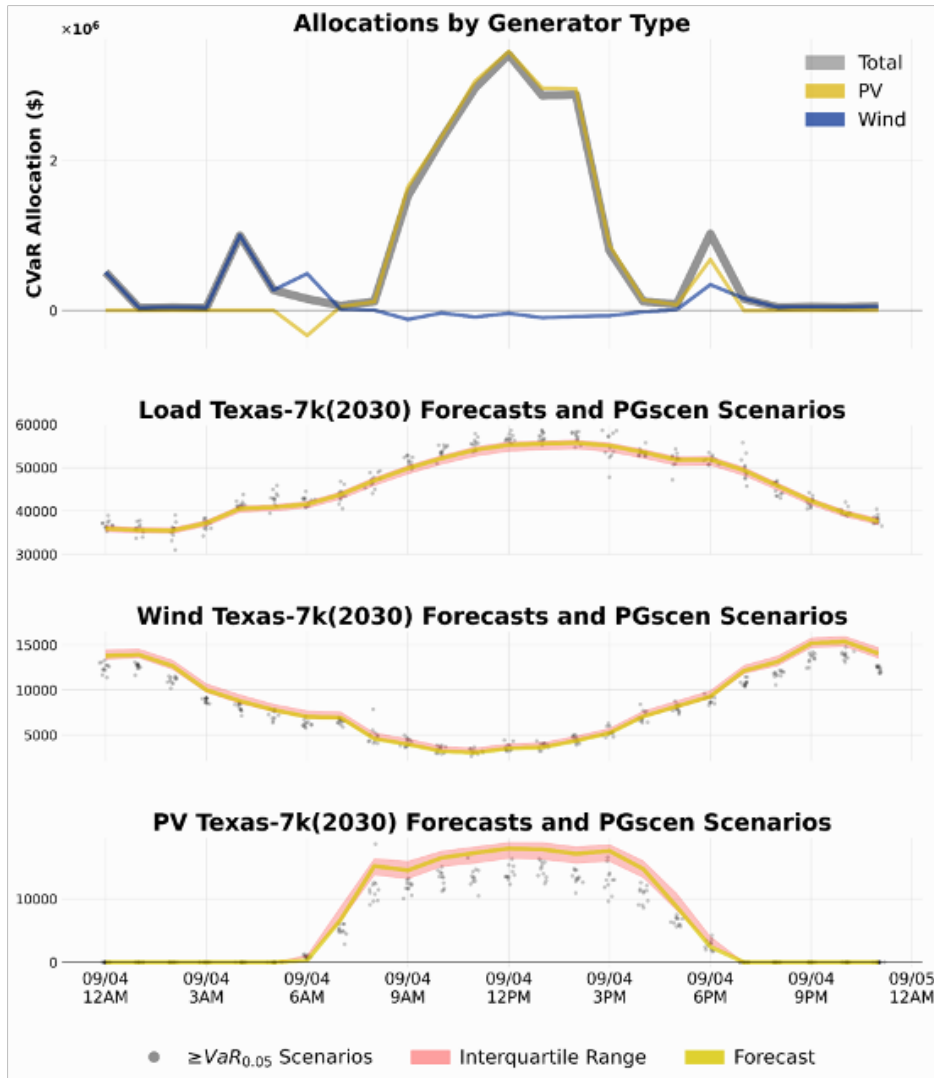


Figure 9: Cohort Allocations for Sept 4 on Texas 2030 grid

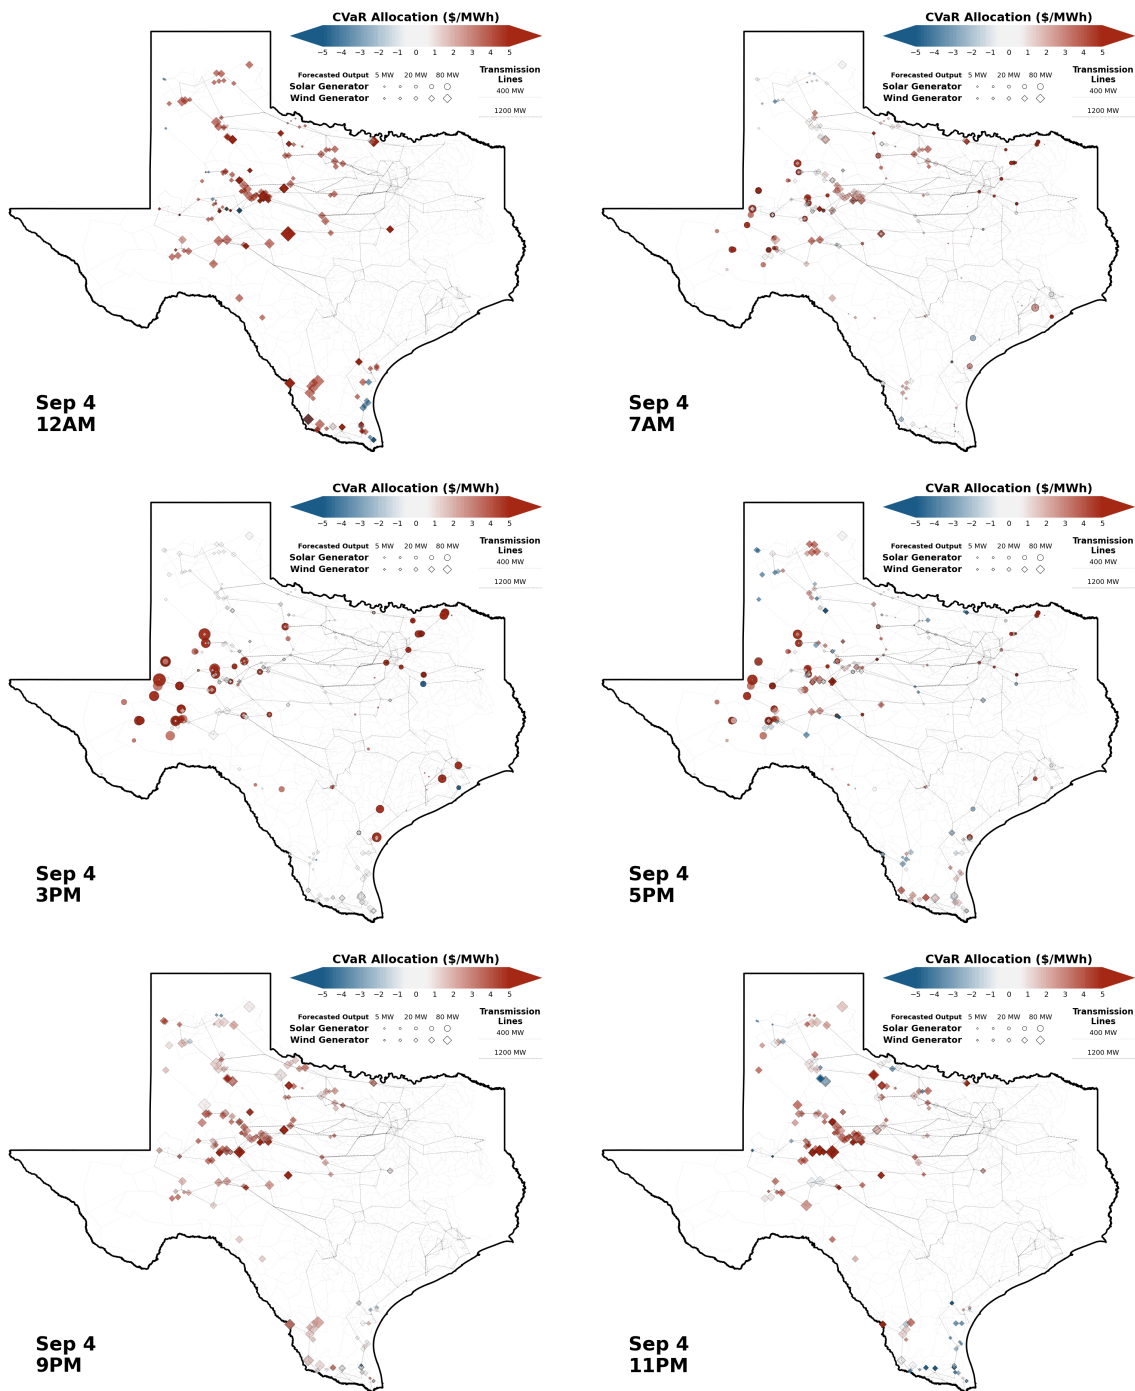


Figure 10: Asset allocations for September 4 on Texas 2030 grid

11pm By the end of the day, risk is concentrated in the west and the southern wind stations provide reliability.

4 Reliability Cost Curves

One way to use the risk allocations is to impose *reliability cost curves* on renewables to account for their stochasticity, and re-do the (deterministic) optimization black box. As a result, unreliable assets move up the bid stack⁴ since they no longer bid at zero marginal cost. The grid software learns to pre-commit more traditional generator backups, reducing tail risk. The goal is to highlight the need for more renewables (and where), not penalize them out. We mention also virtual bidding, that is financial trading in the day ahead and real time electricity markets, which has also been shown to improve commitment decisions under uncertainty: see, for instance, [11].

4.1 Constructing renewables reliability cost curves

The level of cost curve for each renewable generator each hour is scaled by their risk allocation a_i , so higher risk, higher costs curves. Moreover, the probabilistic distribution of power production for each asset is known (numerically) from our Monte Carlo simulations. The “integral” (cumulated sum) of this histogram is the empirical CDF of production from that asset $F_i(p)$, that is, the probability of producing less than p MWh. An example for a wind asset is in Figure 11. This informs the shape of the reliability cost curves that we construct.

Ultimately, the procedure for defining the reliability cost curve $c_i(p)$ for renewable asset i

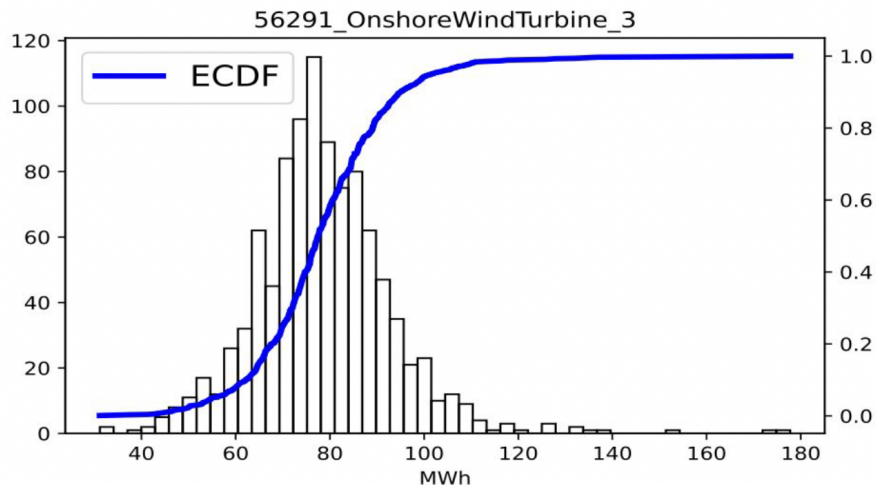
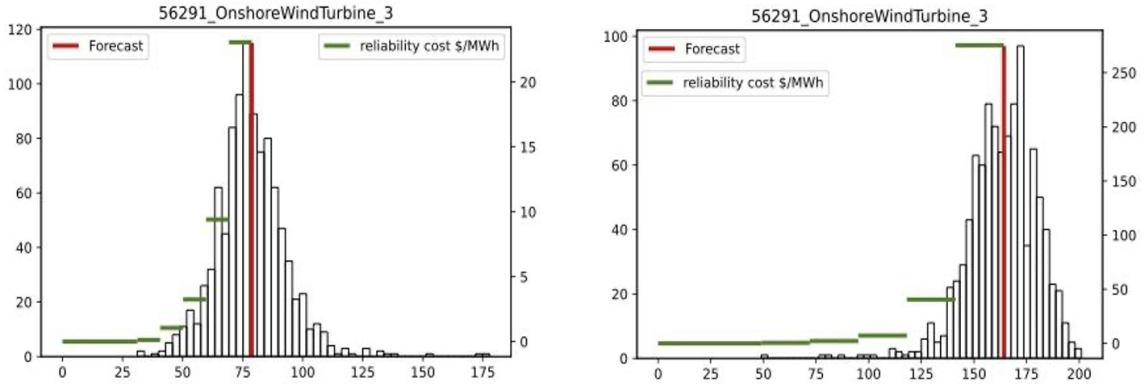


Figure 11: Monte Carlo CDF $F_i(p)$ of a wind asset (units on right axis). The histogram is of production from the wind asset.

⁴We refer to [7] for details on bid stacks.



Left: cost curve at 2018-09-04, 22hr. Total risk allocation is \$338.0.

Right: cost curve at 2018-02-21, 17hr. Total risk allocation is \$6933.1.

Figure 12: Reliability cost curves for wind asset.

with risk allocation a_i (at a given hour) is to take the ECDF F_i up to the asset's forecast for that hour and

$$c_i(p) = \beta \times a_i \times F_i(p) \quad p \leq \text{forecast}, \quad a_i > 0,$$

where $\beta > 0$ is a tuning parameter that converts to appropriate levels of cost, and which can also be interpreted as a universal unreliability-aversion parameter (across all assets and times). In the (rare) case of negative risk allocations, we keep the costs at zero.

Figure 12 shows the constructed cost curves (converted to standard step-bids format) for the wind asset at two different times of the day. The level of cost is on the right axis. On the left, hour 22, the risk allocation is low and the cost curve is just above 20 at the forecast. On the right, hour 17, the risk allocation is much higher and the cost curve peaks at just over 250.

4.2 Unit Commitment re-informed by reliability cost curves

Now that the risk allocations have been used to provide reliability cost curves for the renewable assets whose intermittency most impact extreme system costs, we can re-inform the day ahead unit commitment taking this refined probabilistic information into account. Then we can re-run the real time Monte Carlo simulations and quantify the new excess system risk (CVaR).

An example is shown in Figure 13 with the Texas 2030 grid, on a day in March (top graph). We see that the new unit commitment mitigates the spikes in excess risk at hour 6, 9, 13 and 23, as was intended. Another day in December is shown in the bottom graph of Figure 13, where smaller risk spikes at hours 0, 5, 10 are lessened as well as the bigger spike at hour 17.

It is natural to ask what is the new unit commitment informed by the reliability cost curves, compared with the old where renewables bid at zero marginal cost. This is illustrated in Figures 14 and 15, which show areas where there is an increase in thermal generator use




Figure 13: March 17 2030 (top) and December 5 2030 (bottom) system CVaRs with and without cost curves

in the new UC, and where there was loss of load under the old UC (without cost curves) but not under the new one (with cost curves). In Figure 14 (5am), we see loss of load (purple circles) in some parts of southeastern Texas under the UC where renewables reliability is not taken into account. With the cost-adjusted UC, there is no loss of load (yellow circles), but thermal power use is increased mostly in the areas of Dallas and Austin population centers (red circles). In the afternoon, say 4pm (Figure 15), the loss of load under the traditional UC is pronounced in both the eastern and central parts of the state. With the new UC, these no longer occur, and thermal usage is increased in numerous locations to mitigate renewables intermittency. Our method thus identifies times and places where supply reliability is a potential problem in a future grid.




Figure 14: Re-distribution, December 5 2030 at 5am.




Figure 15: Re-distribution, December 5 2030 at 4pm.

5 Summary & Conclusion

Modern electricity markets face new sources of randomness due to a growing adoption of renewable sources, such as wind and solar power, with near-zero marginal costs of production. As the number of such assets increases, renewables intermittency poses a significant risk to grid stability and costs. We have introduced a method for reliability quantification: ascribing risk and costs to each renewable asset’s contribution to system operational cost. This is achieved by capturing the variability in renewable generation through Monte Carlo scenarios, which are input into an optimal power flow (OPF) model to yield a probabilistic distribution of the excess system-wide operational cost. The extreme high-cost tail is quantified by a monetary system risk measure, which is then allocated back to each asset according to its contribution to system risk. This allocation is based on the concept of capital allocation used in the insurance and finance sectors. The ascribed reliability costs are then utilized to re-order the merit order for ISO optimization by adjusting the bid-in generation costs to reflect reliability risk contribution. This methodology allows the uncertainty of renewables production to be taken into account to better mitigate extreme risks, but without altering the OPF (“black box”) software currently used by grid operators.

A natural future direction in progress is to quantify carbon emissions risk, as covering renewable shortfalls with quick startup thermal generators is costly both in dollar terms and emissions of carbon. Both the current work on costs and forthcoming work on emissions allow grid planners to account for intermittency risks with precise quantitative temporal and spatial information.

Acknowledgements

The authors were partially supported by ARPA-E grants DE-AR0001289 and DE-AR0001390 under the PERFORM program of ARPA-E under the US Department of Energy. The authors thank René Carmona, Michael Ludkovski and Glen Swindle for helpful discussions; Junying Fang, Michal Grzadkowski and Amit Solomon for software development and implementation; and Felix Ackon, Ethan Reese and Dominic Riendeau-Krause for research assistance.

References

- [1] Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [2] Dóra Balog, Tamás Bátyi, Péter Csóka, and Miklós Pintér. Properties and comparison of risk capital allocation methods. *European Journal of Operational Research*, 259(2):614–625, 2017.
- [3] Tomasz R Bielecki, Igor Cialenco, Marcin Pitera, and Thorsten Schmidt. Fair estimation of capital risk allocation. *Statistics & Risk Modeling*, 37(1-2):1–24, 2020.
- [4] Adam Birchfield. Texas A&M PERFORM cases. <https://electricgrids.engr.tamu.edu/texas-am-perform-cases/>.

- [5] Mary B Cain, Richard P O’Neill, and Anya Castillo. History of optimal power flow and formulations. 2012.
- [6] 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.
- [7] Michael Coulon and Sam Howison. Stochastic behaviour of the electricity bid stack: From fundamental drivers to power prices. *he Journal of Energy Markets*, 2:29–69, 2009.
- [8] Michel Denault. Coherent allocation of risk capital. *Journal of Risk*, 4:1–34, 2001.
- [9] Rémi Galarneau-Vincent, Geneviève Gauthier, and Frédéric Godin. Foreseeing the worst: Forecasting electricity DART spikes. *Energy Economics*, 119:–, March 2023.
- [10] Michal Grzadkowski, Carmona Rene, and Xinshuo Yang. Pgscen. <https://github.com/PrincetonUniversity/PGscen/>, 2023.
- [11] Akshaya Jha and Frank A. Wolak. Can forward commodity markets improve spot market performance? evidence from wholesale electricity. *American Economic Journal: Economic Policy*, 15(2):292–330, May 2023.
- [12] Max Junqueira, Luiz Carlos Da Costa Júnior, L.A. Barroso, G.C. Oliveira, Luiz Thome, and Mario Pereira. An aumann-shapley approach to allocate transmission service cost among network users in electricity markets. *Power Systems, IEEE Transactions on*, 22:1532 – 1546, 12 2007.
- [13] Michael Kaltenbrener. An axiomatic approach to capital allocation. *Mathematical Finance*, 15(3):425–437, 2005.
- [14] 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.
- [15] Jacob Mays. Cost allocation and net load variability. *IEEE Transactions on Power Systems*, 33(2):2030–2039, 2017.
- [16] New York Independent System Operator. Reliability and market considerations for a grid in transition. Technical report, NYISO, 2019.
- [17] Anthony Papavasiliou, Shmuel S. Oren, and Richard P. O’Neill. Reserve requirements for wind power integration: A scenario-based stochastic programming framework. *IEEE Transactions on Power Systems*, 26(4):2197–2206, 2011.
- [18] Svetlozar Rachev, Sergio Ortobelli Lozza, Stoyan Stoyanov, and Frank J. Fabozzi. Desirable properties of an ideal risk measure in portfolio theory. *International Journal of Theoretical and Applied Finance*, 11:19–54, 2008.

- [19] Brian Sergi, Cong Feng, Flora Zhang, Bri-Mathias Hodge, Ross Ring-Jarvi, Richard Bryce, Kate Doubleday, Megan Rose, Grant Buster, and Michael Rossol. Arpa-e perform datasets. Technical report, DOE Open Energy Data Initiative (OEDI); National Renewable Energy Laboratory, 2022.
- [20] L. S. Shapley. A value for n-person games. In Harold William Kuhn and Albert William Tucker, editors, *Contributions to the Theory of Games (AM-28), Volume II*, pages 307–318. Princeton University Press, 1953.
- [21] Dirk Tasche. Capital allocation to business units and sub-portfolios: the Euler principle. *arXiv preprint arXiv:0708.2542*, 2007.