Managing Financial Risks for Wind Power Producers in Wholesale Electricity Markets

by

Daniel Weihang Shen

B.S. Mechanical Engineering, University of California, Berkeley, 2020

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2024

Authored by: Daniel Weihang Shen
Department of Electrical Engineering and Computer Science
January 26, 2024

Certified by: Marija D. Ilic
Adjunct Professor of Electrical Engineering and Computer Science and Senior Research Scientist in the Laboratory for Information and Decision Systems, Thesis Supervisor

Accepted by: Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Managing Financial Risks for Wind Power Producers in Wholesale Electricity Markets

by

Daniel Weihang Shen

Submitted to the Department of Electrical Engineering and Computer Science
on January 26, 2024 in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE

ABSTRACT

Wind power plant operators are exposed to financial risk in wholesale electricity markets due to the uncertain nature of wind forecasts, day-ahead electricity prices, and real-time electricity prices. In the event of a shortfall compared to the production forecast, the wind generator may have to repurchase power at a higher price in the real-time market. Based on this consideration, this thesis formulates a mixed-integer quadratic program which uses conditional value at risk to create a hedged "risk-aware" offer curve for the wind generator to submit into the day-ahead electricity market. The formulated program additionally considers specific concerns around the offer optimization process being negatively interpreted as using physical withholding to increase profits. We also exploit the structure of the problem to introduce additional constraints to improve computation time and demonstrate that despite the complexity of mixed-integer variables it can be solved within an acceptable operating timeframe under realistic conditions. We simulate the impacts on financial returns for the generator of applying such an approach for a wind farm in the New York City region; the program can be successfully tuned to adjust the variability in returns based on the agent's preferences, but does not outperform a more naive strategy of simply cutting off the quantity based on a percentile of the forecast distribution. Finally, we provide some discussion on how the act of "active" price creation through these risk-aware offer curves could come into conflict with the current regulatory environment, especially around the concept of exercise of market power, which has long relied on tying fair prices to ones that represent marginal fuel costs for generators.

Thesis supervisor: Marija D. Ilic
Title: Adjunct Professor of Electrical Engineering and Computer Science and Senior Research Scientist in the Laboratory for Information and Decision Systems

Acknowledgments

This thesis, like the coordinated operation of the electric grid, was only possible because of the work and support of many people coming together.

I thank my advisor, Marija Ilic, for her support of my ideas in carrying out this work. It is a treasure to be able to work with somebody who has seen the monumental changes in this industry, both past and present, and is able to help contextualize the importance of grappling with these ideas. Even under the many pressures of the university environment she has been generous with her time and serves as a personal inspiration with her dogged work ethic.

I would also like to thank my parents and brother for their unconditional love, support, and encouragement in my graduate school journey while on the other side of the country.

Finally, I would like to acknowledge the role of ARPA-E's PERFORM program and funding from the National Science Foundation's Graduate Research Fellowship Program in making this work possible through funding and vigorous discussion in the context of industry stakeholders.

Contents

Title page 1

Abstract 3

Acknowledgments 5

List of Figures 9

List of Tables 11

Nomenclature 13

1 Introduction 15
- 1.1 Restructuring: From Vertical Integration to Wholesale Markets . . . 16
  - 1.1.1 The Two-Settlement System . . . 18
- 1.2 Decarbonization and Variable Renewable Energy . . . 19
  - 1.2.1 What's in a Price? . . . 21
- 1.3 Managing Uncertainties . . . 21

2 Problem Framework 23
- 2.1 Optimization Problem Formulation . . . 23
  - 2.1.1 Related Work . . . 24
  - 2.1.2 Risk Measure: Conditional Value at Risk . . . 25
  - 2.1.3 Offer Curve . . . 26

3 Results 35
- 3.1 Curve Interpretation . . . 35
- 3.2 NYISO Performance . . . 36
  - 3.2.1 Data Inputs . . . 37
  - 3.2.2 Regret Comparison . . . 37
- 3.3 Runtime Performance . . . 40

4 Conclusion 43

A Code 45

References 49

List of Figures

1.1 Map of US ISOs/RTOs and Wholesale Markets 17
1.2 Timeline of two-settlement market events in NYISO 19
1.3 Load Forecast Performance, NYISO 20

2.1 Graphical representation of conditional value at risk 25
2.2 Example energy offer curve based on CAISO example 28
2.3 Representation of offer curve optimization variables, row-echelon form 29
2.4 Representation of offer curve optimization variables, special ordered sets form 33

3.1 Offer curves and pairwise visualizations of distributions. 38
3.2 Strategy performance on NYISO data 39
3.3 CVaR MIQP formulation runtime statistics 41

List of Tables

2.1 Example energy bid, CAISO 28

3.1 Simulation parameters for synthetic Gaussian data cases 37
3.2 Runtime performance for CVaR MIQP formulation 40

Nomenclature

Sets and Indices

Ω	Set of price and wind forecasts.
ω	Index over forecast scenarios.
t	Time indexer.
i	Offer curve segment indexer.
N_{seg}	Total number of offer curve segments.
N_{Ω}	Total number of scenarios considered.

Decision Variables

p_{offer}^{ti}	Power offered in segment i (MW).
λ_{offer}^{ti}	Price of the power offered in segment i (\$/MWh).

Auxiliary Variables

$p_{CL}^{t\omega}$	Power cleared in day-ahead under scenario ω.
$\Delta^{t\omega}$	Shortfall between the day-ahead clearing amount and the real-time forecast realization; negative if there is a repurchase in real-time.
$u^{t\omega i}$	Indicator if segment i clears the market.

Random Variables

$\lambda_{DA}^{t\omega}$, $\lambda_{RT}^{t\omega}$	Day-ahead and real-time prices in scenario ω.
$p_{max}^{t\omega}$	Maximum dispatchable wind power in scenario ω.

Parameters

β	Risk factor for conditional value at risk.
M	Arbitrarily large constant for indicator constraints.
ϵ	Arbitrarily small constant for indicator constraints.

Chapter 1

Introduction

The North American power grid is the largest and most complex machine in the world and is integral to almost every aspect of society. The production and consumption of electricity across much of the continent occurs in the context of wholesale electricity markets where buyers and sellers make financial offers and bids for electricity. Unlike other commodities, however, electricity cannot be practically stored at scale and must be consumed as soon as it is generated. To accomodate this degree of coordination required to ensure reliability, most wholesale electricity markets in the United States are administered by a system operator which runs the market by setting prices, facilitating financial settlement and transfers, and choosing where and how generating resources should be run to balance the system, subject to least-cost and reliability objectives. Grid operations and electricity markets have re-entered the public and research spotlight as entities around the world set decarbonization targets in the electricity sector for climate goals, as well as due to exogenous shocks such as Winter Storm Uri in 2021[1] and increased natural gas prices due to the Russian invasion of Ukraine in February 2022.[2] In this context, we provide an optimization procedure for how individual generating agents on the grid can make offers into the day-ahead market to manage their short-term financial uncertainties, as well as some discussion of how this procedure, and bidding behavior in general, interact with the current regulatory environment of electricity markets.

---

1 Although Uri affected grid operations across the United States, its effects were most notable in Texas, where more than 4.5 million homes and business lost power, some for several days. Compounding the crisis was the additional lack of gas and water delivery due to downstream infrastructure effects.

2 The impact of gas prices stemming from the 2022 invasion on electricity prices motivated a session at the 2023 IEEE Power & Energy Society titled "Is Marginal Cost Pricing still the ideal electricity market design?". Panel participants strongly agreed that marginal pricing should still be used, although there was much more mixed discussion on how short-term marginal prices influenced longer-term investment.

1.1 Restructuring: From Vertical Integration to Wholesale Markets

For the majority of the 20th century, most electricity service was provided by vertically-integrated utility monopolies which owned both transmission and generating assets. The costs of construction and operation, along with a premium for investors, was recovered through retail prices approved through a public regulatory body. Beginning in the 1990s, various countries, including the United States, began a process of transitioning electricity production away from a vertically integrated model. This process of "restructuring"[3] was driven by consumer dissatisfaction with increases in retail prices needed to cover high fuel costs and unexpectedly stranded generation assets [1], [2]. The deregulation process of replacing rate-of-return pricing with wholesale competition was i) nominally supported by the hope that future market competition would bring down prices and spur innovation and ii) materially supported by the fact that ratepayers would be able to shift the responsibility for recovering the sunk costs of stranded[4] assets to stockholders.

One of the largest impacts of restructuring was the creation of larger regional wholesale markets for electricity (Fig. 1.1). In several regions, these wholesale markets are managed by an Independent System Operator (ISO) or a Regional Transmission Organization (RTO).[5] These system operators are tasked with maintaining short-term reliability, facilitating competitive electricity markets, and monitoring for exercise of market power. [6] To address the first two tasks, system operators procure energy and ancillary service products through a unit commitment (UC) process. Although the specifics of the market design vary across

[3] The terms "restructuring" and "deregulation" are both used to refer to the changes in the industry in the 90s, and 2000s. It is most correct to only use deregulation to refer to the changes to generation service.

[4] Borenstein and Bushnell note that while many of the nuclear and coal plants were considered "stranded" in the 1990s, by 2007 they were once again profitable due to later increases in natural gas prices [2].

[5] For the purpose of this work, we will refer to a generic "system operator", since both ISOs and RTOs manage markets for economic dispatch.

[6] Market power is "the ability to alter profitably prices away from competitive levels"[3]. See [4] for a description of practical tests for market power.

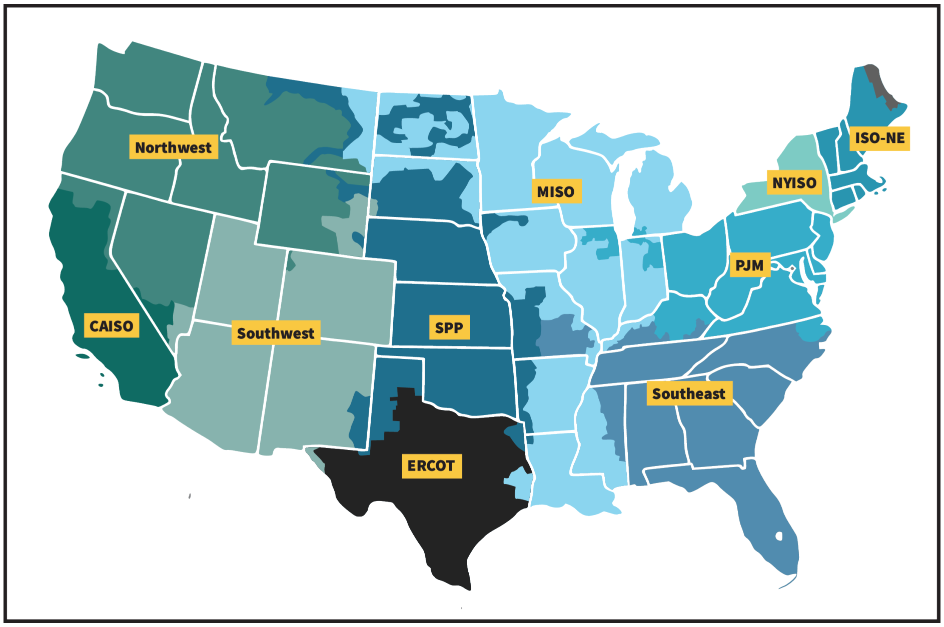

Figure 1.1. A US map of the ISOs/RTOs (CAISO, SPP, ERCOT, MISO, PJM, NYISO, and ISO-NE) and regions where traditional wholesale markets still operate (Northwest, Southwest, and Southeast).
Reproduced from FERC.

system operators, the idealized day-ahead electricity market broadly runs as follows:

1. Generators submit hourly offers specifying their supply functions of price vs. energy produced and loads submit bids specifying their demand functions of price vs. energy consumed. On a longer recurrance (e.g. weeks), generators also specify constraints such as ramp limits from hour-to-hour, minimum run times, and startup costs.

2. The system operator takes the bids and offers and runs a UC engine to determine which generators will start up and what levels they should produce at to minimize production costs or maximize total welfare. The UC problem takes into account transmission constraints and aforementioned generator constraints when producing the dispatch schedule.

3. The system operator issues a set of hourly dispatch levels to generators and a load-consumption "schedule" based on the UC results.

4. In addition to the production and consumption levels, the system operator also generates a set of locational marginal prices against which suppliers will be paid or consumers will be charged for the energy transacted in the day-ahead market.

1.1.1 The Two-Settlement System

The system operator runs the electricity market as a two-settlement system, with a day-ahead and real-time market.[7] Fig. 1.2 gives an overview of the events and the times they occur at for the New York market as a representative example. The day-ahead market is a forward market that is settled before the operating day; the real-time market occurs much closer to the actual operating hour. The two-settlement system is intended to help coordinate an optimal combination of resources in advance of the real-time market, since some generators may have startup times well exceeding the real-time market timeframe of a few hours.

Typically, the day-ahead market operates on a once-a-day cycle and clears all twenty-four hours of the operating day in hourly intervals, and the real-time market runs on an hourly cycle with fifteen minute intervals. In each interval, the system operator calculates a dispatch schedule for each generator (when to turn on/off and how much power will be produced in each interval), as well as a price for the contracted power. Financial payment to generators is based on the quantity contracted in the day-ahead market and the difference in quantity between the day-ahead and real-time contracts.

If a generator is paid λ for power p, then the total payment is:

$$\lambda_{RT}(p_{RT} - p_{DA}) + \lambda_{DA} p_{DA} \tag{1.1}$$

where λ and p represent the clearing price and quantity contracted, respectively.

Since the RT market is settled by considering the deviations in quantities cleared between the two markets, under the assumption of perfect competition, suppliers and consumers have the same performance incentives in both markets as if only the real-time market existed [3]. Put another way, the forward contract resulting from the day-ahead market should not have an impact on the generator's offer behavior in real-time, with the assumption that the RT market is competitive.

[7] Electricity is also traded in bilateral forward markets which are not managed by the system operator; these markets are outside the scope of this work's discussion.

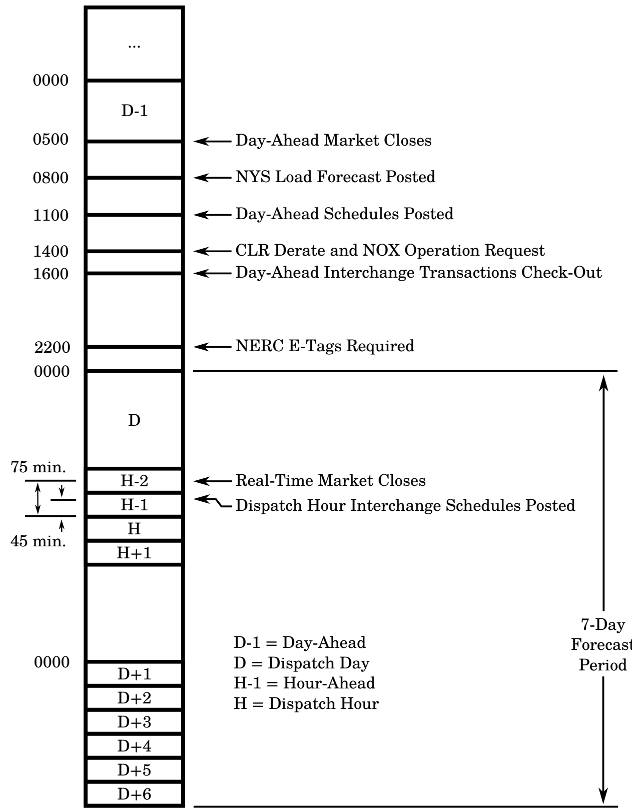

Figure 1.2. Timeline of two-settlement market events in NYISO.
Reproduced from the NYISO Day-Ahead Scheduling Manual [5].

1.2 Decarbonization and Variable Renewable Energy

Decarbonization to meet climate objectives will require a massive buildout of new generating resources, both to replace carbon-emitting generation facilities and accomodate increased demand for electrification as sectors such as heating and transportation shift away from fossil fuels. Much of this carbon-free generation will be in the form of solar and wind resources. This variable renewable energy (VRE) production is tied to intermittent weather conditions; the maximum available output is dependent on weather, although within the weather-dependent production cap, the output may be to an extent curtailed or downwards-dispatched. This is in contrast to conventional generation resources, which are relatively more dispatchable within a known operating range. [8]

Handling forecast uncertainty is not a new problem to the field. System operators have been forecasting demand for decades and are able to consistently achieve forecasts within 3%

[8] There is an often-repeated view is that VRE is not controllable whereas other generators are controllable, but the reality is that certain fossil-based technologies such as combined-cycle plants and block-loaded gas turbines may have nontrivial operating constraints that limit their output within restrictive ranges of nameplate capacity.

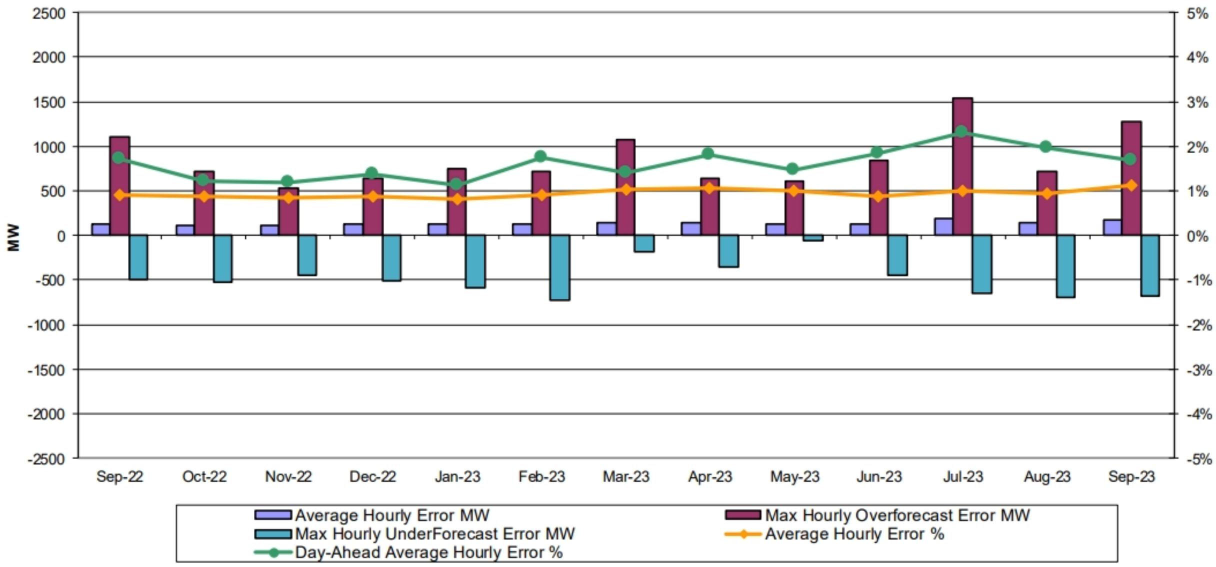

Hourly Error MW: Value of the difference between the hourly average actual load demand and the average hour ahead forecast load demand

Average Hourly Error %: Average value of the ratio of hourly average error magnitude to hourly average actual load demand.

Day-Ahead Average Hourly Error %: Average across all hours of the month of the absolute value of the difference between actual load demand and the Day-Ahead forecast load demand, divided by the actual load demand.

Figure 1.3. Load Forecast Performance, NYISO [8].

of actual demand (Fig. 1.3). This uncertainty has historically been operationally managed by operating reserves [6]. In the case of large imbalances between day-ahead and real-time markets, extra capacity may be redispatched in the real-time market. The increasing addition of VRE resources to the grid has also led system operators to account for VRE production forecasts in grid operations and markets.

The uncertainty of VRE forecasts does not contradict the basic features of restructured, competitive wholesale markets — uncertainty management (albeit moreso on managing uncertainties in demand) has been part of the restructured industry since its conception. However, accounting for the magnitude of production uncertainty under higher levels of VRE penetration does pose new market challenges. As an analagous example, capacity markets, originally introduced to encourage longer-term generation investment in thermal generation capacity, have increasingly incorporated resource contribution ratings based on Effective Load Carrying Capacity (ELCC) and other probabilistic modeling to account for the intermittency of VRE.[9]

[9] Previously, the main consideration for supply variability was forced outages of thermal generation, which some capacity markets accounted for through seasonal derate factors [7].

1.2.1 What's in a Price?

Of particular relevance to this work is the question about how short-term (i.e. day-ahead timescale) markets and prices should accomodate the increased variability from greater penetration of VRE. One of the responsibilities of the system operator is to facilitate competitive and economically efficient[10] electricity markets. This is done in part by calculating and posting prices that reflect the conditions on the system. In an efficient market without nonconvexities and inelastic demand, the canonical assumption is that the marginal price of electricity should be equal to the marginal cost of generation to meet demand at lowest system costs. However, considerations around uncertainty, reliability, and nonconvexities such as startup costs complicate the question of what constitutes an "efficient" price.[11]

Unlike conventional fossil generators which incur fuel costs as part of production, VRE resources have no marginal fuel costs but do impose additional costs on the system if forecast deviations force another generator, such as a natural gas peaker, to start up and compensate for the shortfall in real-time. Zero marginal fuel costs are compatible with marginal prices under convex conditions, but an increase in supply variability due to increased VRE enhances the need for consideration of how prices are "formed", given that fully controllable generators will need to be started and stopped more frequently to balance out the volatility and thus increase the share of costs associated with nonconvexities. Similar to the application of ELCC to capacity markets, the question of what clearing prices should achieve is a preexisting question, [12] which will become more relevant as the amount of VRE on the system increases.

1.3 Managing Uncertainties

System operators mainly manage short-term uncertainties in supply and demand by procuring a variety of operating reserve products. These reserve products are procured for a variety of requirements such as compenstating for the loss of a generator (contingency reserve), variablility in supply/demand at timescales shorter than that of the energy market intervals (regulation reserve), and keeping headroom/footroom available for accomodating ramping conditions in the net load profile (ramping reserve) [6]. This work, however, considers the question of uncertainty management from the individual agent's perspective instead of the system operator. Whereas the system operator's concern for uncertainty management is to

[10] *"Efficiency is the state of having maximized the sum of consumer and producer surplus; productive efficiency means production costs have been minimized... for the given level of output"* [3].

[11] For a brief graphical and mathematical description of the complications of determining optimal marginal prices under nonconvexity of block-loaded generators, refer to Chapter 3-8 of [3].

[12] One such early example of the clearing price problem, before any consideration of renewables, was when NYISO modified its pricing rules as a result of the nonconvexities of block-loaded gas turbine generators [9].

ensure reliability of system service by means of the centrally determined reserve product requirements, the generator's concern for uncertainty of a financial nature. In theory, it is the generator's offer curve that links these two uncertainty considerations: the central reserve requirement changes the total energy dispatched (and thus, the clearing price of energy) through spillover effects, [13] while changes in energy offer prices affect the opportunity cost payout to generators providing reserves.

Currently, market monitoring to mitigate the exercise of market power is primarily based on checks against the generator's short-run marginal costs [11]. This is commonly taken to mean the fuel costs and operations & maintenance costs. However, as noted in Section 1.2.1, the variable costs of a wind generator, when considering system costs in both the day-ahead markets and real-time markets, also have a component associated with possible underdelivery at the tail end of the wind forecast distribution. Allowing individual agents to express risk preferences through active offer curve creation will also necessitate a thorough regulatory consideration (one that is beyond the scope of this thesis) of what constitutes a reasonable market offer, and by extension, a just and reasonable price. [14]

[13] On December 1, 2023, PJM changed its regulation reserve requirement from a percentage of peak & valley load to a fixed amount of 700MW and 525MW during peak and off-peak hours, respectively. Based on this change, it is empirically estimated that a 100MW change in the regulation market had an approximately 400 MWh change on the energy market [10].

[14] It is a common maxim that economic competition will encourage "fair" prices. However, prices are emergent from the techniques of market participants and the rules and assumptions in a market design. *"Prices, and the ways of price making that stand behind them, are never simply facts or things that emerge out of markets, but instead are ongoing objects of struggle [among market stakeholders]... it is in this 'hidden abode' of price making that some of the most intense forms of rent seeking and conflict between market participants now occur"* [11].

Chapter 2

Problem Framework

We consider the problem of a wind generator optimizing its offer into the day-ahead electricity market while taking into account potential real-time penalties for underproduction under a risk measure (conditional value at risk).

2.1 Optimization Problem Formulation

The conventional treatment of generator offers requires a convex offer function consisting of blocks of power offered at discrete marginal prices. We make the following considerations for the wind generator in our problem:

- The agent takes the day-ahead and real-time price scenarios as predicted (price-taker behavior) and also has a set of discrete forecast scenarios of its real-time wind production.
- All power quantities offered at a price below the day-ahead price clear the market and are sold at that day-ahead price.
- In the day-ahead optimization, the agent should not consider any profits from selling or repurchasing electricity in the real-time market to address any concern that the generator is engaging in physical withholding [1] and could potentially be exercising market power.
- If the realized wind power is less in real-time than was contracted for dispatch in the day-ahead market, the generator must buyback the shortfall at the real-time

[1] *"Physical withholding is a decision by a firm not to offer available generation capacity that it owns or controls into the market, when the short-run marginal cost of the capacity is less than or equal to the competitive market price."* [12]

price.

- The generator is downwards dispatchable to 0 MW from any realized wind scenario and such downwards dispatch incurs no operating cost.

2.1.1 Related Work

The question of creating optimal supply and demand curves given technology constraints has been thoroughly examined in the literature since the beginning of restructuring and as focuses on different generation and load technologies have changed. In particular, we situate this work among the following related works:

- [13] describes a mixed integer linear program for a thermal generator calculating day-ahead offer prices and quantities while considering nonconvex operating constraints. It assumes price-taker behavior and a lognormal distribution of electricity prices.

- [14] considers the problem of constructing offer curves when considering the financial risk of correlated real-time prices and wind underdelivery for wind generators, but constructs smooth offer curves instead of the industry-standard piecewise linear offer.

- [15], [16] describe processes for creating piecewise linear bid curves for electricity consumers when considering, shortfall/surplus penalties, albeit with a different curve structure since the market format for demand-side curves differs from supply-side curves.

- [17] considers the scenario where “slow” thermal generators internalize their redispatch costs using conditional value at risk when making offers and evaluates the overall effects on market efficiency

- Finally, this work expands on some of the discussion in [18], which highlights the role of individual agents internalizing the impact of parameters such as operating constraints that are not necessarily visible to the system operator. In this case, the generator’s individual risk preference for financial returns under different future wind and price scenarios (the “internal” parameter) is expressed through changes in the day-ahead offer curve.

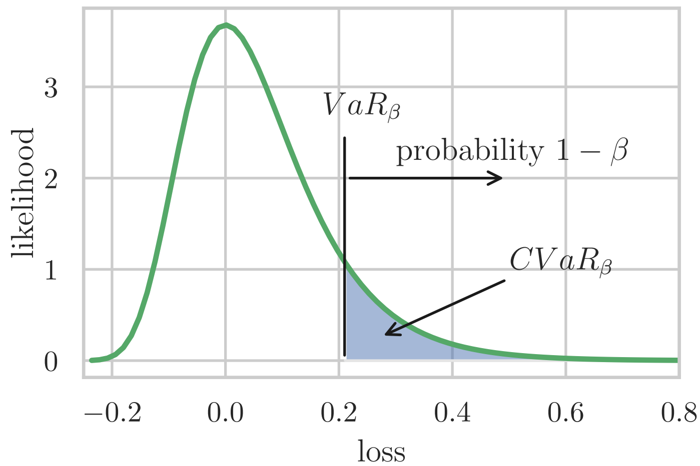

Figure 2.1. Graphical representation of value at risk (VaR) and conditional value at risk (CVaR). Outcomes to the right side of the x axis are less desirable.

2.1.2 Risk Measure: Conditional Value at Risk

In expectation, a certain offer strategy may yield seemingly good profits. However, due to the lack of storage at scale, electricity prices are substantially more volatile than that of other commodities. Hence, it may be necessary to consider a more conservative measure of evaluating the success of a certain strategy. To address the impact of tail events, we use conditional value at risk (CVaR) [19] to formulate our optimization objective.

Value at risk (VaR) and conditional value at risk are popular risk measures which consider the worst-case outcomes in decision making, such as when constructing a portfolio of stocks. For a given cutoff probability β, the VaR_β quantifies a value α for the losses. The probability of incurring a loss above α is equal to $1 - \beta$. The conditional value at risk is equal to the expectation of incurring losses exceeding $\alpha = VaR_\beta$. This is depicted graphically in Fig. 2.1, with $\alpha = 0.2$ and $CVaR_\beta$ equal to the expectation of the blue shaded region.

In addition to its focus on "extreme" tail events to consider the volatile nature of electricity markets, CVaR is also a convex risk measure, making it more well-behaved in optimization problems (VaR is not convex).

Let $f(x, y)$ be the loss associated with a decision x subject to the impact of a random variable y. Let $\rho(y)$ represent the probability density of y. If $\alpha_\beta(x) = VaR_\beta$ and $\phi_\beta(x) = CVaR_\beta$, then [19]:

$$\alpha_\beta(x) = min\{\alpha \in \mathbb{R} : \psi(x, \alpha) \geq \beta\} \tag{2.1}$$

$$\phi_\beta(x) = \frac{1}{1-\beta} \int_{f(x,y) \geq \alpha_\beta(x)} f(x, y)\rho(y)dy \tag{2.2}$$

Where $\psi(x, \alpha)$ is the probability that the loss f does not exceed a threshhold α.

For the context of this thesis, the distribution $\rho(y)$ is not known in closed form. We assume the distribution of price and wind forecast random variables is represented as a collection of samples from an exogenous source, with N_Ω total scenarios. If each scenario y_ω has an associated loss $f(x, y_\omega)$ and assuming $m = (1 - \beta)N_\Omega$ is an integer, then by using the optimization recipe for the m largest entries of a vector [20] we can calculate the $CVaR_\beta(f(x, y)) = \phi_\beta(x)$ with a sample average approximation as:

$$\min_{t,u} \quad mt + \sum_\omega u_\omega \tag{2.3}$$

$$\text{subject to} \quad u_\omega + t \geq f(x, y_\omega) \qquad \omega = 1...N_\Omega \tag{2.4}$$

$$u_\omega \geq 0 \qquad \omega = 1...N_\Omega \tag{2.5}$$

2.1.3 Offer Curve

All ISO-run electricity markets in the United States require generator offers to be in the form of piecewise linear offer curves, where each "piece" specifies a quantity of energy for the dispatch interval. There is much variety in the academic literature about the representation of the curves (e.g. quadratic functions for marginal cost vs. output, continous piecewise linear functions for marginal cost vs. output, discontinous piecewise linear functions for marginal cost vs. output, etc.). For clarity, we briefly present here a description of the offer curves described in this work which are identical to the current offer structure in use for producers.

Table 2.1 is reproduced from the CAISO Business Practices Manual and describes the segments of how generators submit offer curves.[2] For a given operating level segment there

[2] In addition to marginal prices for energy, CAISO and other system operators allow generators to submit startup costs based on their startup condition (i.e. warm, intermediate, and cold) and cost to run at minimum

is a corresponding marginal price the energy in that segment is offered at. Segment 1 corresponds from 70.01MW to 150.00 MW at a price of \$25/MWh, Segment 2 corresponds from 150.01MW to 200.00 MW at a price of \$30/MWh, etc. Fig. 2.2 shows the equivalent marginal-cost offer curve (top) and total cost curve (bottom) for the data in Table 2.1, assuming the unit's minimum load cost is zero. For the purpose of this work, we will use the marginal cost curve representation in the problem formulation (the top graph in Fig. 2.2).

MIQP Formulation: Row Echelon Constraints

The task from the wind generator's point of view is to determine the prices and quantities at which it will offer power into the market. The full expression giving the optimal offer under $CVaR_\beta$, with no consideration of profits in the real-time market (but with consideration of underdelivery penalties) along with the assumptions in Section 2.1 is:

$$\max_{\boldsymbol{x}} \quad CVaR_\beta \left(\lambda^t_{DA}\, p^t_{CL} + \ \lambda^t_{RT}\Delta^t \right) \tag{2.6a}$$

$$\text{subject to} \quad p^{t\omega}_{CL} = \sum_i^{N_{seg}} u^{t\omega i} p^{ti}_{offer} \tag{2.6b}$$

$$max(p^{t\omega}_{CL}) \leq max(p^{t\omega}_{max}) \tag{2.6c}$$

$$\Delta^{t\omega} = min(0, p^{t\omega}_{max} - p^{t\omega}_{CL}) \tag{2.6d}$$

$$u^{t\omega i} \leq 1 - \frac{\lambda^{ti}_{offer} - \ \lambda^{t\omega}_{DA}}{M} \tag{2.6e}$$

$$u^{t\omega i} \geq \epsilon + \frac{\lambda^{t\omega}_{DA} - \lambda^{ti}_{offer}}{M} \tag{2.6f}$$

$$u^{t\omega i} \in \{0, 1\} \tag{2.6g}$$

$$\sum_i^{N_{seg}} u^{t\omega i} \leq \sum_i^{N_{seg}} u^{t(\omega+1)t} \tag{2.6h}$$

$$\lambda^{ti}_{offer} \leq \lambda^{t(i+1)}_{offer} \tag{2.6i}$$

$$u^{(i+1)\omega t} \leq u^{t\omega i} \tag{2.6j}$$

$$\text{where} \quad \boldsymbol{x} = \lambda^{ti}_{offer}, p^{ti}_{offer}$$

Variables λ^{ti}_{offer} and p^{ti}_{offer} specify the price and quantity characteristics of each block in the offer curve (Fig. 2.3). $u^{t\omega i}$ is 1 if a block i clears in a specific scenario and 0 otherwise.

load. This additional information is used to calculate potential uplift payments if the payment from energy sales does not cover these additional costs.

Table 2.1. “Example of energy bid component for a generating unit with a PMin of 70 MW and a PMax of 500 MW” [21]

Segment	Operating Level (MW)	Energy Price ($/MWh)
1	70	25
2	150	30
3	200	35
4	250	40
5	300	45
6	340	50
7	375	55
8	400	60
9	450	65
10	475	75
	500	75

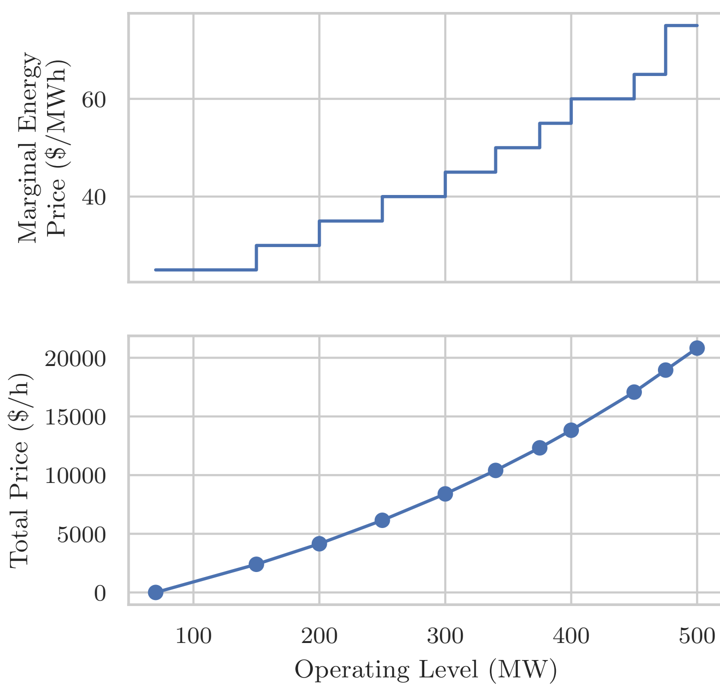

Figure 2.2. Example energy offer curve based on segments in Table 2.1. The top graph is the representation of the marginal cost of each block, while the bottom graph is the total production cost. Both graphs show the same information.

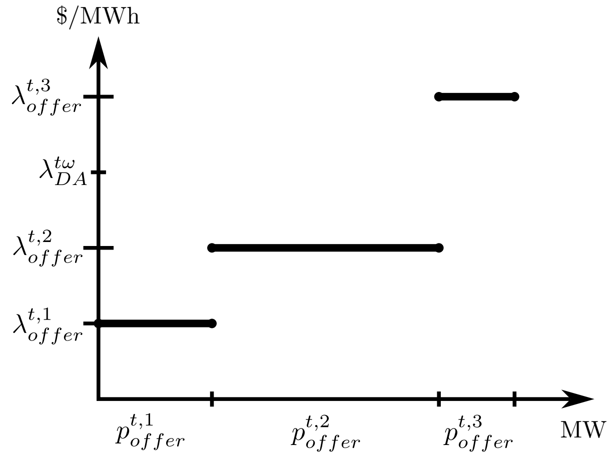

Figure 2.3. Representation of offer curve parameters and corresponding optimization variables under the row-echelon constraint method. Observe that based on the indicated forecast price in scenario ω, only segments 1 and 2 would clear and $[u^{t\omega 1}, u^{t\omega 2}, u^{t\omega 3}] = [1, 1, 0]$

The program can be understood by breaking it down into the following components:

- Calculating power sold in the day-ahead market.
- Calculating penalty for real-time underdelivery.
- Reducing solution space with row-echelon structure.

Calculating power sold in the day-ahead market

The first task is to calculate the power sold in the day-ahead market in each scenario ω. We assume that all power at a price at or below the day-ahead market price of energy $\lambda^{t\omega}_{DA}$ [3] clears the market and is sold at the day-ahead price. [4]

To determine which segments clear the market, we introduce a set of binary auxiliary variables $u^{t\omega i}$ (2.6g) which have a value of 1 if that segment's price is at or below the day-ahead price ($\lambda^{t\omega}_{DA} - \lambda^{ti}_{offer} \geq 0$) and 0 otherwise. For a curve with N_{seg} segments and considering N_{Ω} total scenarios, there would be $N_{seg} * N_{\Omega}$ auxiliary variables $u^{t\omega i}$. Constraints

[3] If the price is at the offer, that implies the offer is marginal and would only partially clear, but it would be impossible to determine the quantity cleared without modeling the demand curve. Thus, for simplicity we assume the entire offer clears.

[4] We also assume here that the agent has information about the energy price component of the locational marginal price; the "total" locational marginal price of energy is a combination of prices from energy, transmission congestion, and transmisson losses. Transmission congestion substantially complicates whether a resource actually clears in "merit order" based solely on energy prices, and so we do not consider its effect here, assuming instead that transmission effects are negligible in the agent's consideration of the clearing process.

(2.6e) and (2.6f) use the big-M trick to enforce this behavior. Once the "active" segments in each scenario are determined, one can sum up the output associated with each segment to find the total amount of power cleared in the day-ahead interval (2.6b).

Finally, we impose an additional constraint on the offer curve that in no scenario can we clear more power than the maximum forecast sample (2.6c).

Calculating penalty for real-time underdelivery

Once we have derived an expression for power cleared in the day-ahead market $p_{CL}^{t\omega}$, we can use it to calculate the real-time penalty, which only applies if there is a shortfall (2.6d); this prevents the strategy from explictly considering withholding day-ahead capacity for real-time profits. To represent this constraint as a linear program, we can transform it into two constraints: one specifying $\Delta^{t\omega} \leq 0$ and another specifying $\Delta^{t\omega} \leq p_{max}^{t\omega} - p_{CL}^{t\omega}$. These two constraints satisfy the spirit of the expression (2.6d) and the problem objective based on the following observations:

- When $p_{max}^{t\omega} - p_{CL}^{t\omega}$ is positive (not all the forecast power cleared in DA), then $\Delta^{t\omega}$ has a tight upper bound of 0.
- When $p_{max}^{t\omega} - p_{CL}^{t\omega}$ is negative (need to buy back in RT), then $\Delta^{t\omega}$ has a tight upper bound of $p_{max}^{t\omega} - p_{CL}^{t\omega}$
- When $\lambda_{RT}^{t\omega}$ is positive, then the solver tries to maximize the value of $\Delta^{t\omega}$ since the overall program is a maximization problem.
- When $\lambda_{RT}^{t\omega}$ is negative, then the problem may be unbounded and the program could try to repurchase energy from the real-time market. In this case, we could add an indicator variable which sets $\Delta^{t\omega}$ to zero. However, since CVaR only considers the worst-case scenarios, such behavior would not affect the objective unless a substantial portion of the scenarios have negative real-time prices. If this were true, then the overall problem becomes trivial, as there is no need to hedge against underdelivery risk as there is no penalty for underdelivery. [5]

Reducing solution space with row-echelon structure

Integer constraints add significant complexity for a solver. We can significantly reduce the feasible space by exploiting the fact that although the overall task to to determine a schedule

[5] Negative prices can result from an overproduction of electricity such as from midday solar or from a lack of supply/demand flexibility combined with generator ramp constraints. CAISO reported a frequency of 11.4% for negative prices in the 15m market for Q2 2023, so such a scenario is not necessarily uncommon [22].

for every hour of the day-ahead market, from the perspective of the optimization problem each time interval's offer curve is not influenced by the solution in other intervals. [6]

Consider the matrix $U^t \in \mathbb{R}^{N_\Omega \times N_{seg}}$ that collects all of the auxiliary variables $u^{t\omega i}$. If we order the rows by increasing day-ahead price in the corresponding scenario ω and constrain each column to represent a block with a higher price than the column preceding it (2.6i), then in each row there must be a cutoff point where all $u^{t\omega i}$ to the left which are smaller than the price $\lambda_{DA}^{t\omega}$ have a value of 1 and all to the right have a value of 0. In other words, the values of $u^{t\omega i}$ in a row are decreasing (2.6j). Since the scenarios are ordered by increasing $\lambda_{DA}^{t\omega}$, we also know that each row must have at least as many active segments compared to the row above it (2.6h). Taken together, these constraints restrict U^t to have a row-echelon structure.

MIQP Formulation: Special Ordered Sets

There was an attempt to formulate a version of the problem using special ordered set constraints to potentially speed up the solver. However, this formulation was unable to fully specify the problem without adding heuristic constraints and so is not guaranteed to give an optimal solution. For this reason, this method is not demonstrated in the results, but the formulation attempt is still presented here for completeness.

Special ordered sets (SOS) are a way to specify constraints that aid the solver software in implementing better branching heuristics for exploring the search space. For a set of variables, a special ordered set of type 1 (SOS1) means that at most one of the variables in the set can be nonzero. We can slightly reformulate the setup as described in Section 2.1.3 to use a SOS1:

[6] This assumption would not hold if there were intertemporal effects such as ramp rate limits or if the price-taker assumption was not true and the quantity sold in an interval influenced the energy prices in another interval.

$$\max_{\boldsymbol{x}} \quad CVaR_\beta \left(\lambda^t_{DA}\, p^t_{CL} + \lambda^t_{RT} \Delta^t \right) \tag{2.7a}$$

$$\text{subject to} \quad p^{t\omega}_{CL} = \sum_{i}^{N_{seg}} u^{t\omega i} p^{ti}_{offer} \tag{2.7b}$$

$$max(p^{t\omega}_{CL}) \leq max(p^{t\omega}_{max}) \tag{2.7c}$$

$$\Delta^{t\omega} = min(0, p^{t\omega}_{max} - p^{t\omega}_{CL}) \tag{2.7d}$$

$$u^{t\omega i} \leq 1 - \frac{\lambda^{ti}_{offer} - \lambda^{t\omega}_{DA}}{M} \tag{2.7e}$$

$$u^{t\omega i} \in \{0, 1\} \tag{2.7f}$$

$$\sum_{i}^{N_{seg}} u^{t\omega i} \leq 1 \qquad \omega = 1 ... N_\Omega \tag{2.7g}$$

$$\lambda^{ti}_{offer} \leq \lambda^{t(i+1)}_{offer} \tag{2.7h}$$

$$\text{where} \quad \boldsymbol{x} = \lambda^{ti}_{offer}, p^{ti}_{offer}$$

Variables λ^{ti}_{offer} and p^{ti}_{offer} specify the price and quantity characteristics of each block in the offer curve, but instead of p^{ti}_{offer} specifying the quantity of energy offered at a given price, it instead gives the total quantity that would clear at a given price (Fig. 2.4). $u^{t\omega i}$ is 1 if the quantity p^{ti}_{offer} clears in a specific scenario and 0 otherwise. In the matrix U^t each row would sum to one or zero; constraint (2.7g) represents the SOS1 constraint.

However, while (2.7e) guarantees that the active segment will have a price lower than the day-ahead price scenario, it does not guarantee that the price in the active segment will be the closest out of all segments' prices to the day-ahead price scenario. Under certain combinations of scenarios, the optimizer will set the active price segment to a looser bound which violates the assumption that all quantity offered at a price below the day-ahead price clears the market. This could be addressed by a heuristic to the objective that penalizes the difference between the active price segment in each scenario and the day-ahead price scenario, but this gives no guarantee of optimality under the original problem intention. Additionally, in informal testing, the SOS1-constrained formulation solved at least three times slower than the row-echelon-constrained formulation.

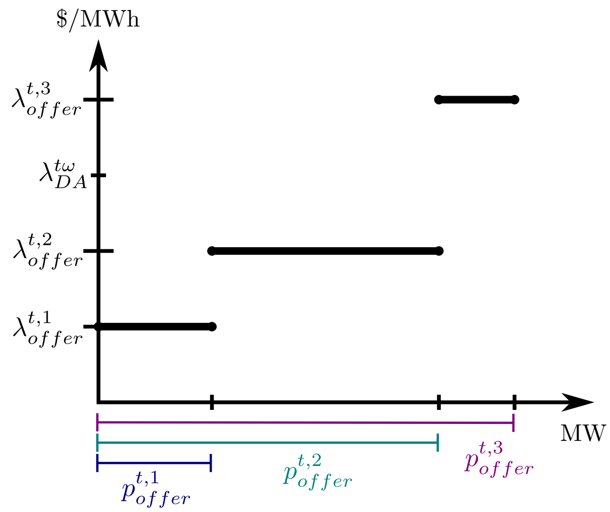

Figure 2.4. Representation of offer curve parameters and corresponding optimization variables under the SOS1 constraint method. Observe that based on the indicated forecast price in scenario ω, an amount equal to p^{t2}_{offer} should clear and $[u^{t\omega 1}, u^{t\omega 2}, u^{t\omega 3}] = [0, 1, 0]$

Chapter 3

Results

This section presents the results of applying the formulation described in Section 2.1.3 to various contexts. We examine the offer curves produced from the program under Gaussian distributions and provide intuition for the effect of the random variables on the curve shapes, examine the performance of the formulation on a real-world example using data from NYISO, and evaluate the runtime performance. The program was formulated in Julia v1.9.2 with the aid of the JuMP library (v1.15.1) and solved using Gurobi v10.0.3. A sample of the Julia code used to formulate the problem is provided in Appendix A. A L2 regularization norm was also added to the price and quantity decisions to improve the conditioning of the program.

3.1 Curve Interpretation

The MIQP described in Section 2.1.3 was run on two different synthetic Gaussian distributions of scenarios to examine the behavior of the program and demonstrate a qualitative method of interpreting the curve shapes. Table 3.1 describes the simulation parameters for the multivariate Gaussian used to generate the two cases. We denote the case where wind and real-time prices are negatively correlated as Case 1, and the case where they are positively correlated as Case 2. Case 1 is the "worst-case" since when a shortfall of wind in real-time is more likely, real-time prices are also more likely to be higher, while Case 2 is a "best-case" since any shortfall will in expectation not have a financial impact, as real-time prices will on average be lower than day-ahead prices. For both cases, we set the program to generate six offer curve segments and then removed segments that are smaller than 1 MW in width to improve readability.

Fig. 3.1 shows the offer curves created by running the MIQP for Cases 1 and 2. In both cases, the right side of the offer curve reflects the nonzero variable costs associated with buyback uncertainty, as would be expected.

For Case 1, increasing β (more risk averse) reduces the amount of total power offered in the day-ahead; the total amount offered is larger than would come from just truncating the marginal wind distribution at the corresponding β alone since there are scenarios where a power shortfall may not result in significant financial loss if the real-time price is also not high in such a scenario.

For Case 2, increasing β decreases the price offered at the right side of the curve when compared to a lower risk aversity (a smaller β value). This result is initially counterintuitive; one would generally expect that a higher risk aversity would result in the uncertain wind portion in the 100 - 125MW range to be priced higher, since that portion will be more likely to be subject to buyback. However, in Case 2, real-time prices and wind are positively correlated, meaning that any wind shortfall will generally be repurchased a lower price in real-time than the price in the day-ahead market; put differently, in expectation there is no financial penalty for underdelivery, and the bulk of the "loss" which is targeted by optimizing CVaR consists of insufficient revenue that clears in the day-ahead market. We can qualitatively justify this claim by examining the joint distribution plot in Fig. 3.1d, which shows the "active" samples of day-ahead price, real-time price, and wind that constitute the tail of the profit distribution that CVaR is calculated from. At the higher β values, the active samples are clustered where day-ahead prices, real-time prices, and wind outputs are all lower than the mean. Since the positive covariance means that there is less penalty for underdelivery, the offer curve is constructed to clear power that may not be deliverable in the real-time market to gain the additional day-ahead revenue.

3.2 NYISO Performance

To investigate the real-world performance of the program, the MIQP was run on real-world price data from the New York Independent System Operator and synthetic wind forecast data corresponding to a wind farm in the NYC price zone. We compare the profits generated when the MIQP is run for varying β against a naive offer strategy that offers at zero marginal price in the day-ahead market, up to a percentile quantity of the forecast wind powers $p^{t\omega}_{max}$. Hourly-level simulations are run for the period of October 1st – 31st, 2019, with this specific date range chosen since ensemble wind forecasts were available in the year 2019 and October empirically had high variability in electricity prices.

Table 3.1. Simulation parameters for synthetic Gaussian data cases

Parameter	Symbol	Value
Mean day-ahead price	μ_{DA}	30 $/MWh
Mean real-time price	μ_{RT}	30 $/MWh
Mean wind forecast	μ_{wind}	100 MW
Variance, day-ahead price	σ^2_{DA}	100
Variance, real-time price	σ^2_{RT}	100
Variance, wind forecast	σ^2_{wind}	100
Covariance*, RT price & wind	$cov_{RT,wind}$	$\pm$ 80
Number of scenarios generated	N_Ω	250

*All other covariances are zero.

3.2.1 Data Inputs

For each operating day, we generate price samples by taking the historical day-ahead and real-time prices in the corresponding hour over the preceeding fifty calendar days. These price scenarios are paired with one hundred wind forecast scenarios generated from historical forecast data using PGscen [1] [23] to create one hundred total scenarios for the hour of interest.

3.2.2 Regret Comparison

For both the “MIQP CVaR” offer strategy and the “forecast percentile” strategy, the corresponding offer strategy is used to create the day-ahead offer curve. After the day-ahead clearing quantity is determined from the actual day-ahead price[2], the shortfall (or excess) wind production relative to the day-ahead clearing quantity is bought-back (sold) at the real-time price.

To account for natural day-to-day variability in revenue, we subtract the total profit from the day-ahead and real-time markets under these two strategies from an idealized profit number calculated by selling all the realized wind production in whichever market is more profitable in hindsight.[3] We refer to this difference as the regret. For example, if the

[1] PGScen is an open-source Python package which uses a graphical model to generate wind, solar, and demand forecasts while preserving conditional dependencies between the three variables.

[2] We assume all power priced at or under the day-ahead price clears the market.

[3] This profit is lower than the absolute maximum profit of offering the generator’s entire nameplate

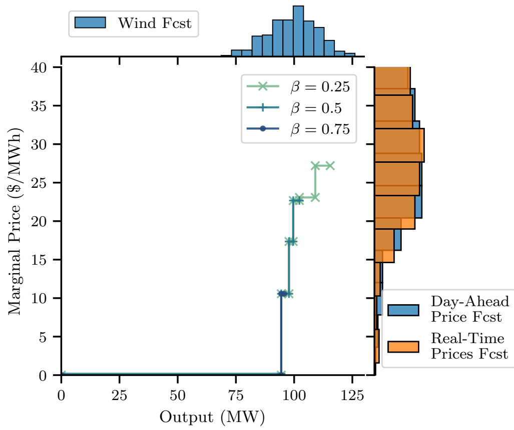

(a) Offer curve for Case 1.

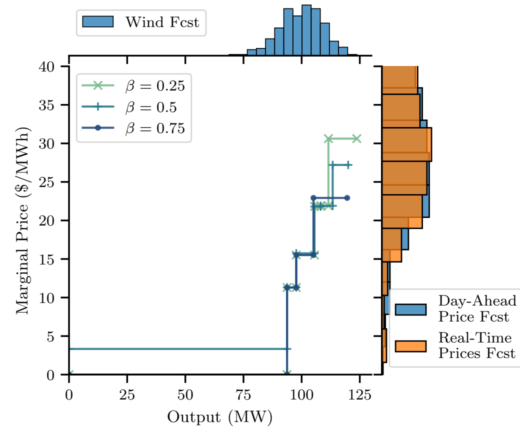

(b) Offer curve for Case 2.

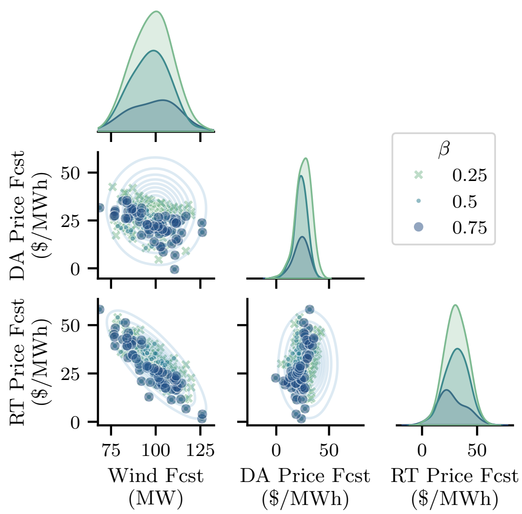

(c) Active samples for Case 1 optimal solution.

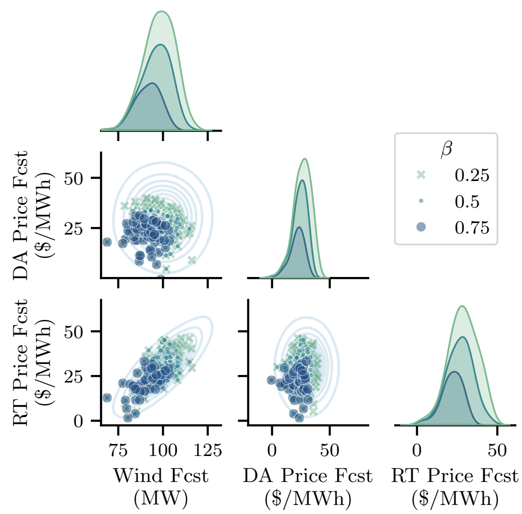

(d) Active samples for Case 2 optimal solution.

Figure 3.1. Offer curves and pairwise visualizations of distributions. The diagonal plots on the lower figures show the marginal distribution of active scenarios corresponding to a single random variable. Countours on the off-diagonal represent the level set of the joint probability density function for the corresponding label on the row and column.

actual day-ahead hourly price was \$10/MWh, the actual real-time price was \$20/MWh, and the actual maximum wind production was 100 MW for that hour, the ideal profit would be \$200 for that hour. If in that hour the MIQP CVaR strategy yielded a total profit of \$150, the regret would be \$50 = \$200 − \$150. Smaller regret values are indicative of a better

capacity in the most profitable market instead of the realized wind power, but in practicality such a strategy could be penalized by the market monitors for misrepresenting actual production.

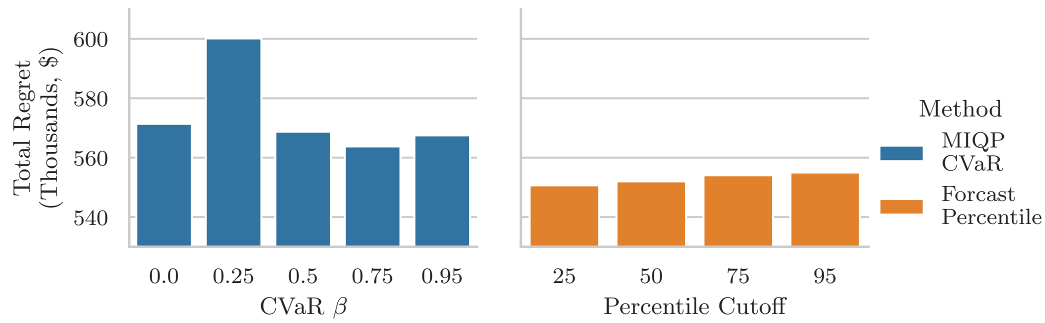

(a) Strategy performance considering the sum of regret.

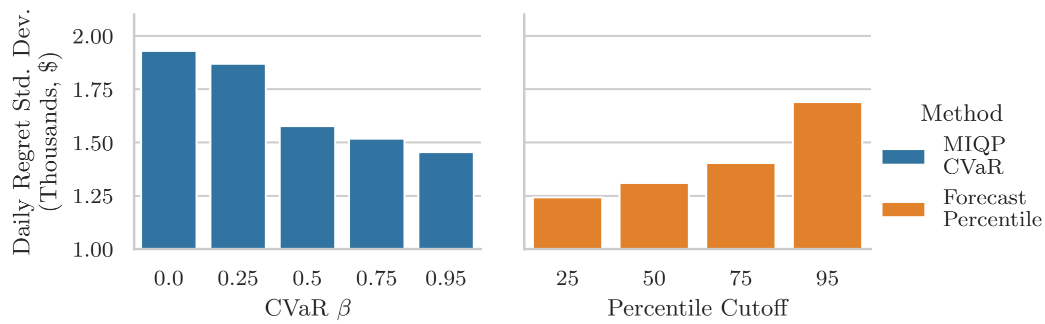

(b) Strategy performance considering the standard deviation of regret.

Figure 3.2. Strategy performance on NYISO data. For the CVaR MIQP approach (blue columns), risk aversity increases from left to right. For the percentile cutoff approach, risk averisty increases from right to left. Lower regret scores are more desirable.

strategy.

Fig. 3.2a shows the total regret over different choices of risk preference β. There is no directional influence of β on the total regret. The MIQP CVaR strategy also underperforms the percentile-based offer strategy in the day-ahead market.

Fig. 3.2b shows the average of the standard deviations of the daily regret. Increasing the β value (more risk averse) decreases the variance of the daily regret, as expected. However, the MIQP CVaR strategy still underperforms the naive 25% and 50% forecast percentile strategies in both total regret and variance of regret, indicating that there is still room for improvement in the MIQP CVaR approach, possibly through improved scenario creation.

Furthermore, under a perfectly characterized forecast distribution (for both prices and wind power), one would expect there to be a Pareto frontier along which variance in market performance could be traded off against the expectation of overall performance (the "bias-variance tradeoff" in statistics and machine learning). The observation that increasing the risk aversity parameter does not lead to decrease in the total regret indicates that there are

Table 3.2. Runtime performance for CVaR MIQP formulation. Runtime is the median out of six test cases.

Curve Segments	Scenarios	W/out Row Echelon Constraints		With Row Echelon Constraints	
		Runtime[1](s)	Standard Deviation[2](s)	Runtime[1](s)	Standard Deviation[2](s)
3	50	0.4	0	0.4	2.9
	100	0.9	0.1	1.3	0.3
	250	3.4	4.1	5.8	0.8
	500	16.1	9.4	22.5	15.9
6	50	4.6	2.4	5.5	4.3
	100	9.6	7.3	35.7	11.5
	250	97.7	61.3	175.5	75.2
	500	1440.7	947.7	1537.2	927.4
9	50	22.8	12.1	29.9	7.8
	100	99.1	197.7	107.5	246.6
	250	2700	n/a	1848.4	n/a
	500	2700	n/a	2700	n/a

[1] Solver timelimit was set to 2700 seconds; cases that terminated with status exceeding the timelimit were recorded as 2700 s.

[2] Standard deviation was not calculated for instances where at least two out of the six cases terminated due to exceeding the optimizer time limit.

additional performance gains to be made in reducing total regret without comprimising the hedging performance of the program. Since the forecasting method for prices is crude, a substantial improvement in performance could probably be achieved through better price predictions and having these price predictions be linked to the wind forecast instead of the two being generated independently.

3.3 Runtime Performance

Integer variables add significant computation complexity for solvers; linear programs can be solved in polynomial time while integer programming is NP-hard. Table 3.2 and Figure 3.3 show the runtime statistics for a set of six test cases. Price and wind scenarios in each test case are generated by sampling a multivariate Gaussian distribution. When the program was run for a curve with nine segments with 250 scenarios, at least half of the cases terminated without an optimal solution because they reached the preset solver time limit of 2700 seconds.

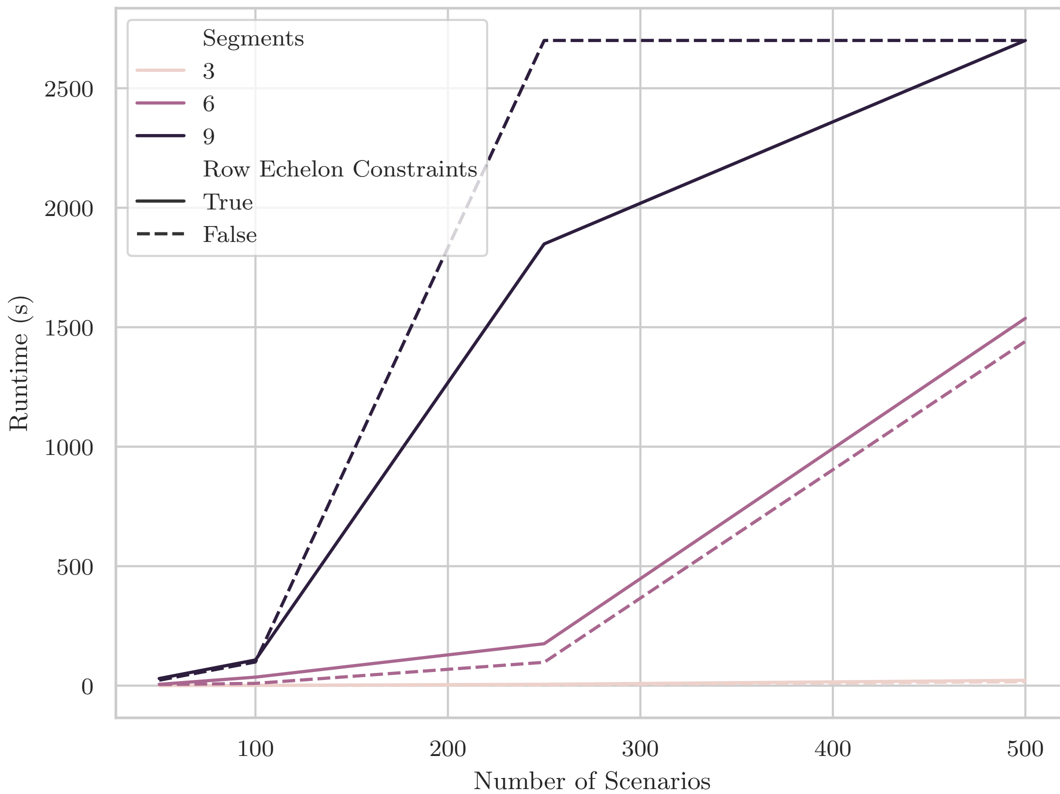

Figure 3.3. CVaR MIQP formulation runtime statistics corresponding to Table 3.2

All computations were performed on a laptop with an Intel Core i5-8300H CPU (2.30GHz) and 31.2 GB of memory. Figure 3.3 shows the trend in median runtime as the problem difficulty increases through adding more curve segments and price & wind scenarios. For curves of three or six segments, adding the additional row-echelon constraints increases the runtime a small amount. However, as the problem becomes more difficult adding the additional row-echelon constraints decreases the solution time. The data is incomplete since at the most difficult level there were several instances in which the solver did not terminate before the time limit, but it is likely that the row-echelon constraints positively interact with Gurobi's branch and bound algorithm to reduce solution time as the problem becomes more difficult with more integer variables.

Day-ahead market offers need to be calculated and submitted within a period on the order of hours. Since the offer for each hour can be calculated independently of the other offers, it is conceivable that with parallel solver instances, a generator could practically utilize this formulation to calculate its day-ahead offers.[4]

[4] The runtime will be strongly dependent on the number of scenarios considered. Such a number is typically determined heuristically based on the application and actuarial experience, but a credible estimate would be on the order of 10^3 [24]

Chapter 4

Conclusion

The electricity industry is faced with the monumental task of integrating significant quantities of variable renewable generation and managing increasingly volatile loads within the next decades. Due to physical constraints and lack of large-scale storage, electricity as a commodity has a high price volatility. This price volatility will likely only increase as VRE penetration increases. In the face of increased volatility, generators may want to hedge their financial risks; this is especially true of VRE generators, which will have a buy-back penalty if they are dispatched for an amount of energy that exceeds their actual real-time output. This thesis presents a mixed-integer quadratic program that utilizes conditional value at risk to determine a piecewise linear offer curve for the day-ahead market which follows the standard format of the piecewise supplier offer used in the current electricity market design (Section 2.1.3).

Mixed-integer programs can be computationally intractable. Previous works instead approximated the problem as one where quantities are fixed and the only decision variable is price. This leads to a mixed-integer linear program (MILP). In this work, we present a mixed-integer quadratic program (MIQP) where both quantity and price are decision variables. Under realistic input complexity the MIQP approach presented is empirically solvable on the timescale of hours that a generator in the day-ahead market would have to make its offer decision (Section 3.3). Further computation speedup could be achieved by reducing the problem to a MILP by fixing a large number of offer quantities and then reducing the number of segments to the standard offer amount by merging approximately colinear segments as in as in [16]. It is unclear how much benefit is gained under specific applications and realistic conditions from not adopting the MILP approximation; future work should consider the tradeoff in solve time with any decreased financial performance, if any. Additionally, if intertemporal effects such as ramp constraints or intertemporal price maker effects are to be considered, then the additional computational complexity would practically necessitate a MILP formulation.

Under some illustrative distributions for price and wind, joint distribution plots marking the scenarios under consideration can qualitatively illustrate how changes in the scenario distributions and generator risk aversity influence the curve shape (Section 3.1).

In a test on historical NYISO price data, the generator can successfully reduce the variance of total profits from electricity sales by increasing its risk aversity through varying the CVaR β parameter. However, the MIQP CVaR strategy does not outperform a naive approach wherein the generator takes a percentile of the forecast as the quantity offered and fixes the price at zero. It is hypothesized that the underperformance is due to insufficiently accurate scenario forecasts as well as a lack of sufficient price volatility in present electricity markets to reward hedging strategies (Section 3.2).

It is an economic orthodoxy that in a competitive market suppliers should price their offers at their marginal costs. In the context of generators in electricity markets, this has sometimes been taken as generators' offer curves should reflect their fuel costs. If, instead, one considers the variable costs of underdelivery in real-time, the MIQP CVaR program presented here can be interpreted as valuing the variable costs associated with buyback risk for wind generators at different levels of their output. Such a valuation effort could tend to increase electricity prices in the day-ahead market if that segment of energy previously offered at zero price was marginal, but is now offered at a positive price. Such an effect could be interpreted as an attempt by a generator to exercise market power, but conversely one could also make the case that this price better reflects the generator's risk preferences which would otherwise be lost if generators were forced to offer at prices reflecting their fuel costs. This financial risk could be financially managed through a series of private instruments instead of through the ISO-run day-ahead markets. However, the market for electricity derivatives is relatively illiquid and relying on a separate market outside of the system operator's wholesale market could also obscure risk information from the system operator.[1] If generators (and loads) create offer curves that reflect some level of strategization which considers the financial risks arising from uncertainty inherent to the technology, new frameworks for monitoring and mitigating market power will also need to be considered to delinate between acceptable hedging and smart bidding behavior and unacceptable attempts to excessively profit off of changing system prices. Key to such consideration will be if the encouragement of individual generators' hedging behavior ultimately has positive impacts for general system reliability and the public interest.

[1] If the generator were to manage its financial risk using private means, the price information and thus the generator's understanding of its financial uncertainty would be obscured from the system operator and other market participants. It is an open question whether representing such risks through an ISO-run market would align with the public interest.

Appendix A

Code

```
using JuMP
using Gurobi
using TensorCast


struct Scenarios
    prices_dayahead::Union{Matrix,Vector}
    prices_realtime::Union{Matrix,Vector}
    wind_powers::Union{Matrix,Vector}
end


function risk_aware_offer(
        scenarios::Scenarios, beta::Number, n_segments::Int;
        regularization=1e-9, row_echelon_constraints=true)
    N = size(scenarios.prices_dayahead)[1]

    # Handle 1D vectors
    prices_dayahead = prices_dayahead[ordering]
    prices_realtime = prices_realtime[ordering]
    wind_powers = wind_powers[ordering]

    # Tx1
    pmax, _ = findmax(wind_powers, dims=1)
    n_cvar = convert(Int, round((1 - beta) * N))
    model = Model()

    set_optimizer(
        model,
        optimizer_with_attributes(Gurobi.Optimizer)
    )

    # Dynamically set big M based on the prices
    M = maximum(abs.(prices_dayahead))
    ϵ = 1e-6
    # Regularization term
```

```
λ = regularization

@variable(model, 0 <= quantities[1:T, 1:n_segments], start = 1)
@variable(model, 0 <= prices[1:T, 1:n_segments])

# Force the offer prices to be increasing so we dont have to manually
# order them later
@constraint(model, [i = 1:n_segments-1],
    prices[:, i] .+ ϵ <= prices[:, i+1])

@constraint(model, quantity_max,
    sum(quantities, dims=2) .<= reshape(pmax, (:, 1)))
# Binary indicator variable with dimensions N x T x n_segments of which
# segments of the offer curve clear.
@variable(model, active[1:N, 1:T, 1:n_segments], binary=true)
# How much power must be repurchased as real-time shortfall, negative if
# purchase is required.
@variable(model, buyback_rt[1:N, 1:T] .<= 0)

# The power quantity cleared in each scenario, final dimension N x T
@cast quantity_cleared[i, j, k] := active[i, j, k] .* quantities[j, k]
@expression(model, quantity_cleared, quantity_cleared)
# Quantity cleared is in the day-ahead market
quantity_cleared_total = dropdims(sum(quantity_cleared, dims=3), dims=3)
@expression(model, quantity_cleared_total, quantity_cleared_total)

@constraint(model, constraint_balance,
    quantity_cleared_total .+ buyback_rt .<= wind_powers)

profits = dropdims(
    sum(
        prices_dayahead .* quantity_cleared_total +
        prices_realtime .* buyback_rt,
        dims=2
    ), dims=2
)

# As a generator, cannot end up purchasing power on balance in real time
@constraint(model, constraint_rt_positive,
    quantity_cleared_total .+ buyback_rt .>= 0)

# Constraints for the indicator variable of which segment is active;
# Non-neg price difference means that segment is active
@cast pricediff[i, j, k] := (
    (prices_dayahead[i, j] - prices[j, k])
)

@expression(model, pricediff, pricediff)
@constraint(model, constraint_active_zero, 1 .+ pricediff ./ M .>= active)
@constraint(model, constraint_active_one, pricediff ./ M .+ ϵ .<= active)
# Impose row-echelon constraints on active segment matrix
if row_echelon_constraints
    for i = 1:N
        for j = 1:(n_segments-1)
            @constraint(
```

```
                    model, active[i, 1, j] >= active[i, 1, j+1]
                )
            end
            if i != N
                @constraint(
                    model, sum(active[i, :, :]) <= sum(active[i+1, :, :])
                )
            end
        end
    end

    @variable(model, quantity_reg >= 0)
    @constraint(model, quantity_reg >= sum(quantities .^ 2))
    @variable(model, price_reg >= 0)
    @constraint(model, price_reg >= sum(prices .^ 2))

    # Take the n worst cast scenarios.
    @variable(model, u[1:N] >= 0)
    @variable(model, t)
    @constraint(
        model, u .+ t >= -profits)
    @objective(model, Min,
        n_cvar * t + sum(u) + λ * (quantity_reg + price_reg))
    return model, prices_dayahead, prices_realtime, wind_powers

end
```

References

[1] S. Blumsack, *Measuring the Benefits and Costs of Regional Electric Grid Integration*, 2007.

[2] S. Borenstein and J. Bushnell, *The U.S. Electricity Industry After 20 Years of Restructuring*, 2015.

[3] Steven Stoft, *Power System Economics: Designing Markets for Electricity*. John Wiley & Sons, Inc., 2002, ISBN: 978-81-265-5391-4.

[4] P. Twomey, R. Green, K. Neuhoff, and D. Newbery, *A Review of the Monitoring of Market Power: The Possible Roles of TSOs in Monitoring for Market Power Issues in Congested Transmission Systems*, Mar. 2005.

[5] New York Independent System Operator, *Manual 11: Day-Ahead Scheduling Manual*, Nov. 2023.

[6] E. Ela, M. Milligan, and B. Kirby, "Operating Reserves and Variable Generation," Tech. Rep. NREL/TP-5500-51978, Aug. 2011. DOI: 10.2172/1023095.

[7] Cynthia Bothwell and Benjamin F. Hobbs, *Crediting Renewables in Electricity Capacity Markets: The Effects of Alternative Definitions upon Market Efficiency*, Jun. 2016.

[8] New York Independent System Operator, "Operations Performance Metrics Monthly Report: September 2023 Report," New York Independent System Operator, Tech. Rep.

[9] U.S. Federal Energy Regulatory Commission, *Order on motion to implement hybrid fixed block pricing rule and requiring tariff filing, acting on related requests for rehearing, and accepting preliminary report*, Apr. 2001.

[10] J. Buchsbaum, C. Hausman, J. Mathieu, and J. Peng, "Spillovers from Ancillary Services to Wholesale Energy Markets," Sep. 2022.

[11] W. Boyd, "Ways of Price Making and the Challenge of Market Governance in U.S. Energy Law," *Minnesota Law Review*, 2020.

[12] "The Guide to Energy Market Manipulation," Global Competition Review, Tech. Rep., 2018.

[13] A. Conejo, F. Nogales, and J. Arroyo, "Price-taker bidding strategy under price uncertainty," *IEEE Transactions on Power Systems*, vol. 17, no. 4, pp. 1081–1088, Nov. 2002, ISSN: 0885-8950. DOI: 10.1109/TPWRS.2002.804948.

[14] H. Shin, D. Lee, and R. Baldick, "An Offer Strategy for Wind Power Producers That Considers the Correlation Between Wind Power and Real-Time Electricity Prices," *IEEE Transactions on Sustainable Energy*, vol. 9, no. 2, pp. 695–706, Apr. 2018, ISSN: 1949-3037. DOI: 10.1109/TSTE.2017.2757501.

[15] Ya'an Liu and Xiaohong Guan, "Purchase allocation and demand bidding in electric power markets," *IEEE Transactions on Power Systems*, vol. 18, no. 1, pp. 106–112, Feb. 2003, ISSN: 0885-8950. DOI: 10.1109/TPWRS.2002.807063.

[16] G. Ruan, H. Zhong, B. Shan, and X. Tan, "Constructing Demand-Side Bidding Curves Based on a Decoupled Full-Cycle Process," *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 502–511, Jan. 2021, ISSN: 1949-3061. DOI: 10.1109/TSG.2020.3012562.

[17] X. Yin, M. D. Ilic, and B. Sinopoli, "Toward design of risk-based real-time dispatch at value," in *2015 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*, Feb. 2015, pp. 1–5. DOI: 10.1109/ISGT.2015.7131810.

[18] M. D. Ilić, J.-Y. Joo, L. Xie, M. Prica, and N. Rotering, "A Decision-Making Framework and Simulator for Sustainable Electric Energy Systems," *IEEE Transactions on Sustainable Energy*, vol. 2, no. 1, pp. 37–49, Jan. 2011, ISSN: 1949-3037. DOI: 10.1109/TSTE.2010.2074217.

[19] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *The Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000, ISSN: 14651211. DOI: 10.21314/JOR.2000.038.

[20] *MOSEK Modeling Cookbook: Release 3.3.0*, Oct. 2023.

[21] California Independent System Operator, *Business Practice Manual for Market Instruments*, Nov. 2023.

[22] California Independent System Operator, *Q2 2023 Report on Market Issues and Performance*, Nov. 2023.

[23] R. Carmona and X. Yang, *Joint Stochastic Model for Electric Load, Solar and Wind Power at Asset Level and Monte Carlo Scenario Generation*, Sep. 2022. arXiv: 2209.13497.

[24] D. Ingram, “How Many Scenarios?” *Risks and Rewards Newsletter*, no. 40, Oct. 2002.