Titre: Title:	A Distributionally Robust Optimization Strategy for Virtual Bidding in Two-Settlement Electricity Markets
Auteur: Author:	Xavier Audet
Date:	2025
Type:	Mémoire ou thèse / Dissertation or Thesis
Référence: Citation:	Audet, X. (2025). A Distributionally Robust Optimization Strategy for Virtual Bidding in Two-Settlement Electricity Markets [Master's thesis, Polytechnique Montréal]. PolyPublie. https://publications.polymtl.ca/71208/

URL de PolyPublie: PolyPublie URL:	https://publications.polymtl.ca/71208/
Directeurs de recherche: Advisors:	Antoine Lesage-Landry
Programme: Program:	Génie énergétique

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets

XAVIER AUDET
Département de génie électrique

Mémoire présenté en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
Génie énergétique

Décembre 2025

POLYTECHNIQUE MONTRÉAL
affiliée à l'Université de Montréal

Ce mémoire intitulé :

A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets

présenté par **Xavier AUDET**
en vue de l'obtention du diplôme de *Maîtrise ès sciences appliquées*
a été dûment accepté par le jury d'examen constitué de :

Keyhan SHESHYEKANI, président
Antoine LESAGE-LANDRY, membre et directeur de recherche
Lesia MITRIDATI, membre externe

DEDICATION

Up the Villa!

ACKNOWLEDGEMENTS

I would like to express my deep gratitude to all the people who contributed, directly or indirectly, to the completion of this Master's thesis.

First, I warmly thank my research supervisor, Antoine, for his trust, his guidance, his scientific rigor, and his support throughout this project. His assistance was essential at every stage of my journey.

I would also like to thank my friends at LORER for their camaraderie, their stimulating discussions, and their valuable help, both at school and at McCarold.

To my friends outside the lab, thank you for your curiosity, your encouragement, and your ability to pull me away from work when it was needed and when it wasn't.

I am deeply grateful to my family for their support, their help, their encouragement to pursue my studies, and for being an infinite source of inspiration.

To Florence, for her patience, her attentiveness, and her constant presence. Thank you for being my greatest supporter.

Finally, I would like to thank my colleagues for their collaboration, for the constructive exchanges that enriched my thinking, and for Wednesday soccer.

To all those who recognize themselves here, thank you from the bottom of my heart.

Je tiens à exprimer ma profonde gratitude à toutes les personnes qui ont contribué, de près ou de loin, à la réalisation de ce mémoire.

Tout d'abord, je remercie chaleureusement mon directeur de recherche, Antoine, pour sa confiance, ses conseils avisés, sa rigueur scientifique et son accompagnement tout au long de ce projet. Son soutien a été essentiel à chaque étape de mon parcours.

Je souhaite également remercier mes amis du LORER, pour leur camaraderie, leurs discussions stimulantes et leur aide précieuse, tant à l'école qu'au McCarold.

À mes amis en dehors du labo, merci pour votre curiosité, vos encouragements et votre capacité à me faire décrocher du travail quand il le fallait et quand il ne le fallait pas.

Je suis profondément reconnaissant envers ma famille, pour leur soutien, pour leur aide, pour leurs encouragements à poursuivre des études et pour être une source d'inspiration semi-infinie.

À Florence, pour sa patience, son écoute et sa présence constante. Merci d'être ma plus grande supportrice.

Enfin, je remercie mes collègues, pour leur collaboration, pour les échanges constructifs qui ont enrichi ma réflexion et pour le soccer du mercredi.

À tous ceux qui se reconnaissent ici, merci du fond du cœur.

RÉSUMÉ

Ce mémoire propose une stratégie de courtage des virtuels sur les marchés de l'électricité de gros, basée sur une approche d'optimisation robuste en distribution (DRO) combinée à une mesure de risque par la valeur conditionnelle à risque (CVaR). Cette stratégie vise à maximiser les profits tout en limitant les risques dans un environnement de marché caractérisé par une forte incertitude et volatilité des prix. La stratégie `DRO-CVaR` est formulée comme un problème d'optimisation convexe, où les quantités d'énergie à soumettre sont déterminées en tenant compte d'un ensemble de distributions probables des écarts de prix (*DART spreads*), défini par une boule de Wasserstein autour de la distribution empirique. Le modèle intègre également la CVaR pour atténuer les risques liés aux fluctuations extrêmes des prix, en équilibrant les pertes attendues et les pertes en queue de distribution selon un facteur d'aversion au risque. La stratégie est mise en œuvre dans un cadre d'optimisation en temps réel, où les données disponibles sont mises à jour quotidiennement. Pour chaque jour de soumission de position de courtage, un ensemble de jours similaires est sélectionné à partir d'une fenêtre mobile de deux ans, en fonction de profils de charge et de capacité de génération hors ligne. Les hyperparamètres du modèle, incluant la taille de l'ensemble de données, le rayon de la boule de Wasserstein, le niveau de risque et les bornes d'incertitude, sont optimisés sur une période d'entraînement de douze mois à l'aide de la librairie `Optuna`, avec pour objectif la maximisation du ratio de Calmar.

La stratégie est ensuite testée sur une période de huit mois à l'aide de données du marché *New York Independent System Operator* (NYISO). Les résultats illustrent que la stratégie `DRO-CVaR` surpasse les autres approches étudiées, i.e., une optimisation stochastique basée sur des scénarios (`SO`), une optimisation stochastique avec CVaR (`SO-CVaR`), une approche robuste en distribution sans CVaR (`DRO`) et une stratégie statique qui vend systématiquement les DART spreads équitablement sur tous les bus (`EW`) en matière de performance ajustée au risque. La stratégie proposée obtient les meilleurs résultats en termes de profit par MWh, de ratio de Sharpe et de ratio de Calmar. Bien qu'elle arrive en deuxième position en termes de profits cumulés, elle se distingue par sa robustesse et sa régularité, n'enregistrant que peu de pertes sur la période testée.

En conclusion, ce mémoire propose une approche combinant DRO et CVaR, laquelle permet de concevoir une stratégie de courtage des virtuels efficace et résiliente dans les marchés de l'électricité de gros. Les perspectives futures incluent l'évaluation de la stratégie sur d'autres marchés ainsi que la résolution d'un écoulement de puissance linéarisé (DC-OPF) sous divers

scénarios de charge, de génération et de topologie pour construire une distribution empirique des prix de congestion qui seront utilisés dans le cadre des enchères des produits de congestion. Cette approche permettra de déplacer l'incertitude des prix vers les paramètres influençant l'exploitation du réseau.

ABSTRACT

This Master's thesis proposes a data-driven strategy for virtual bidding in two-settlement wholesale electricity markets, designed to address price uncertainty and volatility inherent to these markets. The approach combines distributionally robust optimization (DRO) with conditional value-at-risk (CVaR) to create a risk-aware bidding framework that maximizes expected returns while mitigating downside risk. The virtual bidding problem is formulated as a convex optimization model where bid quantities are optimized under an ambiguity set defined by a Wasserstein ball around the empirical distribution of historical price spreads between the day-ahead market (DAM) locational marginal price (LMP) and the real-time market (RTM) LMP (DART spreads). CVaR is incorporated to control tail risk according to a predefined confidence level and risk-aversion factor, ensuring robustness against extreme price fluctuations. The proposed strategy combining DRO and CVaR (`DRO-CVaR`) operates in an online fashion, updating daily with data from a rolling two-year window of historical data. For each bidding day, a subset of similar past days is selected based on load profiles and offline generation capacity, ensuring that the model reflects current system conditions. Hyperparameters, including the dataset size, Wasserstein radius, risk-aversion level, and uncertainty bounds, are tuned over a 12-month training period using `Optuna` to maximize the Calmar ratio.

The strategy is then tested on eight months of NYISO market data. Results show that `DRO-CVaR` outperforms four benchmarks: scenario-based stochastic optimization (`SO`), stochastic optimization with CVaR (`SO-CVaR`), distributionally robust optimization without CVaR (`DRO`), and an equally weighted static portfolio (`EW`). It ranks first in Sharpe and Calmar ratios as well as in profit per MWh, while maintaining consistent gains and limiting losses. Although it ranks second in cumulative profit, the fact that it is the only strategy that never has cumulative profits dip below zero makes it a compelling approach for risk-adjusted performance.

In conclusion, combining DRO and CVaR enables the design of an effective and resilient virtual bidding strategy under high market uncertainty. Future work includes extending the approach to other markets and integrating a linearized optimal power flow (DC-OPF)-based congestion price modelling under varying load, generation, and topology scenarios to trade financial transmission rights (FTRs). This approach will allow us to shift price uncertainty toward controllable network parameters.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF SYMBOLS AND ACRONYMS

Acronyms

CAISO	California Independent System Operator
CVaR	Conditional Value-at-Risk
DAM	Day-Ahead Market
DART	Day-Ahead minus Real-Time
DC-OPF	Direct Current Optimal Power Flow
DRO	Distributionally Robust Optimization
EMD	Earth Mover Distance
ERCOT	Electric Reliability Council of Texas
EW	Equally Weighted
FERC	Federal Energy Regulatory Commission
FTR	Financial Transmission Rights
ISO	Independent System Operator
ISO-NE	ISO New England
LMP	Locational Marginal Price
LSTM	Long Short-Term Memory
MDD	Maximum Drawdown
MISO	Midcontinent Independent System Operator
NESO	National Energy System Operator
NYISO	New York Independent System Operator
OPF	Optimal Power Flow
PJM	PJM Interconnection LLC
RNN	Recurrent Neural Network
RTO	Regional Transmission Organization
RTM	Real-Time Market
SCED	Security-Constrained Economic Dispatch
SCUC	Security-Constrained Unit Commitment
SO	Stochastic Optimization
SPP	Southwest Power Pool
VaR	Value-at-Risk

Symbols

α	Confidence level for CVaR (risk level)
$\boldsymbol{\alpha}_{\min}, \boldsymbol{\alpha}_{\max}$	Dual variables to the generator limits constraint
β	Threshold in VaR definition
$\boldsymbol{\gamma}$	Variable vector in uncertainty set reformulation
Γ	Similarity metric between days
$\delta_{\mathbf{s}_t^d}$	Dirac measure at point $\mathbf{s}_t^d$
ϵ	Wasserstein ball radius (ambiguity set parameter)
ζ	Penalty parameter for weekday/weekend distinction
η_j	Scaled return on day j
γ	Dual variable to the power balance constraint
$\gamma^{\mathrm{E}}, \gamma^{\mathrm{L}}, \gamma^{\mathrm{C}}$	Components of LMP (Energy, Losses, Congestion)
θ	Offline generation capacity (MW)
Λ	Maximum bound for DART spread components
λ	Auxiliary variable in DRO reformulation
μ_k^d	Scalar variable in uncertainty set reformulation
$\boldsymbol{\mu}_{\min}, \boldsymbol{\mu}_{\max}$	Dual variables associated to transmission limits constraint
$\boldsymbol{\nu}_{t,k}^d$	Vector variable in uncertainty set reformulation
$\boldsymbol{\pi}$	Joint probability distribution (transport plan)
ρ	Risk-aversion factor
σ	Dual variable to the loss function constraint
σ_η	Standard deviation of scaled daily returns
τ	CVaR threshold variable
Φ	Set of probability distributions
$\boldsymbol{\omega}$	Loss distribution factors
$\mathcal{B}_\epsilon(\hat{P}_{\mathcal{D}})$	Wasserstein ball of radius ϵ around $\hat{P}_{\mathcal{D}}$
$\mathcal{D}$	Dataset of historical DART spreads
$\mathcal{K}$	Index set $\{1, 2\}$ for CVaR formulation
$\mathcal{L}$	Set of transmission lines
$\mathcal{N}$	Set of buses in the power network
$\mathcal{P}$	Set of probability distributions
$\mathcal{S}$	Uncertainty set for DART spreads
$\mathcal{T}$	Set of hours in a day $\{0, 1, \ldots, 23\}$
a_k	Coefficient in risk-averse formulation
b_k	Coefficient in risk-averse formulation
C	Calmar ratio

$c(\cdot)$	Generation cost function
$\mathbf{C}$	Matrix in cone uncertainty set formulation
d	Index for historical days in dataset
$\mathbf{d}$	Vector in cone uncertainty set formulation
d_W	Wasserstein distance
$\mathbf{g}$	Demand vector
$\mathbf{f}$	Loss coefficients vector
$f(\mathbf{q}, \mathbf{s})$	Loss function
$\mathbf{g}$	Generation vector
$\underline{\mathbf{g}}, \overline{\mathbf{g}}$	Generator lower and upper bounds
i	Index for buses
$\mathbf{I}$	Identity matrix
j	Index for bidding days
J	Total number of trading days
k	Index in CVaR formulation
L	Total energy bidding limit (MWh)
ℓ	Total system losses
ℓ^0	Base losses (constant term in loss function)
$\text{LMP}_{i,t}^{\text{DA}}$	Day-ahead locational marginal price
$\text{LMP}_{i,t}^{\text{RT}}$	Real-time locational marginal price
p	Order of the Wasserstein distance
$\mathbf{p}$	Load profile vector
$\mathbb{P}$	Probability distribution
$\hat{P}_{\mathcal{D}}$	Empirical probability distribution
$\mathbf{q}_t$	Bid quantity vector at time t (MWh)
$\mathbb{Q}$	Probability distribution
r_j	Daily return on day j
$R_{\text{annualized}}$	Annualized return
$\mathbf{s}_t$	DART spread vector at time t
$s_{i,t}$	DART spread at bus i and time t
$\mathbf{s}_t^d$	Historical DART spread for day d at time t
S	Sharpe ratio
$\overline{\mathbf{S}}$	Vector of transmission line flow limits
t	Time index (hour)
$\mathbf{T}$	Power Transfer Distribution Factors (PTDF matrix)
$\mathbf{U}$	Cone uncertainty set

v_0	Initial portfolio value
v_j	Portfolio value on day j
v_{peak}	Peak portfolio value before drawdown
v_{trough}	Trough portfolio value during drawdown
$\mathbb{W}$	Probability distribution in ambiguity set
$\mathbf{x}^d$	Auxiliary variable in DRO-CVaR reformulation
y^d	Excess loss variable in CVaR reformulation
$\mathbf{1}$	Vector of ones
$\|\cdot\|$	Norm (context-dependent)
$\|\cdot\|_1$	ℓ_1 norm (sum of absolute values)
$\|\cdot\|_2$	ℓ_2 norm (Euclidean norm)
$\|\cdot\|_*$	Dual norm

CHAPTER 1 INTRODUCTION

Modern power systems operate under increasingly complex and uncertain conditions, driven by factors such as the electrification of transport, the growing number of data centres, and the impacts of extreme weather events. System operators must maintain the balance between electricity supply and demand in real time, ensuring reliable grid operations within physical and security limits. This task has become more challenging due to variability in both demand and supply introduced by large-scale renewable integration, distributed generation, and dynamic consumption patterns. Maintaining this equilibrium requires mechanisms that allocate resources efficiently while preserving system security and reliability.

To support operational coordination and promote economic efficiency, liberalized power systems employ organized electricity markets that structure decision-making across multiple timescales. Among these, the day-ahead and real-time markets (DAM and RTM) play a central role in determining both the operational schedule and the financial settlement of energy transactions. Prices established in these markets, known as locational marginal prices (LMPs), reflect the marginal cost of delivering electricity at specific nodes, accounting for transmission constraints and network losses. While this market structure enables efficient short-term resource allocation, it is also subject to discrepancies between forecasted and realized conditions that can lead to significant price volatility across space and time.

Financial mechanisms such as virtual bidding have been developed to mitigate these inefficiencies by promoting price convergence between the day-ahead and real-time markets [61]. Virtual bidders act as financial participants that arbitrage expected price differences between the DAM and RTM (DART spreads). Virtual supply pushes down day-ahead prices when market participants expect that they will exceed real-time prices, while virtual demand boosts day-ahead prices when market participants expect the opposite, thus tightening the spread, and enhancing overall market liquidity [54, 65]. However, because these prices are affected by numerous uncertain and interdependent factors, designing bidding strategies that are both profitable and risk-aware remains a challenge.

In this context, distributionally robust optimization (DRO) provides a promising framework for addressing uncertainty in electricity markets [53]. By considering all probability distributions within an ambiguity set, DRO ensures robustness against errors in the underlying distribution, an advantage when historical data is limited, noisy, or when its distribution is unknown [36]. This approach leverages convex duality to reformulate worst-case expectation problems into tractable convex problems, enabling data-driven implementations [71].

This Master's thesis applies DRO to the virtual bidding problem in two-settlement electricity markets, introducing a data-driven approach that explicitly accounts for distributional uncertainty in DART spreads.

The remainder of this chapter introduces the fundamental concepts underpinning this work. Sections 1.1-1.4 describe the operational and financial structure of electricity markets, including the day-ahead and real-time markets, locational marginal pricing, and DART spreads. Section 1.5 discusses virtual bidding and its role in promoting price convergence and Section 1.6 then presents distributionally robust optimization and its relevance to electricity market applications. Finally, the main contributions, publications, and presentations associated with this research are summarized in Section 1.7.

1.1 Day-Ahead Market

We frame our work within North American wholesale electricity markets, where two-settlement systems are the standard design for energy trading. The day-ahead market (DAM) is the forward component of the two-settlement electricity market where energy transactions are scheduled one day prior to physical delivery. In this market, participants, i.e., generators, load-serving entities, and virtual traders, submit bids and offers based on forecasted load and generation for each hour of the following day, specifying the price and quantity at which they are willing to sell or purchase electricity. The independent system operator (ISO) or the regional transmission organization (RTO), e.g., the New York Independent System Operator (NYISO), Electric Reliability Council of Texas (ERCOT), and California Independent System Operator (CAISO), clears the market by solving optimization problems such as security-constrained unit commitment (SCUC) [14] and security-constrained economic dispatch (SCED) [12] to determine the least-cost generation schedule that respects system constraints such as transmission limits, generator capacities, and power flow equations. These decisions are made through automated algorithms and market rules codified in ISO manuals and tariffs [49, 50], with oversight from the Federal Energy Regulatory Commission (FERC). The result is a set of LMPs and scheduled quantities for each node in the network.

The DAM serves two primary purposes: (i) to provide a financially binding schedule that improves operational planning, and (ii) to signal expected real-time conditions through market prices. As noted in [33], the day-ahead market functions as a hedging mechanism against price volatility in the real-time market. Its design promotes system reliability and economic efficiency by allowing participants to lock in prices before actual delivery.

1.2 Real-Time Market

The real-time market, also known as the balancing market, operates close to the time of physical electricity delivery and serves as the mechanism through which ISOs ensure continuous system balance. While the DAM provides a forward schedule for generation and demand, real-time conditions often deviate from these forecasts due to factors such as unexpected demand changes, renewable output intermittency, or generator outages. The RTM enables ISOs to make real-time dispatch adjustments to maintain reliability and system security.

The RTM functions as a balancing mechanism that corrects for any discrepancies between scheduled and actual system conditions. ISOs monitor grid frequency, transmission flows, and generation availability in real time and issue dispatch instructions based on updated state estimations. The real-time market clearing process computes new LMPs that reflect the actual state of the system.

The RTM is therefore directly linked to the DAM through a two-settlement structure. Virtual financial settlements are based on discrepancies between day-ahead schedules and real-time outcomes: any difference in actual generation or consumption from the day-ahead commitment is settled at real-time prices. This mechanism ensures economic consistency between both markets and encourages accurate scheduling and forecasting.

1.3 Locational Marginal Pricing

A locational marginal price represents the cost of delivering one additional megawatt-hour (MWh) of electricity to a specific bus in the transmission network, considering both generation and congestion costs. LMPs are determined by solving a linearized optimal power flow (DC-OPF) problem that minimizes system cost subject to physical and operational constraints [40]. They consist of three main components: the marginal energy cost, the marginal congestion cost, and the marginal loss component.

ISOs can extract LMPs from a DC-OPF with power flow sensitivities called power transfer distribution factors (PTDFs) and line loss sensitivities called loss factors such as defined in [40] and reformulated in [20]. Let the power system be modelled as the graph $(\mathcal{N}, \mathcal{L})$ where $\mathcal{N} \subset \mathbb{N}$ denotes the set of buses, $\mathcal{L} \subset \mathcal{N} \times \mathcal{N}$ represents the set of transmission lines connecting these buses. Let $c : \mathbb{R}^{|\mathcal{N}|} \to \mathbb{R}$ be the cost function assumed to be monotonically increasing and the decision variables $\mathbf{g} \in \mathbb{R}^{|\mathcal{N}|}$ be the power generated by generator, and $\ell \in \mathbb{R}$ be the total losses of the system. The problem parameters are the power demand $\mathbf{e} \in \mathbb{R}^{|\mathcal{N}|}$, the loss coefficients $\ell^0 \in \mathbb{R}$ and $\mathbf{f} \in \mathbb{R}^{|\mathcal{N}|}$, the loss distribution factors $\boldsymbol{\omega} \in \mathbb{R}^{|\mathcal{N}|}$, the PTDF

matrix $\mathbf{T} \in \mathbb{R}^{|\mathcal{L}|\times|\mathcal{N}|}$, and the physical limits $\overline{\mathbf{S}} \in \mathbb{R}^{|\mathcal{L}|}$, and $\underline{\mathbf{g}}, \overline{\mathbf{g}} \in \mathbb{R}^{|\mathcal{N}|}$.

$$\begin{aligned}
\min_{\mathbf{g}\in\mathbb{R}^{|\mathcal{N}|},\ell\in\mathbb{R}} \quad & c(\mathbf{g}), && && \text{(1.1a)}\\
\text{s.t.} \quad & \mathbf{1}^\top(\mathbf{g}-\mathbf{e}) = \ell, && (\gamma), && \text{(1.1b)}\\
& \ell = \ell^0 + \mathbf{f}^\top(\mathbf{g}-\mathbf{e}), && (\sigma), && \text{(1.1c)}\\
& -\overline{\mathbf{S}} \le \mathbf{T}(\mathbf{g}-\mathbf{e}-\boldsymbol{\omega}\ell) \le \overline{\mathbf{S}}, && (\boldsymbol{\mu}_{\min}, \boldsymbol{\mu}_{\max}), && \text{(1.1d)}\\
& \underline{\mathbf{g}} \le \mathbf{g} \le \overline{\mathbf{g}}, && (\boldsymbol{\alpha}_{\min}, \boldsymbol{\alpha}_{\max}), && \text{(1.1e)}
\end{aligned}$$

where the dual variables associated to the primal formulation of the DC-OPF (1.1) are $\gamma \in \mathbb{R}$ from power balance constraint (1.1b), $\sigma \in \mathbb{R}$ from the loss component constraint (1.1c), $\boldsymbol{\mu}_{\min}, \boldsymbol{\mu}_{\max} \in \mathbb{R}^{|\mathcal{L}|}$ from the transmission limits constraint (1.1d), and $\boldsymbol{\alpha}_{\min}, \boldsymbol{\alpha}_{\max} \in \mathbb{R}^{|\mathcal{N}|}$ from the generators capacity constraint (1.1e). The dual formulation of (1.1), from which it is possible to retrieve the LMPs, is:

$$\begin{aligned}
\max_{\gamma,\sigma,\boldsymbol{\alpha}_{\min},\boldsymbol{\alpha}_{\max},\boldsymbol{\mu}_{\min},\boldsymbol{\mu}_{\max}} \quad & \gamma\mathbf{1}^\top\mathbf{e} + \sigma(\ell^0 - \mathbf{f}^\top\mathbf{e}) + \boldsymbol{\mu}_{\max}{}^\top(\overline{\mathbf{S}} + \mathbf{T}\mathbf{e}) \\
& \quad + \boldsymbol{\mu}_{\min}{}^\top(\overline{\mathbf{S}} - \mathbf{T}\mathbf{e}) + \boldsymbol{\alpha}_{\min}{}^\top\underline{\mathbf{g}} - \boldsymbol{\alpha}_{\max}{}^\top\overline{\mathbf{g}} && \text{(1.2a)}\\
\text{s.t.} \quad & \gamma\mathbf{1} + \sigma\mathbf{f} + \boldsymbol{\mu}_{\max}{}^\top\mathbf{T} - \boldsymbol{\mu}_{\min}{}^\top\mathbf{T} + \boldsymbol{\alpha}_{\min} - \boldsymbol{\alpha}_{\max} = c && \text{(1.2b)}\\
& \gamma + \sigma - (\boldsymbol{\mu}_{\max} - \boldsymbol{\mu}_{\min})^\top\mathbf{T}\boldsymbol{\omega} = 0 && \text{(1.2c)}\\
& \boldsymbol{\mu}_{\min}, \boldsymbol{\mu}_{\max}, \boldsymbol{\alpha}_{\min}, \boldsymbol{\alpha}_{\max} \ge 0 && \text{(1.2d)}
\end{aligned}$$

The LMPs are decomposed in three components: the marginal costs of energy ($\boldsymbol{\gamma}^{\mathrm{E}}$), of losses ($\boldsymbol{\gamma}^{\mathrm{L}}$), and of congestion ($\boldsymbol{\gamma}^{\mathrm{C}}$), defined as:

- $\boldsymbol{\gamma}^{\mathrm{E}} := \gamma\mathbf{1}$;
- $\boldsymbol{\gamma}^{\mathrm{L}} := \sigma\mathbf{f}$;
- $\boldsymbol{\gamma}^{\mathrm{C}} := \boldsymbol{\mu}_{\max}{}^\top\mathbf{T} - \boldsymbol{\mu}_{\min}{}^\top\mathbf{T}$.

where $\mathbf{1} \in \mathbb{R}^{|\mathcal{N}|}$ is the all-one vector. For node i and time t, the day-ahead LMP and real-time LMP can be expressed as:

$$\begin{aligned}
\mathrm{LMP}^{\mathrm{DA}}_{i,t} &= \gamma^{\mathrm{E,\ DA}}_{i,t} + \gamma^{\mathrm{L,\ DA}}_{i,t} + \gamma^{\mathrm{C,\ DA}}_{i,t},\\
\mathrm{LMP}^{\mathrm{RT}}_{i,t} &= \gamma^{\mathrm{E,\ RT}}_{i,t} + \gamma^{\mathrm{L,\ RT}}_{i,t} + \gamma^{\mathrm{C,\ RT}}_{i,t}.
\end{aligned}$$

where the superscripts DA and RT represent when (1.1) is solved. Differences in LMPs across network nodes signal transmission congestion or losses and therefore drive price differentiation within the system. These differences provide incentives for efficient generation decisions at each location. Both the DAM and RTM produce nodal LMPs, but are calculated at a different time, with different information available. This creates a variation between these two price sets that defines the day-ahead − real-time (DART) spread, which is exactly what virtual bidders are trying to arbitrage.

1.4 DART Spreads

The DART spread is defined as the difference between the day-ahead and real-time LMPs at a given node and hour:

$$s_{i,t} = \text{LMP}_{i,t}^{\text{RT}} - \text{LMP}_{i,t}^{\text{DA}} \tag{1.3}$$

where $s_{i,t}$ denotes the spread at node i of the network and time t. A positive DART spread indicates that the real-time price exceeds the day-ahead price, whereas a negative spread indicates the opposite. DART spreads are inherently stochastic and influenced by many factors including forecast errors, transmission congestion, renewable intermittency, and strategic participant behaviour [28, 41, 47].

For speculative virtual bidders, DART spreads represent profit opportunities: buying virtually in the DAM and selling in the RTM when $s_{i,t} > 0$, or conversely selling in the DAM and buying in the RTM when $s_{i,t} < 0$. However, the high volatility and non-stationary nature of spreads make prediction and risk management challenging, motivating the use of optimization methods such as stochastic or distributionally robust optimization for virtual bidding.

1.5 Virtual Bidding

Virtual bidding, also called convergence bidding [61], allows market participants to arbitrage price discrepancies between the DAM and RTM without physically generating or consuming electricity. Participants submit purely financial bids, virtual supply or demand, at specific nodes and hours. Even though the product is purely financial, ISOs treat the virtual supply or load as an actual MWh injection or withdrawal at the specified nodes when solving the OPF. If the bid clears in the DAM, it is automatically offset in the RTM, producing a profit (or loss) proportional to the realized DART spread [33, 34].

By exploiting price differentials, virtual bidders help align day-ahead and real-time prices, promoting market efficiency [31]. Yet, the design of virtual bidding is not without challenges. Studies have shown that, under certain conditions, these financial positions can distort dispatch outcomes or even increase system costs [54], and that strategic behaviour under uncertainty may undermine social welfare [65]. Because virtual bidders are motivated solely by profit rather than system reliability, such actions can exacerbate operational challenges. These concerns highlight that while virtual bidding can enhance price convergence, it also introduces risks that call for prudent strategies and strong market regulations [11].

In U.S. markets such as the Midcontinent Independent System Operator (MISO), PJM Interconnection LLC (PJM), ISO New England (ISO-NE) [32], Southwest Power Pool (SPP), CAISO, NYISO and ERCOT illustrated on Figure 1.1, virtual bidding is a well-established mechanism for financial participants. In contrast, most European markets, including the OMI Group, responsible for the Iberian Peninsula [51], and the National Energy System Operator (NESO) in the United Kingdom [24], mainly operate through day-ahead and intra-day markets, with no equivalent real-time market like in the United States. While intra-day markets in Europe allow some arbitrage opportunities, they do not directly allow for virtual bidding strategies like the DART spread does due to the absence of a real-time market for price adjustments. As such, virtual bidding, as formulated in this work, is most applicable to markets with both day-ahead and real-time market structures, such as those found in the United States.

1.6 Distributionally Robust Optimization

Distributionally robust optimization (DRO) is a mathematical programming framework for decision-making under uncertainty that bridges the gap between stochastic and robust optimization [36]. Instead of assuming a known probability distribution, DRO defines an ambiguity set provided by a family of plausible distributions around the empirical data, and optimizes for the worst-case expectation within that set. This provides protection against distributional uncertainty while avoiding the over-conservatism of classic robust optimization [35, 45].

A common way to define the ambiguity set is through the Wasserstein distance, which measures the difference between probability distributions in terms of the minimal "transport cost" required to transform one distribution into another, as illustrated in Figure 1.2.

Formally, let $\mathcal{S} \subseteq \mathbb{R}^n, n \in \mathbb{N}$, be the support of the random vectors and let $\mathcal{P}(\mathcal{S})$ be the set of all probability distributions on $\mathcal{S}$. For two probability distributions $\mathbb{P}, \mathbb{Q} \in \mathcal{P}(\mathcal{S})$, the

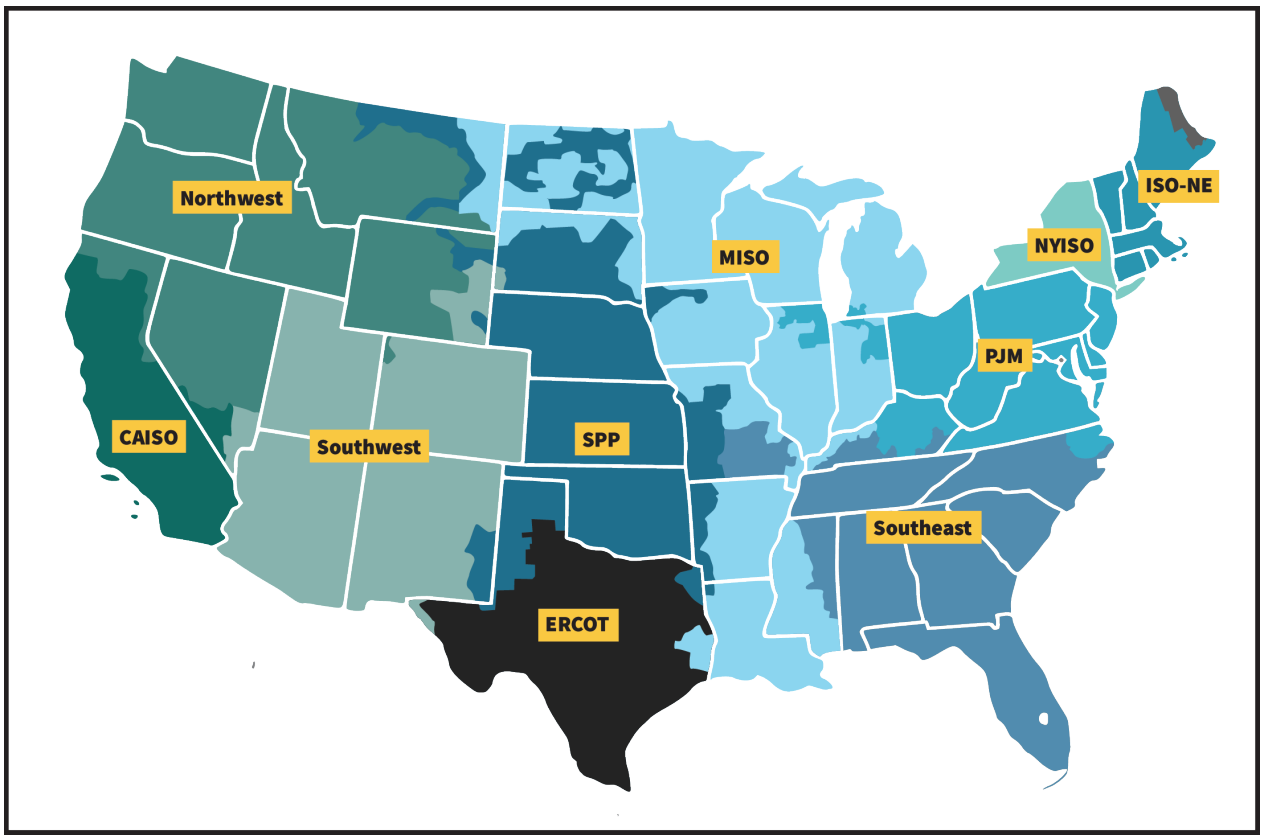

Figure 1.1 Map of U.S. ISOs and RTOs. Source: FERC [27].

Wasserstein distance of order $p \geq 1$ is defined as:

$$d_{\mathrm{W}}(\mathbb{P}, \mathbb{Q}) = \left(\inf_{\pi \in \Pi(\mathbb{P}, \mathbb{Q})} \int_{\mathcal{S} \times \mathcal{S}} \|\mathbf{s}_1 - \mathbf{s}_2\|^p \, \mathrm{d}\pi(\mathbf{s}_1, \mathbf{s}_2) \right)^{1/p}, \tag{1.4}$$

where $\Pi(\mathbb{P}, \mathbb{Q})$ denotes the set of all joint distributions on $\mathcal{S} \times \mathcal{S}$ with marginals $\mathbb{P}$ and $\mathbb{Q}$. Intuitively, the Wasserstein distance represents the minimum amount of "mass transport" required to morph one probability distribution into another [45, 52].

By constructing a Wasserstein ball of radius $\epsilon > 0$ centred around the empirical distribution $\hat{\mathbb{P}}_{\mathcal{D}}$, defined as

$$\mathcal{B}_\epsilon(\hat{\mathbb{P}}_{\mathcal{D}}) = \left\{ \mathbb{Q} \in \mathcal{P}(\mathcal{S}) \,\middle|\, d_{\mathrm{W}}(\mathbb{Q}, \hat{\mathbb{P}}_{\mathcal{D}}) \leq \epsilon \right\}, \tag{1.5}$$

where the empirical distribution $\hat{\mathbb{P}}_{\mathcal{D}}$ is given by

$$\hat{\mathbb{P}}_{\mathcal{D}} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \delta_{\mathbf{s}^d}, \tag{1.6}$$

with $\delta_{\mathbf{s}^d}$ denoting the Dirac measure at the observed sample $\mathbf{s}^d$, and $\mathcal{D}$ representing the

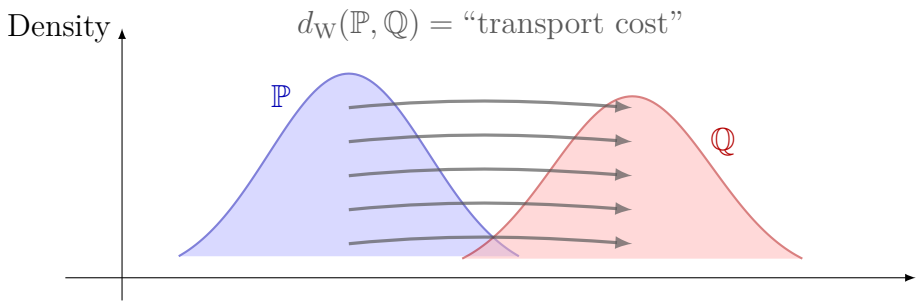

Figure 1.2 Illustration of the Wasserstein distance $d_{\mathrm{W}}(\mathbb{P},\mathbb{Q})$ between the probability distributions $\mathbb{P}$ and $\mathbb{Q}$, represented as the minimal "mass transport" required to move one distribution into the other.

dataset of historical observations. A larger ϵ implies greater protection against distributional shifts at the cost of solution conservatism, while a smaller ϵ yields solutions closer to the empirical mean [45].

Building on this construction, the general DRO problem seeks decisions that minimize the worst-case expected cost over all distributions in the ambiguity set [59]. Formally, let $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^n$ denote the decision vector constrained to the set $\mathcal{X}$ and $h(\mathbf{x},\mathbf{s})$ the cost function. The DRO model is:

$$\inf_{\mathbf{x}\in\mathcal{X}} \sup_{\mathbb{Q}\in\mathcal{B}_\epsilon(\hat{\mathbb{P}}_{\mathcal{D}})} \mathbb{E}[h(\mathbf{x},\mathbf{s})]. \tag{1.7}$$

In the context of virtual bidding, DRO is particularly appealing because electricity market prices exhibit heavy tails, regime shifts, and structural uncertainties. Traditional stochastic models that assume a fixed distribution often fail to capture these dynamics, resulting in unstable strategies. By contrast, a DRO framework builds its ambiguity set around the empirical distribution of observed data, enabling a data-driven approach that accounts for distributional uncertainty. This allows market participants to design risk-aware and more reliable bidding strategies under volatile conditions [13, 18, 55].

1.7 Contributions

This work proposes an uncertainty- and risk-aware virtual bidding methodology to account for two-settlement electricity market uncertainty and volatility. Our approach is data-driven and accounts for the lack of generally available system models. Specifically, we formulate the virtual bidding problem as a DRO with a Wasserstein distance-based ambiguity set [45] and provide a tractable reformulation that can be efficiently solved. The strategy is obtained by

identifying the optimal parameters of the model over a 12-month training set, and is implemented in an online fashion, updating daily with new market data to remain responsive to evolving conditions. The effectiveness of our approach is shown by using historical data from NYISO on an 8-month testing set, showing that DRO-based strategies outperform traditional methods, namely, the equally weighted selling portfolio and scenario-based stochastic optimization approaches, especially in instances with high market volatility and uncertainty. The specific contributions of our work are as follows:

- We formulate a distributionally robust optimization model that explicitly accounts for market uncertainty and volatility.
- We propose an online bidding strategy with an integrated pipeline for efficiently tuning model parameters, making the strategy adaptable to changing market and network conditions.

While similar studies using NYISO data exist, our case study provides new insights by comparing our bidding strategy against other uncertainty-aware models and existing benchmarks, highlighting the advantages in terms of both profitability and risk mitigation. This validation based on cumulative profits, profits per invested MWh, Sharpe ratio, and Calmar ratio, demonstrates the practical applicability and effectiveness of our proposed approach.

1.7.1 Publications

The main project presented in this Master's thesis is published as an article in the *Sustainable Energy, Grids and Networks* journal [5]:

- **Xavier Audet**, Kliti Qako, and Antoine Lesage-Landry, "A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets," Sustainable Energy, Grids and Networks, p. 101904, 2025.

1.7.2 Presentations

The following presentations have showcased this research at conferences and events:

- *A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets*, 12th Bulk Power System Dynamics and Control Symposium IREP'2025, June 2025, Sorrento, Italy.
- *A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets*, Optimization Days, May 2025, Montréal, Québec, Canada.

- *A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets*, 5e Journée des étudiant(e)s du GERAD, February 2025, Montréal, Québec, Canada.

1.8 Organization of the Master's Thesis

The remainder of this Master's thesis is organized as follows. Chapter 2 presents a comprehensive review of the literature relevant to electricity market design, virtual bidding strategies, and decision-making under uncertainty. This chapter establishes the theoretical and practical context for the research and identifies the methodological gap addressed by this work. Chapter 3 formulates the virtual bidding problem within a DRO framework. It introduces the mathematical model, explains the integration of the CVaR risk measure, and provides a tractable reformulation that enables efficient computations. Chapter 4 develops an online bidding strategy based on the proposed model. This chapter details the data selection process, the rolling-window approach for adapting to changing market conditions, and the hyperparameter tuning procedure designed to optimize risk-adjusted performance. Chapter 5 reports the numerical experiments conducted on historical NYISO data. It compares the proposed strategy against benchmark models, analyzes daily and cumulative profits, and evaluates performance using risk-adjusted metrics such as the Sharpe and Calmar ratios. Finally, Chapter 6 summarizes the main contributions of the Master's thesis, discusses limitations, and outlines promising directions for future research, including extensions to other markets and congestion hedging instruments.

CHAPTER 2 LITERATURE REVIEW

This chapter surveys the literature surrounding three main aspects: electricity market design, virtual bidding strategies, and distributionally robust optimization. Section 2.1 reviews the foundational structure of two-settlement electricity markets and the economic principles underlying locational marginal pricing, establishing the context in which virtual bidders operate. Section 2.2 then examines existing virtual bidding strategies, with particular emphasis on forecasting methodologies, machine learning approaches, and the empirical behaviour of market participants. Finally, Section 2.3 focuses on the rapidly growing body of work on DRO in power systems, emphasizing its applications to unit commitment, optimal power flow, renewable generation scheduling, and bidding strategies for physical market participants. Finally, Section 2.4 summarizes the main insights from the literature and identifies the research gap addressed in this Master's thesis.

2.1 Market Design

The design of electricity markets has been the subject of research aimed at balancing economic efficiency, reliability, and transparency in power system operations. Since the restructuring of electricity industries in the 1990s, following policy reforms such as FERC Orders No. 888 and 2000 [25, 26], a large body of literature has examined the theoretical and practical foundations of competitive market design. These studies have primarily focused on how to coordinate decentralized decision-making through market mechanisms while respecting the physical and operational constraints of the grid.

A central theme in the literature is the structure of two-settlement systems, where the day-ahead and real-time markets operate in tandem. The work of [33] established the multi-settlement framework as a way to achieve consistency between forward and real-time prices, mitigate market power, and ensure operational feasibility. Subsequent analyses, such as those in [31], examined how this design enables efficient hedging of short-term uncertainty and supports price convergence between settlements. Other authors have also explored the role of day-ahead commitments in improving system reliability and providing signals for resource scheduling [19].

The choice of pricing mechanism, particularly the use of LMPs, is another cornerstone of market design literature. Studies such as [39] formalized LMPs as the marginal value of power injections at each node, derived from the dual variables of the OPF problem. This approach

ensures that prices internalize congestion and losses while preserving system security. Other works, including [17], emphasized how nodal and temporal price signals promote efficient dispatch and guide long-term investment decisions by revealing the true economic cost of network constraints.

Recent research has increasingly focused on the integration of stochastic and risk-based perspectives into market design. For instance, [41] reviews how uncertainty and volatility challenge pricing structures, motivating new approaches to risk management and market efficiency. Similarly, [28] and [47] underscore the role of forecast errors and extreme events in shaping market outcomes, which has led to the growing interest of incorporating probabilistic models, forecasting tools, and robust optimization techniques into market design.

From a system operator's perspective, the literature recognizes that market design and system operations are inherently coupled. As shown in [19], optimization-based market clearing models, such as the SCUC and SCED, ensure consistency between physical feasibility and market outcomes. The relationship between operational constraints and market prices, where LMPs are shadow prices of grid constraints, as illustrated in Section 1.3, has been widely analyzed as the mechanism linking power system engineering and economic principles [39].

The literature highlights several ongoing challenges in electricity market design. These include addressing the increasing variability from renewable energy, mitigating strategic bidding behaviour, and improving coordination between short-term markets and long-term planning [11, 17, 41]. Moreover, while the theoretical basis of two-settlement systems is well established, their empirical performance, particularly in capturing uncertainty and promoting efficient price formation, remains an active research topic, motivating the exploration of methods such as DRO in market design and bidding strategies. While this work does not propose changes to market design, it is inherently shaped by its structure. In particular, we focus on virtual bidding within the existing two-settlement framework, as discussed in the following section.

2.2 Virtual Bidding Strategies

Virtual bidding, also known as convergence bidding [61], has attracted significant attention in the literature due to its role in linking the day-ahead and real-time markets through arbitrage and price convergence. By allowing financial participants to trade on expected price differentials without owning physical assets, virtual bidding provides liquidity and promotes efficient price formation [31]. However, it also introduces strategic and operational challenges associated with uncertainty, forecast errors, and risk exposure [11].

Early contributions to the study of virtual bidding focused on its economic and policy implications. The work of [11] provides one of the first comprehensive analyses of the benefits and drawbacks of virtual bidding, emphasizing how it enhances price convergence between the day-ahead and real-time markets but can also amplify volatility if not properly monitored. In a subsequent analysis, [31] explores the design of regulatory frameworks that ensure market efficiency while limiting manipulation or congestion gaming. These studies highlight that while virtual bidding contributes to efficient price clearing, it also creates incentives for speculative behaviour. These studies motivate research into quantitative approaches for strategy formulation and risk management.

In this direction, a few studies have formulated a virtual bidding strategy as an optimization problem. The work in [43] employs stochastic programming to model price uncertainty. They proposed a two-stage stochastic optimization model that enables physical market participants with virtual bidding capacity to jointly optimize generation and virtual trading profits. Their framework considers the coupling between physical and financial positions and highlights the importance of probabilistic modelling in managing exposure to real-time deviations. Also, [34] proposes four alternative models for clearing two-settlement markets under uncertainty: a multi-player stochastic equilibrium, ISO co-optimization using stochastic unit commitment, a sequential equilibrium based on deterministic day-ahead scheduling with real-time rebalancing, and an extension of the sequential model with self-scheduling slow-start generators. Their findings show that errors from deterministic day-ahead scheduling can be mitigated by virtual bidders with self-scheduling by slowstart generators, reinforcing the role of virtual bidding as a mechanism for improving price convergence and operational feasibility.

The integration of machine learning and data-driven methods into virtual bidding has also been explored. In [61], the authors propose a data-driven convergence bidding strategy based on reverse engineering of real-world market behaviour using historical bidding and pricing data. Applied to the CAISO market, their approach first focuses on maximizing net profit by capturing price spikes through an optimization model that determines bid prices capable of exploiting rare but highly profitable deviations in day-ahead prices. It then introduces a dynamic node labelling process that classifies each node as suitable for demand bids, supply bids, both, or neither, based on its historical profitability potential. Finally, the strategy combines these results with insights from reverse-engineered bidding patterns to select the most appropriate method for each node and hour. Similarly, [6] introduced an online learning algorithm for virtual trading in electricity markets. Their approach models the bidding process as a sequential decision problem under uncertainty, where strategies are updated dynamically based on observed day-ahead and real-time prices. The framework employs regret-minimization algorithms to ensure that cumulative performance approaches

that of an ideal strategy with perfect foresight. In addition to maximizing cumulative profits, the authors incorporate risk-adjusted metrics such as the Sharpe ratio to evaluate robustness against volatility. This adaptive learning process demonstrates that historical electricity price data contain exploitable patterns, enabling profitable virtual bidding decisions. Finally, [38] leverages deep learning to forecast DART spreads and integrates price-sensitivity models into portfolio optimization, outperforming models that ignore such sensitivity.

The stochastic nature of day-ahead and real-time prices has also motivated the development of forecasting models for DART spreads. The work of [28] focuses on predicting electricity price spikes through a supervised learning framework that incorporates forward-looking, backward-looking, and seasonal features. Their approach employs four models, namely, logistic regression, random forests, gradient boosting trees, and deep neural networks to anticipate extreme deviations, allowing traders or decision-makers to mitigate exposure to adverse spread events.

Advances in deep learning have further refined price spread modelling by capturing temporal dependencies. For instance, the study of [37] developed a machine-learning-driven virtual bidding framework integrating neural networks and gradient boosting to forecast DART spreads. Their use of recurrent neural networks (RNNs), particularly long short-term memory (LSTM) models, enables dynamic adaptation to inter-hour dependencies, outperforming traditional regression methods. Moreover, their modelling of price sensitivity to net virtual bids through constrained gradient boosting trees represents a step toward explicitly accounting for the market impact of virtual traders.

2.3 DRO in Power Systems

DRO has emerged as a compelling approach to tackle electricity market problems due to its ability to balance robustness and performance by considering an ambiguity set of probability distributions rather than presuming access to a single known distribution. This is particularly relevant in electricity markets, where the joint distribution of uncertain variables such as loads, renewable outputs, and price spreads is difficult, if not impossible, to characterize. DRO methods construct ambiguity sets consistent with observed data to provide protection against distributional error while leveraging available historical information.

The theoretical foundations of modern DRO have been significantly advanced in recent years. Contributions such as [35, 36, 45] establish tractable reformulations for ambiguity sets based on ϕ-divergences, moments, and Wasserstein distances. These works have enabled DRO to become computationally viable for large-scale optimization problems by showing how worst-

case expectations can be expressed through convex reformulations.

In power systems, DRO has been applied across a broad range of operational and market-facing problems where uncertainty in loads, renewable injections, and prices is consequential. For short-term scheduling and commitment in small-scale systems, [70] develops a Kullback–Leibler divergence–based ambiguity set and a two-level decomposition method to derive a distributionally robust unit commitment formulation for microgrids, demonstrating improved cost–risk tradeoffs relative to conventional stochastic approaches. For optimal power flow and chance-constrained dispatches, [69] provides a conic reformulation of the distributionally robust chance-constrained OPF that leverages moment information to obtain tractable convex programs which guard against worst-case distributional deviations. In a similar way, [3] derives exact reformulations for chance-constrained energy–reserve dispatch under Wasserstein ambiguity sets.

In the context of economic dispatch with high renewable penetration, [58] proposes a Wasserstein-based DRO formulation for look-ahead economic dispatch that explicitly accounts for temporal coupling and operational constraints, improving robustness to distributional shifts in renewable forecasts. In electricity markets, [10] uses DRO in a transactive energy framework for multi-microgrid energy management, using data-driven ambiguity sets to balance market interactions and uncertainty. Similar integrated approaches have been proposed for virtual power plant (VPP) bidding strategies, where DRO and CVaR jointly address price uncertainty and tail risk [67]. For market participation and bidding strategies, [30] develops a bi-level distributionally robust bidding framework for a wind–storage aggregator, incorporating uncertainty in both wind generation and market prices via an ambiguity set informed by historical forecast errors. Their findings indicate that by optimizing against the worst-case distribution within this set, the aggregator significantly reduces downside risk, particularly during adverse price or generation outcomes. The model achieves this through an ARMA-based forecast correction combined with a DRO approach, which effectively tailors bids to better withstand extreme events and market fluctuations. Finally, [55] investigates DRO-based trading strategies for renewable energy producers, demonstrating that distributional ambiguity plays a central role in shaping realistic profit–risk trade-offs. The author argues that traditional robust optimization is often overly conservative in this context, while classical stochastic models depend heavily on the presumed accuracy of the forecast distribution. DRO, by contrast, provides a middle ground that is both data-driven and risk-aware, making it a natural fit for markets with increasingly stochastic and intermittent supply.

Despite this growing literature, the application of DRO to electricity market decision problems for speculative market participants remains limited. In particular, virtual bidding has

received little attention in the DRO community. Virtual bidding strategies in the existing literature typically rely on point forecasts, simple regression models, or classical stochastic approaches to model DART spreads. To the best of the authors' knowledge, before our work [5] was published, no published work incorporates distributionally robust optimization into virtual bidding.

2.4 Literature Summary

While prior work has explored stochastic optimization, bi-level formulations, and machine learning for virtual bidding, these approaches often rely on fixed probabilistic models or static strategies, which can struggle under the high volatility and structural uncertainty of electricity markets [34, 43]. Data-driven methods have shown promise, from reverse-engineering real-world bidding patterns [61] to online learning algorithms that adapt sequentially to market feedback [6], and deep learning models integrated with price-sensitive portfolio optimization [38]. These studies highlight the importance of leveraging historical data and dynamic adaptation to evolving conditions. In parallel, DRO has emerged as a powerful tool for decision-making under uncertainty, offering protection against uncertain distributions without the conservatism of traditional robust models [30]. In this work, we address the challenge of uncertainty and volatility in virtual bidding by formulating the problem within a DRO framework using CVaR. Our approach is data-driven, constructing ambiguity sets around empirical DART spread distributions, and incorporates an online adaptation mechanism that adjusts to market dynamics daily.

CHAPTER 3 PROBLEM FORMULATION

Building on the insights from the literature, we propose a DRO framework combined with CVaR to design a risk-aware virtual bidding strategy. This formulation explicitly accounts for uncertainty in DART spreads and tail risk, providing a tractable model that can adapt to evolving market conditions. Virtual bidding allows market participants to submit financial bids without physically delivering or consuming electric energy. In a two-settlement electricity market, which consists of a day-ahead and a real-time market, virtual bids aim to capitalize on the price differences between these two markets [33]. Participants place bids in the day-ahead market as if they were buyers (loads) or sellers (generators) and then settle their positions in the real-time market. In doing so, virtual bidders help align the day-ahead prices with real-time ones, thereby improving market efficiency and reducing price volatility.

Virtual bids are cleared if the buy price is greater than or equal to the price set by the grid operator in the day-ahead market or if the sell price is less than or equal to the day-ahead market price. Unlike most U.S. ISO/RTO markets, where virtual bids are submitted at the nodal level (i.e., specific buses), NYISO [48] operates virtual bidding at the zonal level. In CAISO [9], MISO [44], PJM [56], ISO-NE [32], SPP [63], and ERCOT [21] participants must specify both a price and a quantity for a particular bus in the grid for each hour of the day. By contrast, NYISO aggregates its network into eleven geographical load zones (A–K), and virtual bids are placed on these zones rather than individual nodes. This zonal structure simplifies bidding but reduces granularity compared to nodal markets. Our case study focuses on NYISO because its zonal bidding design reduces computational complexity while still reflecting the uncertainty and volatility challenges of U.S. electricity markets. Figure 3.1 illustrates the NYISO load zones used for virtual bidding.

To avoid the computational burden of mixed-integer programming, we adopt a price-taker approach, where we always submit a high price for buying and a low price for selling. This approach ensures that our bids clear every time and allows us to focus on optimizing the energy quantity rather than the bid price itself. Because our strategy does not seek to influence market prices but rather takes them as given, this price-taker assumption is reasonable and effective for simplifying the bidding process. Previous studies have also adopted a similar approach to reduce computational complexity in energy markets [61].

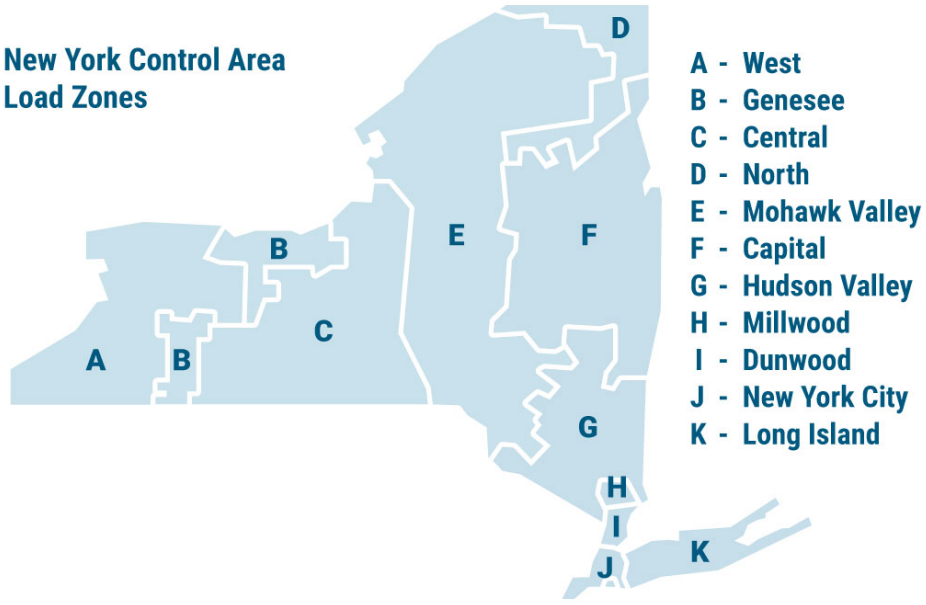

Figure 3.1 Map of load zones in NYISO. Source: U.S. Energy Information Administration [66].

3.1 Virtual Bidding Modelling

Let the power system be modelled as the graph $(\mathcal{N}, \mathcal{L})$ where $\mathcal{N} \subset \mathbb{N}$ denotes the set of buses and $\mathcal{L} \subset \mathcal{N} \times \mathcal{N}$ represents the set of transmission lines connecting these buses. The quantity vector $\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}$ collects the bid energy quantities (in MWh) submitted at all buses at a specific hour $t \in \mathcal{T} = \{0, 1, ..., 23\}$ of the considered day. To ensure operational feasibility, we impose a limit $L > 0$ on the total energy that can be bid across all buses in each hour, expressed as:

$$\|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}.$$

We suppose that the limit L ensures that our bids do not affect prices in the considered market. The impact of virtual bids on prices (price-makers) is a topic for future work. It is explored, for example, in [37].

Let $\mathcal{S} \subset \mathbb{R}^{|\mathcal{N}|}$ be the uncertainty set of the random variable $\mathbf{s}_t \in \mathcal{S}$ representing the DART spreads, for all buses at time t. We model the spread $\mathbf{s}_t$ as a random variable with distribution $\mathbb{P} \in \mathcal{P}$, where $\mathcal{P}$ is the set of probability distribution on $\mathcal{S}$. The virtual bidding problem

takes the following form:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}} \quad \mathbb{E}^{\mathbb{P}} \left[\sum_{t \in \mathcal{T}} -\mathbf{s}_t^\top \mathbf{q}_t \right], \tag{3.1a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}, \tag{3.1b}$$

where the expectation is taken with respect to $\mathbf{s}_t$.

Given the complexity of the market, i.e., its dependency on many factors such as the weather, the availability of energy, and consumer behaviour, the distribution $\mathbb{P}$ is unknown. In this work, we are interested in a data-driven approach. Using a set $\mathcal{D}$ of historical DART spreads $\mathbf{s}_t^d$, the distribution $\mathbb{P}$ can be approximated with the discrete empirical probability distribution:

$$\hat{\mathbb{P}}_{\mathcal{D}} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \delta_{\mathbf{s}_t^d},$$

where $\delta_{\mathbf{s}_t^d}$ is the Dirac measure giving a probability mass $\frac{1}{|\mathcal{D}|}$ to each observed data point $\mathbf{s}_t^d \in \mathcal{D}$. This expression reflects our estimation of the distribution based on the observed data. This enables us, for example, to approximate (3.1) as:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}} \quad \mathbb{E}^{\hat{\mathbb{P}}_{\mathcal{D}}} \left[-\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t \right] := -\frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \left[\sum_{t \in \mathcal{T}} \mathbf{s}_t^{d\top} \mathbf{q}_t \right], \tag{3.2a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}, \tag{3.2b}$$

which we later use as a benchmark in our numerical study. We refer to the scenario-based, single-stage stochastic optimization problem (3.2) as the stochastic optimization (`SO`) model.

3.2 Volatility- and Uncertainty-Aware Bidding Model

Volatility in electricity markets refers to the significant fluctuations in prices over time, which introduce risks for market participants [47]. Concurrently, variability in market conditions encompasses broader uncertainties, such as changes in generation, demand, or external factors affecting the distribution of prices [41]. In our approach, we address both challenges using a combination of CVaR and DRO, respectively. CVaR focuses on tail risks, providing protection against extreme price movements [68]. By minimizing potential losses in the worst-case scenarios, CVaR helps manage the impact of sharp price drops. In parallel, DRO considers a

range of possible market conditions by defining an ambiguity set, which captures variability in the market's underlying distribution. This ensures that our virtual bidding strategy remains robust not only to known price risks but also to shifts and uncertainties in future market conditions.

3.2.1 Volatility Mitigation

CVaR, or expected shortfall, is a risk measure commonly used to evaluate investment strategies [64] that captures the tail risk beyond a given confidence level $\alpha \in (0, 1]$. The CVaR is defined as the average of the worst-case outcomes beyond a specified confidence level [60], e.g., the average return of the worst 10%. To account for the volatility in our model, we integrate the CVaR in the objective function of the single-stage stochastic problem (3.2).

Following [60], let $f(\mathbf{q}, \mathbf{s})$ denote the loss associated with a decision vector $\mathbf{q} \in \mathbb{R}^{|\mathcal{N}|}$ and a random vector $\mathbf{s}$ defined on $\mathcal{S}$ with probability measure $\mathbb{P} \in \mathcal{P}$ and corresponding density $p(\mathbf{s})$. The CVaR at confidence level $\alpha \in (0, 1]$ can be defined formally as an integral over the tail of the loss distribution:

$$\mathrm{CVaR}_\alpha(\mathbf{q}) = \frac{1}{1-\alpha} \int_{\{\mathbf{s} | f(\mathbf{q},\mathbf{s}) \geq \mathrm{VaR}_\alpha(\mathbf{q})\}} f(\mathbf{q}, \mathbf{s})\, p(\mathbf{s})\, \mathrm{d}\mathbf{s}, \tag{3.3}$$

where $\mathrm{VaR}_\alpha(\mathbf{q}) = \min\{\beta \in \mathbb{R} \mid \Pr_{\mathbb{P}}(f(\mathbf{q}, \mathbf{s}) \leq \beta) \geq \alpha\}$ is the value-at-risk. This integral is equivalent to taking the conditional expectation over the tail, which allows us to write CVaR in the more familiar expectation-based form:

$$\mathrm{CVaR}_\alpha(\mathbf{q}) = \frac{1}{1-\alpha}\, \mathbb{E}^{\mathbb{P}}\Big[f(\mathbf{q}, \mathbf{s}) \,\Big|\, f(\mathbf{q}, \mathbf{s}) \geq \mathrm{VaR}_\alpha(\mathbf{q}) \Big]. \tag{3.4}$$

Finally, [60] shows that CVaR can equivalently be reformulated as a minimization problem over a threshold $\tau \in \mathbb{R}$, which is particularly useful for optimization since it yields a convex formulation:

$$\mathrm{CVaR}_\alpha(\mathbf{q}) = \min_{\tau \in \mathbb{R}} \; \tau + \mathbb{E}^{\mathbb{P}} \left[\frac{1}{\alpha} \max \left\{ -\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t - \tau,\; 0 \right\} \right], \tag{3.5}$$

where τ acts as a quantile of the loss distribution beyond which tail risk is considered. The resulting formulation minimizes a combination of the expected loss and the CVaR of the loss. The reformulation in (3.5) is particularly useful because the objective is convex, which makes it useful in practice. This property makes CVaR optimization computationally tractable and suitable for large-scale problems. Let $\alpha \in (0, 1]$ be the confidence level used for the CVaR, and representing the portion of tail risk considered. Let $\rho \in [0, 1]$ be the risk-aversion factor

which we use to balance between minimizing the expected loss and the CVaR. The problem formulation becomes: :

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}} \quad \rho \mathbb{E}^{\mathbb{P}} \left[-\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t \right] + (1-\rho) \mathrm{CVaR}_\alpha \left[-\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t \right], \tag{3.6a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall\, t \in \mathcal{T}. \tag{3.6b}$$

Substituting in the CVaR definition (3.5), we obtain:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}} \quad \rho \mathbb{E}^{\mathbb{P}} \left[-\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t \right] + (1-\rho) \min_{\tau \in \mathbb{R}} \mathbb{E}^{\mathbb{P}} \left[\tau + \frac{1}{\alpha} \max \left\{ -\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t - \tau,\, 0 \right\} \right], \tag{3.7a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall\, t \in \mathcal{T}. \tag{3.7b}$$

Problem (3.7) can be re-written as the following linear problem using the auxiliary variable y^d [45], which represents the excess loss above threshold τ in each scenario $d \in \mathcal{D}$. To implement the CVaR-based formulation in practice, we approximate the expectation under $\mathbb{P}$ using its empirical distribution derived from a finite set of scenarios $\mathcal{D}$, the same process as to obtain (3.2), leading to

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}, \tau \in \mathbb{R}, y^d \in \mathbb{R}} \quad -\rho \left(\frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} \mathbf{s}_t^{d\top} \mathbf{q}_t \right) + (1-\rho) \left(\tau + \frac{1}{|\mathcal{D}|\alpha} \sum_{d \in \mathcal{D}} y^d \right), \tag{3.8a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}, \tag{3.8b}$$

$$y^d + \tau + \sum_{t \in \mathcal{T}} \mathbf{s}_t^{d\top} \mathbf{q}_t \geq 0, \qquad \forall d \in \mathcal{D}, \tag{3.8c}$$

$$y^d \geq 0, \qquad \forall d \in \mathcal{D}. \tag{3.8d}$$

The resulting model is one that incorporates risk and scenarios into the optimization process. We refer to (3.8) as the `SO-CVaR` model and we use it later to benchmark our approach.

3.2.2 Distributional Uncertainty Modelling

DRO provides a framework for decision-making under uncertainty by considering a set of possible probability distributions rather than relying on the empirical distribution like in [6]. This approach aims to ensure that solutions remain robust against variability in the underly-

ing data. We recall here important concepts from Section 1.6 and adapt them to the virtual bidding notation. To define our ambiguity set, we use the Wasserstein distance $d_{\mathrm{W}}(\mathbb{W}, \mathbb{V})$, which quantifies the dissimilarity between probability distributions $\mathbb{W}, \mathbb{V} \in \mathcal{P}$. Specifically, let $\mathbf{s}_1$ and $\mathbf{s}_2$ be the marginals of $\mathbb{W}$ and $\mathbb{V}$, respectively, and $\Pi(\mathbb{W}, \mathbb{V})$ denotes the set of all joint distributions on $\mathcal{S} \times \mathcal{S}$. The order-$p$ Wasserstein distance [52] with respect to some norm $\|\cdot\|$ is :

$$d_{\mathrm{W}}(\mathbb{W}, \mathbb{V}) = \left(\min_{\pi \in \Pi(\mathbb{W}, \mathbb{V})} \int_{\mathcal{S} \times \mathcal{S}} \|\mathbf{s}_1 - \mathbf{s}_2\|^p \, \mathrm{d}\pi(\mathbf{s}_1, \mathbf{s}_2) \right)^{1/p}.$$

By constructing a Wasserstein ball of radius $\epsilon > 0$ centred around the empirical distribution:

$$\mathcal{B}_\epsilon(\hat{\mathbb{P}}_{\mathcal{D}}) := \left\{ \mathbb{W} \in \mathcal{P} \mid d_{\mathrm{W}}(\mathbb{W}, \hat{\mathbb{P}}_{\mathcal{D}}) \leq \epsilon \right\},$$

as illustrated on Figure 3.2, we can delineate a set of plausible distributions that captures the uncertainty inherent to the data. In this context, we formulate a DRO model to maximize the virtual bidding benefits under the worst-case distribution within the Wasserstein ball $\mathcal{B}_\epsilon(\hat{\mathbb{P}}_{\mathcal{D}})$. This ambiguity set includes all distributions whose Wasserstein distance from the empirical distribution does not exceed ϵ.

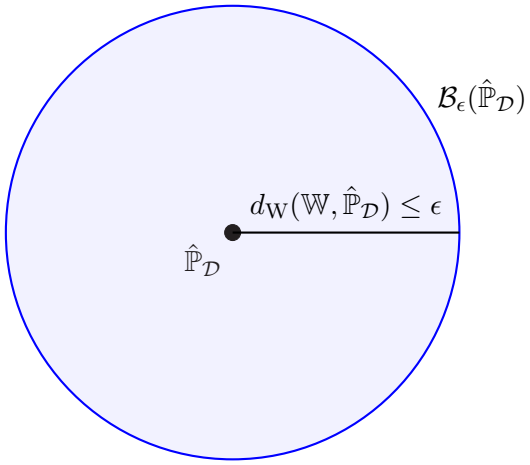

Figure 3.2 Wasserstein ball representation

The distributionally robust reformulation of (3.1) is given by:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}} \max_{\mathbb{W} \in \mathcal{B}_\epsilon(\hat{\mathbb{P}}_{\mathcal{D}})} \quad \mathbb{E}^{\mathbb{W}} \left[-\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t \right], \tag{3.9a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}. \tag{3.9b}$$

Before we proceed to the reformulation of the inner maximization of (3.9), we define, for some vector $\mathbf{x} \in \mathbb{R}^{|\mathcal{N}|}$, the dual norm $\|\cdot\|_*$ of the norm $\|\cdot\|$ as:

$$\|\mathbf{x}\|_* = \sup_{\|\mathbf{y}\| \leq 1} \mathbf{x}^\top \mathbf{y}.$$

In our context, we choose the Euclidean norm in the Wasserstein distance. For the Euclidean norm, the dual norm is equal to the norm itself, i.e., $\|\mathbf{q}_t\|_* = \|\mathbf{q}_t\|_2$. The 2-norm is selected because it promotes a strategy that distributes bids across the different buses on the grid. We also choose the order $p = 1$ that corresponds to the Wasserstein-1 distance, also known as the earth mover distance. It is computationally simpler compared to higher-order Wasserstein distances.

Because the objective function of (3.9) is linear, we can re-express the inner maximization as [45, Remark 6.6]:

$$\min_{\lambda \in \mathbb{R}} \quad -\frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} \mathbf{s}_t^{d\top} \mathbf{q}_t + \epsilon\lambda, \tag{3.10a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_* \leq \lambda, \qquad \forall t \in \mathcal{T}. \tag{3.10b}$$

Substituting (3.10) in (3.9) leads to:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}, \lambda \in \mathbb{R}} \quad -\frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \sum_{t \in \mathcal{T}} \mathbf{s}_t^{d\top} \mathbf{q}_t + \epsilon\lambda, \tag{3.11a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_2 \leq \lambda, \qquad \forall t \in \mathcal{T}, \tag{3.11b}$$

$$\|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}. \tag{3.11c}$$

Problem (3.11) provides a distributionally robust virtual bidding model that is readily solvable with off-the-shelf solvers. For a linear problem like ours, the DRO reformulation is equivalent to the original problem with an added regularizer [13]. We refer to (3.11) as the `DRO` model and we use it as well to benchmark our approach.

3.2.3 Integrated Volatility-and Uncertainty-Aware Virtual Bidding

We now present our integrated bidding model. We combine the CVaR to account for price volatility and DRO to handle uncertainty in the market's underlying distribution in a data-driven way. CVaR protects against extreme price changes, while DRO ensures robustness to shifts in market conditions. This combined model addresses both difficulties in a single

problem. Following the approach shown in [45], we can write the risk-averse problem (3.7) as:

$$\min_{\mathbf{q}_t \in \mathbb{R}^{|\mathcal{N}|}, \tau \in \mathbb{R}} \quad \mathbb{E}^{\mathbb{P}}\Big[\max_{k \in \mathcal{K}} a_k \left(\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t\right) + b_k \tau\Big], \tag{3.12a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}, \tag{3.12b}$$

where $\mathcal{K} = \{1, 2\}$, $a_1 = -\rho$, $a_2 = -\rho - \frac{1-\rho}{\alpha}$, $b_1 = 1 - \rho$, $b_2 = (1-\rho)(1-\frac{1}{\alpha})$. Although the exact distribution $\mathbb{P}$ remains unknown, we propose using box constraints following [18] to define the uncertainty set $\mathcal{S}$ as the range within which the uncertain DART spread vector $\mathbf{s}_t$ is allowed to vary. We set:

$$\mathcal{S} = \{\mathbf{s}_t \in \mathbb{R}^{|\mathcal{N}|} : |s_{i,t}| \leq \Lambda, \quad \forall i \in \mathcal{N}, t \in \mathcal{T}\}, \tag{3.13}$$

where each component of $\mathbf{s}_t$, $s_{i,t}$, is bounded by the maximum value $\Lambda > 0$, ensuring that all elements of $\mathbf{s}_t$ remain within this fixed range as illustrated in Figure 3.3.

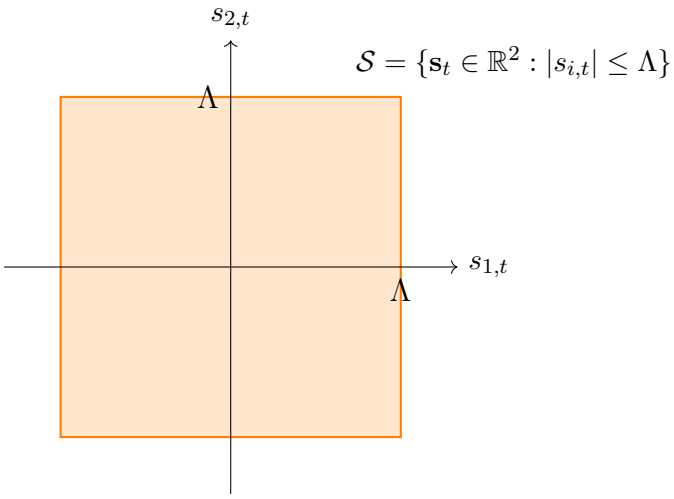

Figure 3.3 Box constraints representation in $\mathbb{R}^2$

The distributionally robust counterpart of (3.12) with respect to the Wasserstein ambiguity set $\mathcal{B}_\epsilon(\hat{\mathbb{P}}_\mathcal{D})$ is given by:

$$\min_{\mathbf{q}_t, \tau} \max_{\mathbb{W} \in \mathcal{B}_\epsilon(\hat{\mathbb{P}}_\mathcal{D})} \mathbb{E}^{\mathbb{W}}\Big[\max_{k \in \mathcal{K}} a_k \left(\sum_{t \in \mathcal{T}} \mathbf{s}_t^\top \mathbf{q}_t\right) + b_k \tau\Big], \tag{3.14a}$$

$$\text{s.t.} \quad \|\mathbf{q}_t\|_1 \leq L, \qquad \forall t \in \mathcal{T}. \tag{3.14b}$$

Using [45, (27)] and setting the vector $\boldsymbol{\nu}_{t,k}^d \in \mathbb{R}^{|\mathcal{N}|}$ and the scalar $\mu_k^d \in \mathbb{R}$ as variables associated to the uncertainty set (3.13) as in [18] yields the readily implementable form:

$$\min_{\mathbf{q}_t,\tau,\lambda,x^d,\boldsymbol{\nu}_{t,k}^d,\mu_k^d} \quad \lambda\epsilon + \frac{1}{|\mathcal{D}|}\sum_{d\in\mathcal{D}} x^d, \tag{3.15a}$$

$$\text{s.t.} \quad b_k\tau + a_k\left(\sum_{t\in\mathcal{T}} {\mathbf{s}_t^d}^\top \mathbf{q}_t\right) + \sum_{t\in\mathcal{T}} {\boldsymbol{\nu}_{t,k}^d}^\top \mathbf{s}_t^d + \Lambda\mu_k^d \le x^d, \qquad \forall d\in\mathcal{D}, k\in\mathcal{K}, \tag{3.15b}$$

$$\|\boldsymbol{\nu}_{t,k}^d - a_k\mathbf{q}_t\|_2 \le \lambda, \qquad \forall d\in\mathcal{D}, k\in\mathcal{K}, t\in\mathcal{T}, \tag{3.15c}$$

$$\mu_k^d \ge \|\boldsymbol{\nu}_{t,k}^d\|_1, \qquad \forall d\in\mathcal{D}, k\in\mathcal{K}, t\in\mathcal{T}, \tag{3.15d}$$

$$\|\mathbf{q}_t\|_1 \le L, \qquad \forall t\in\mathcal{T}. \tag{3.15e}$$

The detailed derivation to obtain (3.15) given the uncertainty set (3.13) is provided next in Section 3.2.4. This formulation provides a distributionally robust optimization framework that balances profit maximization and risk mitigation in addition to accounting for the uncertainty. The choice of the Euclidean norm in (3.15) aligns with the approach used to obtain (3.11). We refer to (3.15) as the `DRO-CVaR` model.

3.2.4 Uncertainty Set Demonstration

To clarify how the box uncertainty set (3.13) fits into the general DRO framework of [45, (27)], we provide a brief demonstration showing that (3.15) is a special case of [45]'s cone uncertainty set formulation. This establishes the consistency of our model with the general theoretical results while illustrating the practical simplifications afforded by box constraints. For the sequel, we only consider the variables and the parameters involved in the constraints that differ between (3.15) and [45, (27)]. The cone uncertainty set presented in [45, (27)] is defined by:

$$\mathcal{U} = \{\mathbf{s}_t \in \mathbb{R}^{|\mathcal{N}|} : \mathbf{C}^\top\mathbf{s}_t \le \mathbf{d}, \quad \forall i\in\mathcal{N}, t\in\mathcal{T}\}, \tag{3.16}$$

where $\mathbf{C}^\top \in \mathbb{R}^{2|\mathcal{N}|\times|\mathcal{N}|}$ and $\mathbf{d} \in \mathbb{R}^{2|\mathcal{N}|}$. The box uncertainty set of (3.15) is defined by:

$$\begin{aligned}\mathcal{S} &= \{\mathbf{s}_t \in \mathbb{R}^{|\mathcal{N}|} : |s_{i,t}| \leq \Lambda, \quad \forall i \in \mathcal{N}, t \in \mathcal{T}\} \\ &= \{\mathbf{s}_t \in \mathbb{R}^{|\mathcal{N}|} : s_{i,t} \leq \Lambda,\ -s_{i,t} \leq \Lambda \quad \forall i \in \mathcal{N}, t \in \mathcal{T}\},\end{aligned}$$

where each component of $\mathbf{s}_t$, $s_{i,t}$, is bounded by a maximum value $\Lambda > 0$, ensuring that all elements of $\mathbf{s}_t$ remain within this fixed range. This is a special case of (3.16) where:

$$\mathbf{C} = \begin{pmatrix}\mathbf{I} \\ -\mathbf{I}\end{pmatrix}, \qquad \mathbf{d} = \Lambda \begin{pmatrix}\mathbf{1} \\ \mathbf{1}\end{pmatrix}.$$

Let $\boldsymbol{\gamma} \in \mathbb{R}^{2|\mathcal{N}|}$. From [45, (27c)], we have :

$$\begin{aligned} & \boldsymbol{\gamma}^\top(\mathbf{d} - \mathbf{C}\mathbf{s}_t) \geq 0 \\ \Leftrightarrow \qquad & -\boldsymbol{\gamma}^\top \begin{pmatrix}\mathbf{I} \\ -\mathbf{I}\end{pmatrix} \mathbf{s}_t + \boldsymbol{\gamma}^\top \Lambda \begin{pmatrix}\mathbf{1} \\ \mathbf{1}\end{pmatrix} \geq 0.\end{aligned}$$

This allows us to re-express the last two terms of (3.15b) in terms of:

$$\boldsymbol{\nu} = -\begin{pmatrix}\mathbf{I} \\ -\mathbf{I}\end{pmatrix}^\top \boldsymbol{\gamma} = \begin{pmatrix}-\mathbf{I} & \mathbf{I}\end{pmatrix}\boldsymbol{\gamma}, \quad \mu = \boldsymbol{\gamma}^\top \begin{pmatrix}\mathbf{1} \\ \mathbf{1}\end{pmatrix}.$$

Let

$$\boldsymbol{\gamma} = \begin{pmatrix}\boldsymbol{\alpha} \\ \boldsymbol{\beta}\end{pmatrix},$$

where $\boldsymbol{\alpha} \in \mathbb{R}^{|\mathcal{N}|}$ and $\boldsymbol{\beta} \in \mathbb{R}^{|\mathcal{N}|}$. We can rewrite $\boldsymbol{\nu}$ and μ: as $\boldsymbol{\nu} = -\boldsymbol{\alpha} + \boldsymbol{\beta}$ and $\mu = \mathbf{1}^\top(\boldsymbol{\alpha} + \boldsymbol{\beta})$. Recalling (3.15d) and substituting the above definitions, we obtain:

$$\begin{aligned} & \mu \geq \|\boldsymbol{\nu}\|_1 \\ \Leftrightarrow \qquad & \mathbf{1}^\top(\boldsymbol{\alpha} + \boldsymbol{\beta}) \geq |-\boldsymbol{\alpha}_1 + \boldsymbol{\beta}_1| + ... + |-\boldsymbol{\alpha}_n + \boldsymbol{\beta}_n| \\ \Leftrightarrow \qquad & \sum_{i\in\mathcal{N}} (\boldsymbol{\alpha}_i + \boldsymbol{\beta}_i) \geq \sum_{i\in\mathcal{N}} |\boldsymbol{\alpha}_i - \boldsymbol{\beta}_i|.\end{aligned}$$

This is always true if $\boldsymbol{\gamma} \geq 0$, which coincides with [45, (27e)]. Finally, we note that (3.15d):

$$\begin{aligned} & \|\boldsymbol{\nu}\|_2 \leq 0 \\ \Leftrightarrow \qquad & \|-\boldsymbol{\alpha} + \boldsymbol{\beta}\|_2 \leq 0 \\ \Leftrightarrow \qquad & \|\mathbf{C}^\top\boldsymbol{\gamma}\|_2 \leq 0,\end{aligned}$$

is equivalent to [45, (27d)]. Note that the condition $\|\boldsymbol{\nu}\|_2 \leq 0$ forces $\boldsymbol{\nu} = 0$, which is trivially satisfied in this case because the box set imposes symmetric bounds that cancel out in the dual representation. This demonstrates that the box uncertainty set is a special case of the general conic uncertainty set used in [45, (27)].

CHAPTER 4 ONLINE VIRTUAL BIDDING STRATEGY

We devise a rolling virtual bidding strategy centred on our model (3.15). Our strategy is based on a dataset $\mathcal{D}$ updated daily to determine bids for the next day. To implement our model within a bidding strategy, we begin by selecting the most relevant data for each day that will define the dataset $\mathcal{D}$. Next, we optimize the hyperparameters over a training period to ensure the best strategy performance.

4.1 Selection of Similar Days

A key challenge in data-driven optimization is selecting historical data that best represents future conditions. In our study, we carefully select relevant historical data by identifying dates with load and generation profiles similar to the target bidding day. For each target date, we consider a two-year rolling window of past data and calculate the similarity between the target bidding date and each day within the historical window.

Let the similarity metric $\Gamma : (\mathbb{R}^{24} \times \mathbb{R}) \times (\mathbb{R}^{24} \times \mathbb{R}) \to \mathbb{R}$ measure the similarity between two days. The load profile $\mathbf{p} \in \mathbb{R}^{24}$ collects the total network load (across all buses) for each timestep over the day in a vector, while the offline generation capacity $\theta \in \mathbb{R}$ reflects the total megawatts (MW) of generating capacity forecasted to be offline for the ISO throughout the day. The similarity metric combines the Euclidean distance between load profiles and the absolute difference in offline capacities in addition to a penalty $\zeta > 0$ for splitting weekdays with weekends. The Γ metric is defined as:

$$\Gamma\left((\mathbf{p}_{\text{target}}, \theta_{\text{target}}), (\mathbf{p}, \theta)\right) = 2\|\mathbf{p}_{\text{target}} - \mathbf{p}\|_2 + |\theta_{\text{target}} - \theta| + \zeta,$$

where $\zeta = 0$ if both days are either weekdays or weekends, and $\zeta = 1000$ otherwise. The fixed value of ζ ensure the penalty significantly influences Γ without excluding the possibility that a weekend day is more similar to a target day than any weekday, and vice versa. The Γ metric ranks past days by similarity, with smaller values indicating more similar days. Finally, we remark that other comparison methods could be used in our approach.

4.2 Strategy Execution

We use Γ to build the relevant historical data set $\mathcal{D}_j$ to be used by our model at day j. We discuss the number of data to include in $\mathcal{D}_j$ in Section 4.3. We then solve (3.15) using $\mathcal{D}_j$ and

commit the resulting bids for day j. For example, consider bids to be placed for $j = \text{May } 13^{\text{th}}$ on May 12^{th}. Let us assume $|\mathcal{D}| = 3$, evaluating Γ on the month of data preceding j, we determine that the most similar days to May 13^{th} are April 30^{th}, May 3^{rd}, and May 7^{th} and we form $\mathcal{D}_j$ accordingly. The process is visualized in Figure 4.1.

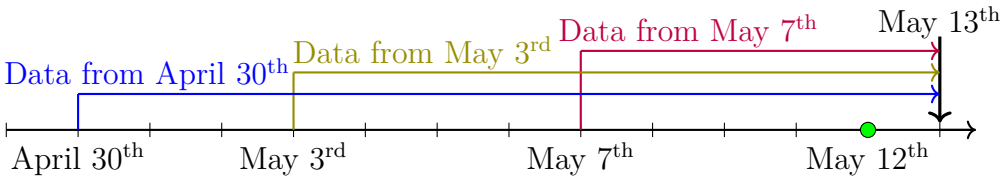

Figure 4.1 Strategy execution example for bids to be placed on May 12^{th} for the next day

Finally, we repeat this process for all bidding days $j \in \{1, 2, ..., J\}$, where $J \in \mathbb{N}$ is the bidding horizon, adding each day one new data point in the pipeline and removing the oldest one. The proposed strategy is deployed as an online process, updating daily both the dataset $\mathcal{D}_j$ and the broader dataset from which $\mathcal{D}_j$ is constructed. This broader dataset consists of a rolling two-year window that is updated each day to incorporate the most recent data while discarding the oldest observations. Within this window, $\mathcal{D}_j$ is dynamically selected based on the similarity of past days to the current one. This ensures that the model continuously adapts to the most relevant information while accounting for real-world changes such as evolving demand, climate variability, and grid modifications.

4.3 Strategy Tuning

In this section, we describe our bidding strategy and tuning method, which we also apply to the benchmark models for a fair comparison. The hyperparameters of each model are optimized on a training set comprising 12 months of historical DART spreads from NYISO from February 1^{st}, 2023, to January 31^{st}, 2024 with the rolling window going back as far as February 1^{st}, 2021 when tuning starts. The online bidding strategy is then evaluated on a separate 8-month testing set illustrated in Figure 4.2.

Problem (3.15) includes hyperparameters ϵ, ρ, $|\mathcal{D}|$, and Λ. These are tuned using the `Optuna` library [1], which is designed for efficient hyperparameter optimization. Unlike exhaustive grid search or simple random sampling, `Optuna` employs a model-based blackbox optimization

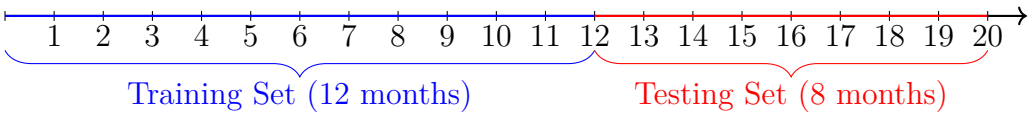

Figure 4.2 Training and testing sets

(BBO) approach [4]. Its default algorithm, the tree-structured Parzen estimator (TPE), is a Bayesian optimization method that constructs probabilistic models of promising and less promising regions in the hyperparameter space and selects new candidates by maximizing an acquisition function [7]. Its objective is to maximize the annualized Calmar ratio C. The Calmar ratio measures the trade-off between risk and return by comparing the annualized return of a bidding strategy to its maximum drawdown and provides an indication of the risk-adjusted performance of the strategy [42]. We first initialize a portfolio that has a value of \$1M, denoted as $v_0 = 10^6$.

The daily returns r_j for each day $j \in \{1, 2, ..., J\}$ are calculated as the total daily profit or loss from all bids, given by:

$$r_j = \sum_{t \in \mathcal{T}} -\mathbf{s}_t^\top \mathbf{q}_t.$$

The returns are then scaled relative to the value of the portfolio from the previous day v_{j-1} to obtain the scaled return η_j, capturing the proportionate growth or decline in the portfolio:

$$\eta_j = \frac{r_j}{v_{j-1}}. \tag{4.1}$$

The portfolio value is updated iteratively by adding the daily return:

$$v_j = v_{j-1} + r_j.$$

The annualized return $R_{\text{annualized}}$ is then calculated as the cumulative return for the year, scaled to a full 365-day period:

$$R_{\text{annualized}} = \left(\prod_{j=1}^{J} (1 + \eta_j) \right)^{\frac{365}{J}} - 1, \tag{4.2}$$

where J is the total number of trading days. The maximum drawdown (MDD) is defined as the largest peak-to-trough decline during the trading period:

$$\text{MDD} = \max \left(\frac{v_{\text{peak}} - v_{\text{trough}}}{v_{\text{peak}}} \right), \tag{4.3}$$

where v_{peak} is the highest portfolio value before a drawdown, and v_{trough} is the lowest value reached during the drawdown. Finally, combining (4.2) and (4.3), the Calmar ratio C is computed as:

$$C = \frac{R_{\text{annualized}}}{\text{MDD}}. \tag{4.4}$$

In our numerical example, these calculations are performed using the `empyrical` library [22]. The use of the Calmar ratio in our analysis allows for a robust evaluation of the portfolio's performance in high-risk environments, such as those observed in electricity markets.

The Calmar ratio is chosen to tune our model because it is a recognized tool for measuring risk-adjusted performance of investment strategies, particularly by focusing on downside risk [2, 15, 73]. This makes it a suitable objective to select hyperparameters during tuning. Each model is tuned on the same training set for 1000 trials with the quantity limit constraint set to $L = 400$ MWh. The tuning is parallelized and performed on 128 Intel(R) Xeon(R) Platinum 8375C CPUs @ 2.90GHz with 4.2 TB of memory, using the computing power of *Lab Castor* from Hydro-Québec. We let `Optuna` try values for the hyperparameters in the ranges provided in Table 4.1 with the objective of maximizing the Calmar ratio over the 12-month training set. It takes `Optuna` around a week to compute the 1000 trials for our strategy.

Table 4.1 Hyperparameter range for `Optuna` on the 12-month training set

Hyperparameter	**min**	**max**
$\lvert\mathcal{D}\rvert$	2	100
ϵ	5	50
ρ	0.2	0.8
Λ	2000	5000

For `SO-CVaR` and `DRO-CVaR`, the hyperparameter α is fixed to 10% to represent a predefined level of risk aversion, commonly used to focus on the worst 10% of outcomes. This approach ensures interpretability and stability, as optimizing α could lead to overly conservative or risk-seeking behaviours. Additionally, fixing α simplifies the hyperparameter optimization process, reducing its dimensionality and leveraging contextual information on acceptable risk levels in the electricity trading sector.

It is worth noting that hyperparameter tuning should not be viewed as a one-time procedure. Electricity market conditions evolve due to changes in demand patterns, renewable generation, congestion levels, and operational rules, all of which can affect the statistical properties of the price spread. For this reason, the tuning process presented in this chapter should be

performed periodically so that the online bidding strategy remains calibrated to the most recent data and reflects current market behaviour. Regular recalibration ensures that the model continues to generalize well and preserves its robustness as new operating regimes emerge.

CHAPTER 5 NUMERICAL RESULTS

This section presents and analyzes the results of our `DRO-CVaR` strategy. The strategy is tested on a 8-month period with hyperparameters obtained from the tuning process introduced in Section 4.3. We remark that the eight months of data utilized for testing was not used for hyperparameter tuning and, as such, the dataset is considered as out-of-sample.

5.1 Case Study Description

The case study is based on data from NYISO, which operates an electricity market with 11 geographic zones. These zones are aggregated from a larger number of buses, each representing a point of generation or consumption. Virtual bids are submitted for each hour before the virtual bidding market closing time (5 AM), and the system clears the DAM a day prior to the RTM. The case study considers placing bids for every hour of every day over an 8-month testing period, from February 1$^{\text{st}}$, 2024, to October 1$^{\text{st}}$, 2024.

We simulate bidding for every hour of every day from February 1$^{\text{st}}$, 2024 to October 1$^{\text{st}}$, 2024, resulting in 244 daily profit observations. All strategies, including our benchmarks, are subject to the same hourly bidding quantity constraint of $L = 400$ MWh.

5.2 Benchmarking our Strategy

We compare our `DRO-CVaR` strategy against four alternative approaches:

- `SO`: a scenario-based stochastic optimization approach provided in (3.2);
- `SO-CVaR`: a risk-averse version of `SO` provided in (3.6);
- `DRO`: the distributionally robust formulation without CVaR provided in (3.11) and;
- `EW`: an equally-weighted portfolio benchmark.

The equally-weighted strategy (`EW`) places identical sell quantities across all zones and hours, assuming, consistent with NYISO market patterns, that DAM prices tend to exceed RTM prices. It is widely used in practice and is regarded as a standard benchmark [23, 29, 57]. In this work, it serves as a static baseline illustrating the benefits of data-driven bidding.

5.3 Analysis of Daily and Cumulative Profits

Daily and cumulative profits for all strategies are displayed in Figures 5.1a and 5.1b respectively.

Analyzing Figures 5.1a and 5.1b, it can be observed that combining distributionally robust optimization with CVaR tends to achieve both profitability and risk mitigation. Strategies employing the CVaR, namely `SO-CVaR` and `DRO-CVaR`, show reduced downside risk, although only `DRO-CVaR` consistently maintains positive cumulative profits.

The `DRO`-based strategy (without CVaR) achieves the highest cumulative profit, confirming that robustness against distributional shifts can significantly improve long-run performance. However, this comes at the cost of larger downside deviations, meaning that `DRO` can experience considerable daily losses. We observe that `DRO-CVaR` never suffers the largest daily loss and is regularly amongst the most important earners for any given day and yields the second-best results in terms of cumulative profits at the end of the test period. This indicates that the strategy is more robust and reliable compared to the benchmarks, albeit at the cost of the loss of some higgh-profit opportunities.

The `DRO-CVaR`-based strategy has the most consistent profits, as losses are limited, leading to large cumulative profits. It is the only strategy that never has its cumulative profits dip below zero, which demonstrates better robustness than the other strategies.

All four data-driven strategies (`SO`, `SO-CVaR`, `DRO`, `DRO-CVaR`) outperform the equally-weighted baseline (`EW`) by a large margin. This emphasizes the benefits of exploiting price patterns and uncertainty-aware optimization rather than using static heuristics.

5.4 Profit Distribution and Volatility Analysis

Figure 5.2 illustrates the distribution of the daily profits of our strategy compared to all the benchmarks. We note that our strategy is less volatile than the other observed strategies. They are all centred slightly right of zero, which indicates profitability, but the main difference is that the number of negative profit occurrences is diminished, although at the cost of missing out higher positive profit opportunities. Relative to `EW`, `SO`, `SO-CVaR`, and `DRO`, our strategy eliminates many of the severe losses that the benchmarks occasionally suffer. This confirms the expected benefit of integrating CVaR within a distributionally robust model. While `DRO-CVaR` misses some of the very large positive outliers that other strategies achieve, this is the expected behaviour of a risk-averse model. The strategy prioritizes consistent gains and controlled risk exposure over capturing occasional large opportunities.

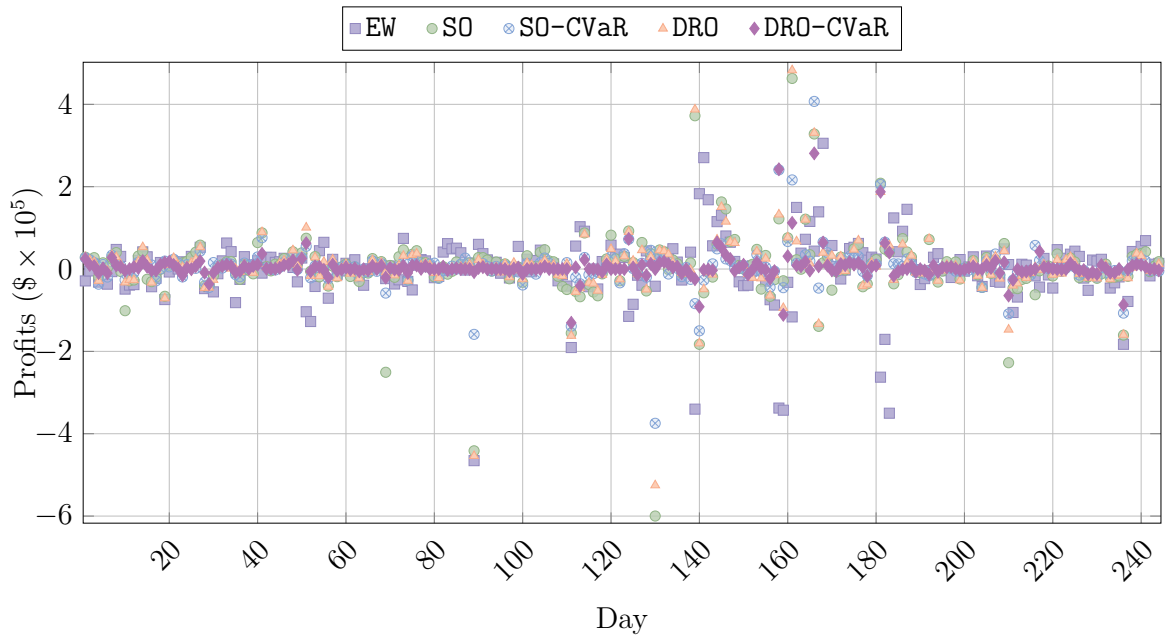

(a) Daily profits over the 8-month testing set

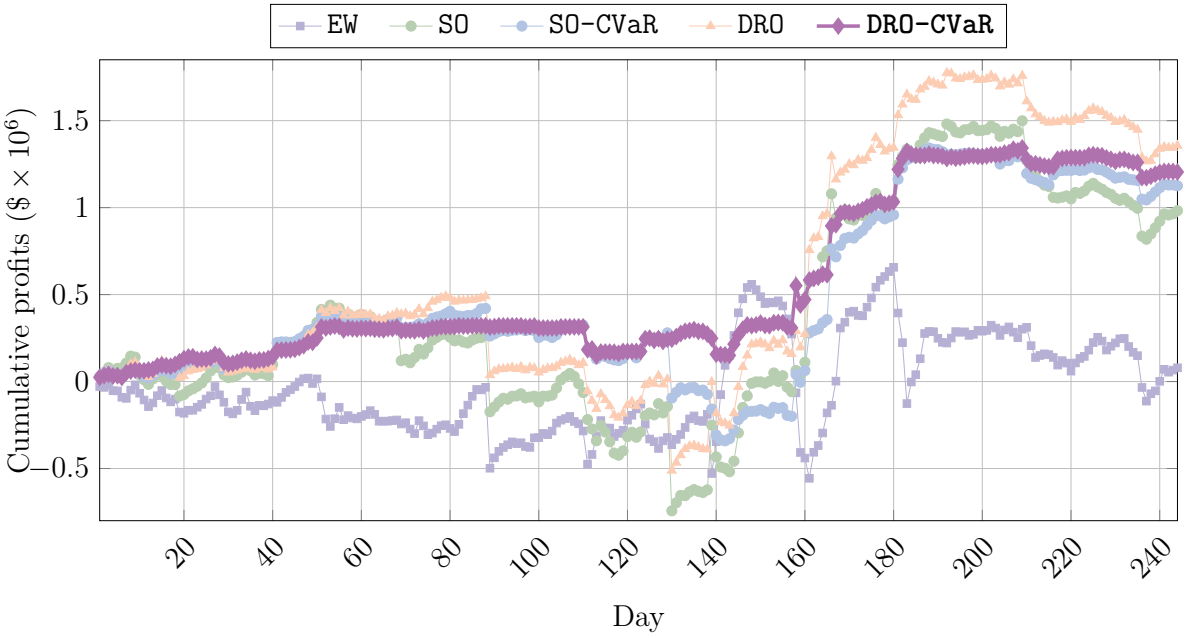

(b) Cumulative profits over the 8-month testing set

Figure 5.1 Profits over the 8-month testing set

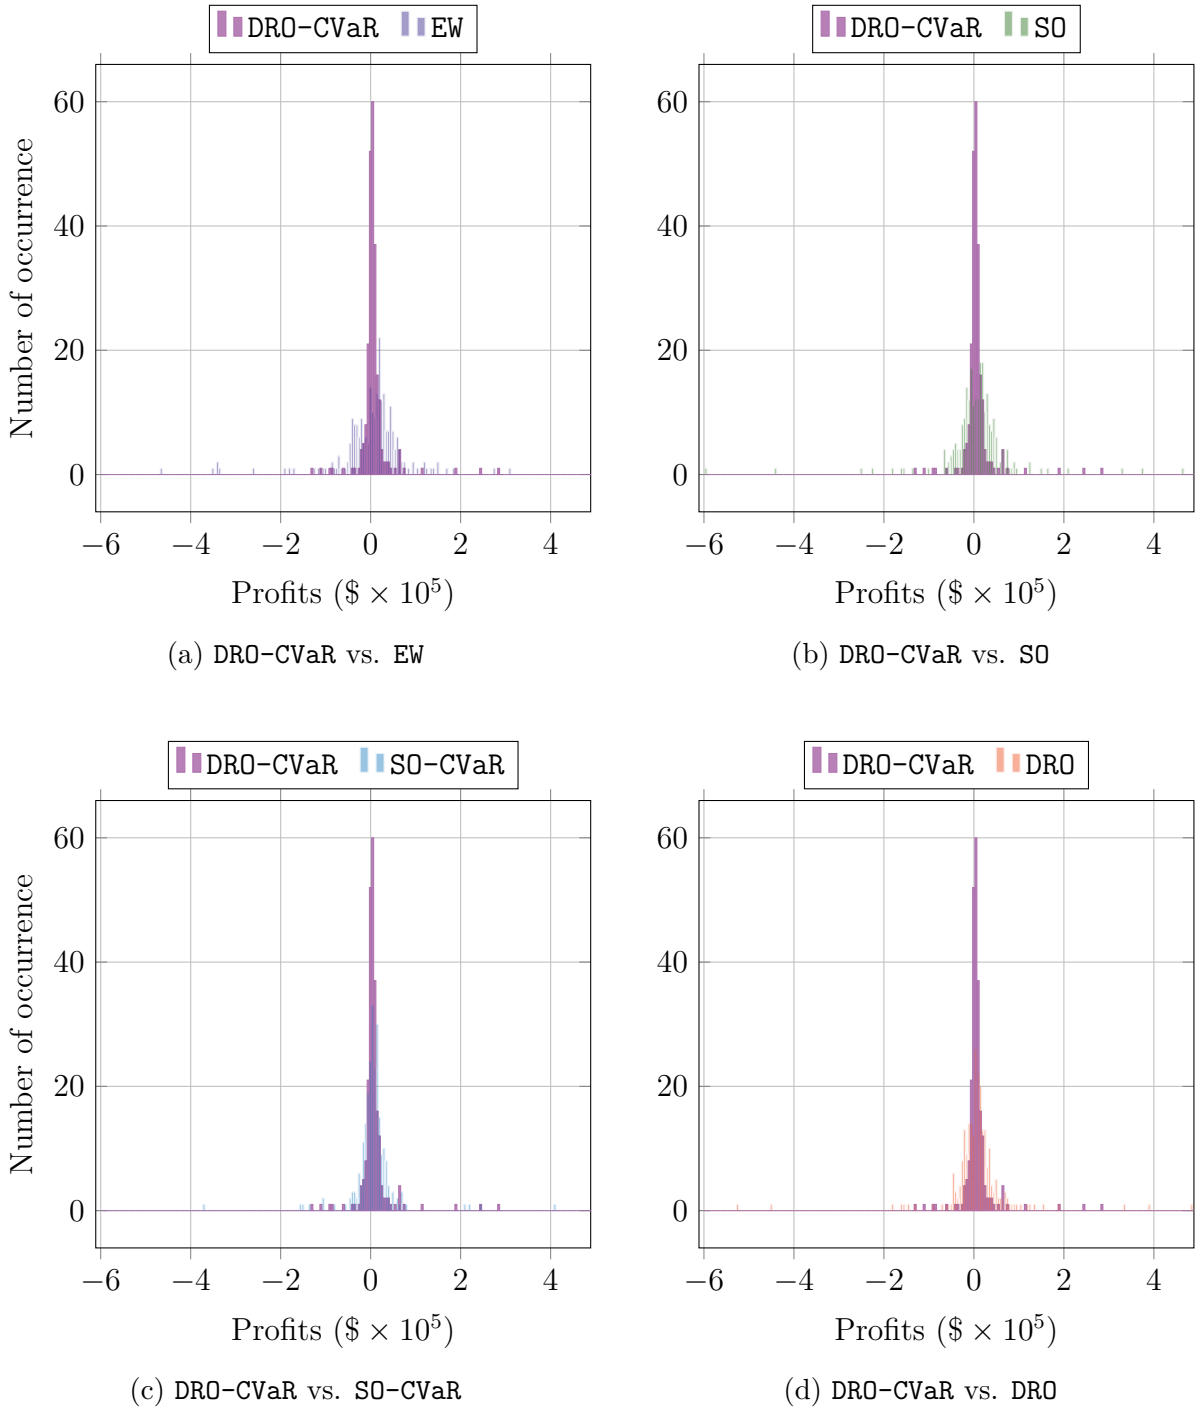

(a) DRO-CVaR vs. EW

(b) DRO-CVaR vs. SO

(c) DRO-CVaR vs. SO-CVaR

(d) DRO-CVaR vs. DRO

Figure 5.2 Comparison of daily profits over the 8-month testing set

5.5 Profits Distribution Analysis

Figure 5.3 illustrates the distribution of profits across the different strategies, presenting the interquartile ranges, medians, whiskers, and outliers. It can be observed that the CVaR-based strategies exhibit reduced profit volatility, as evidenced by the narrower interquartile ranges and shorter whiskers. Among them, `DRO-CVaR` has the smallest whiskers and the lowest interquartile spread, confirming it is the least volatile option. This figure reinforces the earlier conclusion: `DRO-CVaR` provides the best balance between risk and return, minimizing extreme losses without sacrificing overall profitability.

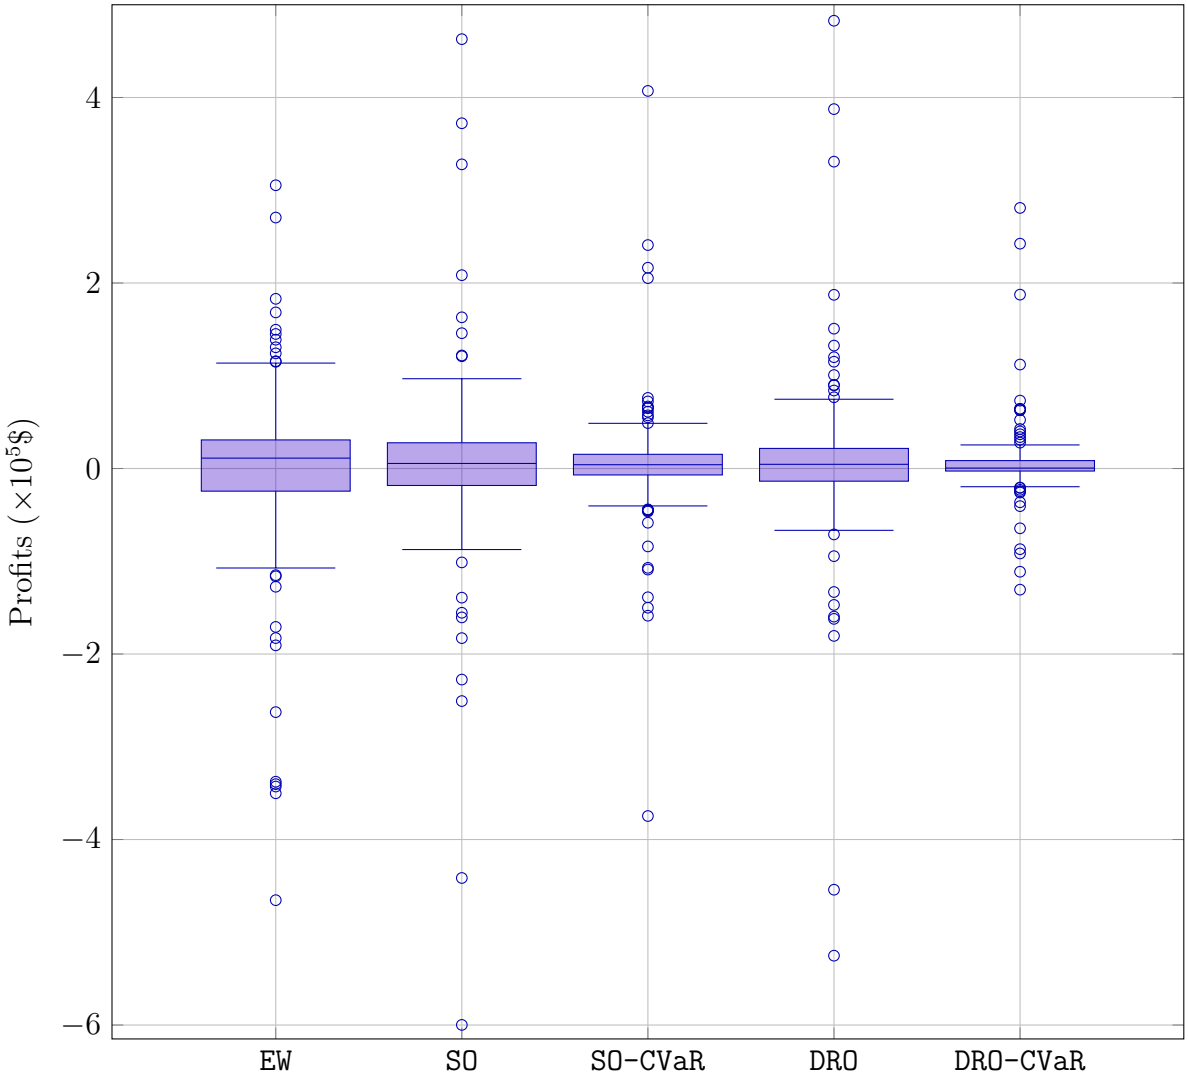

Figure 5.3 Boxplot of daily profits over the 8-month testing period

5.6 Efficiency Analysis

In evaluating the performance of the proposed strategies, it is important to consider not just their total profits but also the profits relative to the quantity of bids placed. Each strategy operates under a consistent set of bidding constraints; however, some pursue a more aggressive bidding approach, bidding the maximum quantity at every timestep, while others adopt a more conservative strategy, opting not to bid to the limit. Figure 5.4 presents the scaled profits for each strategy, which are calculated as the cumulative profit at the end of the 8-month period divided by the total quantity of virtual bids. This measure provides insight into the profitability per MWh, illustrating how effectively each strategy translates bid quantities into realized gain.

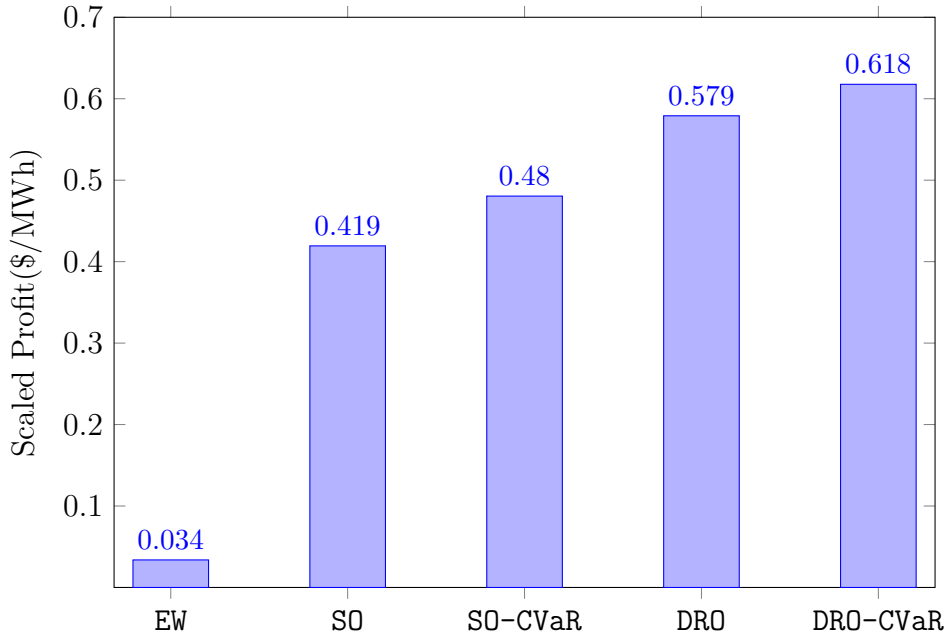

Figure 5.4 Scaled profit over the 8-month testing period

Figure 5.4 shows that the best-performing strategy in terms of scaled profit is the one based on `DRO-CVaR`. The `DRO`-based strategy also demonstrates strong scaled profit, highlighting the importance of uncertainty considerations in increasing profits per MWh. The moderate scaled profits of `SO-CVaR` suggest it could be improved with additional robustness features while the very low cumulative gain-per-MWh of `EW` reflects a poor performance of the strategy.

5.7 Performance in Terms of Risk-Adjusted Metrics

To evaluate performance in a risk-adjusted context, two other metrics are used: the Sharpe and Calmar ratios. The Sharpe ratio measures the return earned per unit of risk, allowing us to assess how each strategy balances reward against volatility [62].

5.7.1 Sharpe Ratio

To calculate the Sharpe ratio S, we start with an initial portfolio value of $v_0 = 10^6$. The Sharpe ratio is then calculated as the mean scaled returns $\bar{\eta}$ over the total number of trading days $J = 244$ days (8 months) and the standard deviation σ_η of the scaled daily returns η_j as defined in (4.1) :

$$S = \frac{\bar{\eta}}{\sigma_\eta}\sqrt{J}.$$

The Sharpe ratio allows us to compare the reward versus risk consistently across strategies and periods [6]. A higher Sharpe ratio indicates a more efficient strategy as it achieves greater returns per unit of risk. In Figure 5.5, we illustrate the Sharpe ratios for each strategy.

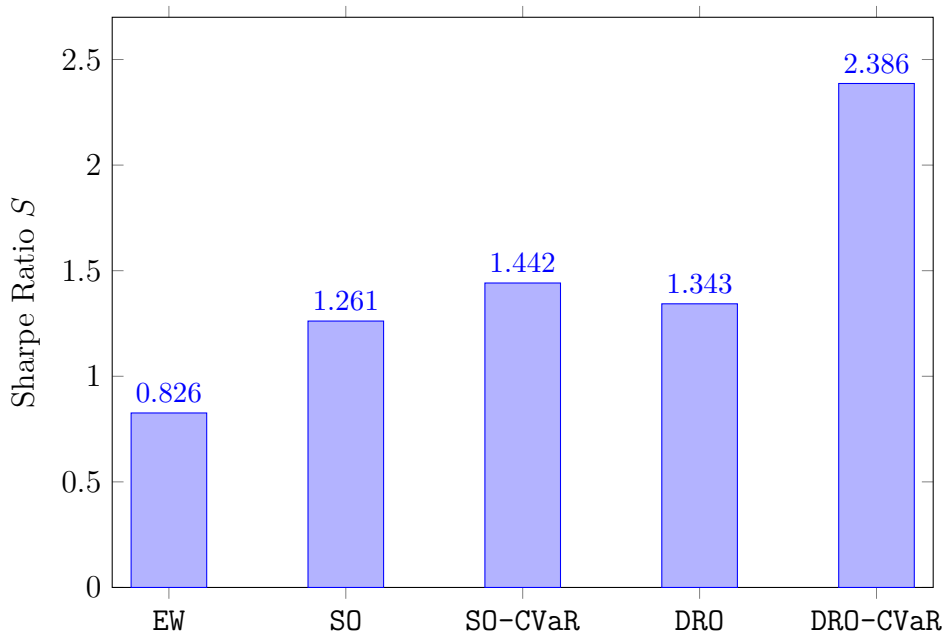

Figure 5.5 Sharpe ratios over the 8-month testing period

Figure 5.5 shows that the `DRO-CVaR` strategy significantly outperforms the other benchmarks, achieving a Sharpe ratio of 2.386. This indicates a highly efficient strategy with greater returns per unit of risk. Figure 5.5 validates that strategies incorporating CVaR achieve

higher risk-return outcomes with SO-CVaR having the second best Sharpe ratio with 1.442. This analysis underscores the effectiveness of DRO-CVaR in balancing risk and profits.

5.7.2 Calmar Ratio

While the Sharpe ratio focuses on volatility, the Calmar ratio evaluates returns relative to maximum drawdown (MDD), thereby reflecting the severity of worst-case performance. Figure 5.6 provides Calmar ratios defined in (4.4) for all strategies. A higher Calmar ratio indicates a strategy's ability to deliver returns that sufficiently compensate for its worst historical losses. To ensure fairness and avoid any bias, the tuning process for all data-driven strategies was performed by maximizing the Calmar ratio.

The Calmar ratio analysis highlights the superior performance of the DRO-CVaR strategy, which achieves a remarkably high Calmar ratio of 9.570 as shown in Figure 5.6. In comparison, the SO-CVaR and DRO-based strategies have Calmar ratios of 2.199 and 2.116, respectively, indicating good performance though limited compared to DRO-CVaR. This reaffirms that our strategy is the top-performing one in terms of risk-adjusted returns.

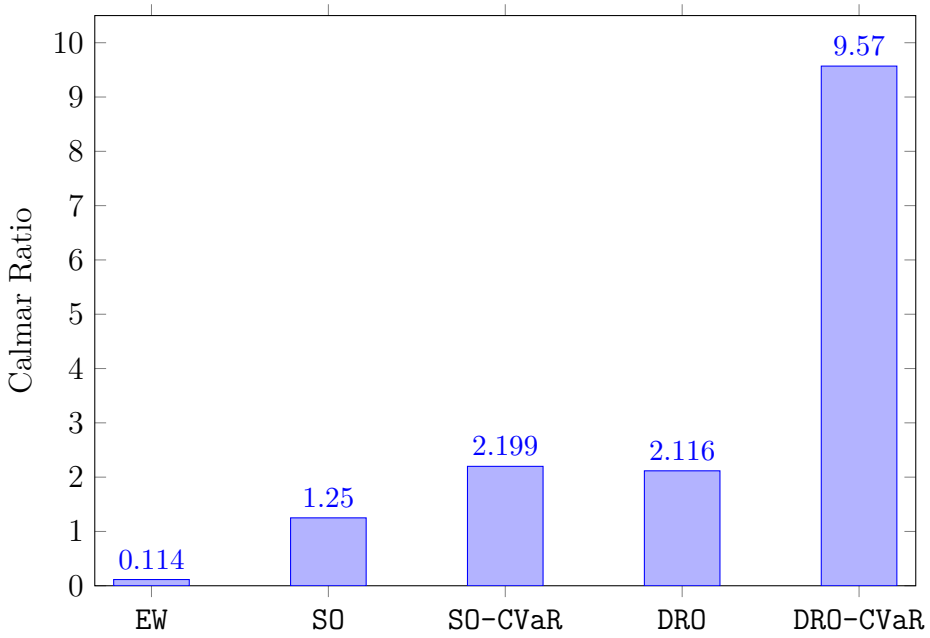

Figure 5.6 Calmar ratios over the 8-month testing period

5.8 Results Summary

This numerical study demonstrates that incorporating both distributional robustness and tail-risk considerations yields a virtual bidding strategy that is profitable and resilient to adverse market conditions. Across the 8-month out-of-sample testing period, `DRO-CVaR` strikes the best balance between expected returns and downside protection among all evaluated strategies.

First, the daily and cumulative profit trajectories show that `DRO-CVaR` avoids the severe losses frequently experienced by `EW`, `SO`, `SO-CVaR`, and `DRO`, while still capturing a substantial portion of the profitable opportunities. Although the `DRO` model achieves the highest cumulative profit, it does so at the cost of much higher volatility and occasional large negative deviations. In contrast, `DRO-CVaR` never experiences a cumulative drawdown below zero and remains one of the top-performing strategies on most trading days.

Second, the distributional analysis of daily profits confirms that `DRO-CVaR` substantially compresses the left tail of the profit distribution. It eliminates many of the extreme negative outcomes present in the benchmark strategies, confirming that CVaR effectively mitigates worst-case price risk. As expected for a risk-averse formulation, this is accompanied by a reduced frequency of very large positive outliers. Still, the improvement in downside protection clearly dominates the small sacrifice in upside potential.

Third, the boxplot analysis highlights the significantly reduced variability of the `DRO-CVaR` strategy. It exhibits the smallest interquartile range, the shortest whiskers, and the fewest extreme values, all indicative of tighter and more predictable daily performance. Among all tested strategies, `DRO-CVaR` achieves the most favorable risk–return profile, combining strong median profits with the lowest exposure to extreme losses.

Finally, all optimization-based strategies substantially outperform the static equally-weighted benchmark, emphasizing the value of data-driven bidding and of modelling both distributional uncertainty and price volatility. The best values are displayed in bold characters in Table 5.1.

Table 5.1 Performance metrics per strategies on the 8-month testing period

Metric	`EW`	`SO`	`SO-CVaR`	`DRO`	`DRO-CVaR`
Cumulative Profit($)	79,167	982,299	1,125,487	**1,356,383**	1,204,584
Scaled Profit($/MWh)	0.034	0.419	0.48	0.579	**0.618**
Sharpe Ratio	0.826	1.261	1.442	1.343	**2.386**
Calmar Ratio	0.114	1.250	2.199	2.116	**9.570**

Overall, the results validate the core idea of this work: incorporating CVaR into a DRO framework reduces downside risk while maintaining strong returns. The `DRO-CVaR` strategy illustrates that robust and risk-aware data-driven methods can improve virtual bidding outcomes in modern electricity markets characterized by variability and uncertainty.

CHAPTER 6 CONCLUSION

This work introduces a data-driven strategy for virtual bidding in two-settlement electricity markets that combines DRO with CVaR. The proposed `DRO-CVaR` strategy addresses two key challenges: uncertainty in price distributions and exposure to extreme losses. By integrating a Wasserstein distance-based ambiguity set with a CVaR objective, the model provides a tractable convex formulation that can be solved efficiently. The strategy is implemented in an online setting, tuned using hyperparameter optimization, and validated on NYISO data.

6.1 Summary

The `DRO-CVaR` model is designed to capture both distributional uncertainty and tail risk, resulting in a convex optimization problem that can be solved efficiently using commercial solvers. To ensure adaptability to changing market conditions, the model is embedded in an online strategy that updates daily based on a rolling dataset of similar historical days. A hyperparameter tuning pipeline developed using `Optuna` to maximize the Calmar ratio is employed.

The empirical validation demonstrates the effectiveness of the proposed approach. Using a 12-month training set for hyperparameter tuning and an 8-month out-of-sample testing period, the `DRO-CVaR` strategy is benchmarked against three data-driven alternatives: `SO`, `SO-CVaR`, and `DRO` also tuned accordingly, as well as a static benchmark: `EW`. The results show that `DRO-CVaR` consistently delivered better risk-adjusted performance, ranking first in Sharpe ratio and Calmar ratio. In terms of efficiency, `DRO-CVaR` also delivered the highest scaled profit per MWh, indicating that it translated its bidding quantities into realized gains more effectively than the other strategies. Although `DRO` achieved the highest cumulative profit, `DRO-CVaR` was second best while maintaining positive cumulative returns throughout the testing period, confirming the value of incorporating CVaR for downside protection. Overall, the numerical case study highlights that combining distributional robustness with tail-risk control yields a strategy that is both profitable and resilient under uncertain market conditions.

6.2 Limitations

Despite the strong performance of the proposed framework, several limitations must be acknowledged.

First, the strategy relies heavily on historical data to construct ambiguity sets and tune hyperparameters. This assumption presumes that past price behaviour remains representative of future conditions. However, electricity markets are subject to structural changes such as shifts in generation mix, change in the grid topology, and regulatory reforms. These changes can significantly modify price dynamics, reducing the predictive value of approaches based on past observations. Although we incorporate load and generation capacity forecasts, significant unforeseen changes in the network could substantially reduce strategy performance. A potential solution is to use adaptive ambiguity sets that update dynamically as new data becomes available, mitigating the risk of model obsolescence [8].

Second, the current model assumes price-taking behaviour and does not account for market-impact effects. In NYISO, because virtual bidding is only done the zonal level, the impact of individual bidders is limited, but in other markets where the virtual bids are placed at the nodal level, the market impact is assumed to be greater. In reality, large virtual bids can influence clearing prices, especially in less liquid markets. Ignoring this interaction may lead to suboptimal or even infeasible strategies. Incorporating market equilibrium models or bi-level optimization frameworks could help capture the feedback between bidding decisions and price formation. While this adds complexity, recent advances in mathematical programming with equilibrium constraints (MPEC) [16] and game-theoretic formulations [46] provide tractable approaches for modelling strategic interactions.

Finally, the tuning process, while effective, is computationally intensive and performed offline. In this work, the hyperparameter tuning was done only once for the 8-month testing set because of the computational burden it represents, but the tuning could be done more frequently. Market conditions evolve continuously, and static hyperparameter settings may degrade over time. A fully online reinforcement learning-based tuning mechanism would be ideal but is often impractical in real-world settings. Instead, lightweight adaptive strategies such as population-based training (PBT) offer a more feasible compromise, enabling incremental hyperparameter adaptation during operation [72].

6.3 Future Research

The work presented in this Master's thesis opens several avenues for future research. We intend to extend our backtesting to include more diverse datasets, such as other markets and products, to validate the generalizability, and the effectiveness of our approach.

A particularly promising direction is the extension of the `DRO-CVaR` methodology to congestion hedging instruments, such as financial transmission rights (FTRs). This is the focus of

the author's ongoing and future research. The FTR-focused research would integrate DC optimal power flow simulations to construct a distribution of congestion prices under diverse operating conditions. A regularization method is needed to ensure continuity of the congestion price function with respect to system conditions, then enabling the use of the continuous mapping theorem and justifying the convergence of the empirical price distribution to its true distribution. Building on this theoretical foundation, a distributionally robust optimization model could be developed to determine optimal monthly FTR positions, incorporating the same principles of robustness and risk aversion used in the present Master's thesis. The test system, inspired by a 300-bus representation of NYISO, demonstrates that such an OPF-derived approach can provide a promising basis for congestion hedging.

REFERENCES

[1] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2019.

[2] Saud Almahdi and Steve Y Yang. An adaptive portfolio trading system: A risk-return portfolio optimization using recurrent reinforcement learning with expected maximum drawdown. *Expert Systems with Applications*, 87:267–279, 2017.

[3] Adriano Arrigo, Christos Ordoudis, Jalal Kazempour, Zacharie De Grève, Jean-François Toubeau, and François Vallée. Wasserstein distributionally robust chance-constrained optimization for energy and reserve dispatch: An exact and physically-bounded formulation. *European Journal of Operational Research*, 296(1):304–322, 2022.

[4] Charles Audet and Warren Hare. *Derivative-Free and Blackbox Optimization*. Springer Series in Operations Research and Financial Engineering. Springer Cham, 1 edition, 2017. Hardcover ISBN: 978-3-319-68912-8 (2017-12-13), Softcover ISBN: 978-3-319-88680-0 (2018-09-04).

[5] Xavier Audet, Kliti Qako, and Antoine Lesage-Landry. A distributionally robust optimization strategy for virtual bidding in two-settlement electricity markets. *Sustainable Energy, Grids and Networks*, page 101904, 2025.

[6] Sevi Baltaoglu, Lang Tong, and Qing Zhao. Algorithmic bidding for virtual trading in electricity markets. *IEEE Transactions on Power Systems*, 34(1):535–543, 2018.

[7] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. *Advances in Neural Information Processing Systems*, 24, 2011.

[8] Dimitris Boskos, Jorge Cortés, and Sonia Martínez. High-confidence data-driven ambiguity sets for time-varying linear systems. *IEEE Transactions on Automatic Control*, 69(2):797–812, 2023.

[9] California Independent System Operator. California ISO, 2024. Accessed: 2024-10-14.

[10] Yongsheng Cao, Demin Li, Yihong Zhang, Qinghua Tang, Amin Khodaei, Hongliang Zhang, and Zhu Han. Optimal energy management for multi-microgrid under a trans-

active energy framework with distributionally robust optimization. *IEEE Transactions on Smart Grid*, 13(1):599–612, 2021.

[11] Metin Celebi, Attila Hajos, and Philip Q Hanser. Virtual bidding: The good, the bad and the ugly. *The Electricity Journal*, 23(5):16–25, 2010.

[12] Hong Chen. Security constrained economic dispatch (SCED) overview. Technical report, PJM Interconnection and Alberta Electric System Operator, n.d. Prepared for AESO.

[13] Ruidi Chen and Ioannis Ch Paschalidis. Distributionally robust learning. *Foundations and Trends® in Optimization*, 4(1-2):1–243, 2020.

[14] Yonghong Chen, Feng Pan, Feng Qiu, Alinson S Xavier, Tongxin Zheng, Muhammad Marwali, Bernard Knueven, Yongpei Guan, Peter B Luh, Lei Wu, et al. Security-constrained unit commitment for electricity market: Modeling, solution methods, and future challenges. *IEEE Transactions on Power Systems*, 38(5):4668–4681, 2022.

[15] Jaehyung Choi, Hyangju Kim, and Young Shin Kim. Diversified reward-risk parity in portfolio construction. *Studies in Nonlinear Dynamics & Econometrics*, 2024.

[16] Gonzalo Constante-Flores, Antonio J. Conejo, and S. Constante-Flores. Solving certain complementarity problems in power markets via convex programming. *Top*, 30(3):465–491, 2022.

[17] Peter Cramton. Electricity market design. *Oxford Review of Economic Policy*, 33(4):589–612, 2017.

[18] Ningning Du, Yankui Liu, and Ying Liu. A new data-driven distributionally robust portfolio optimization method based on Wasserstein ambiguity set. *IEEE Access*, 9:3174–3194, 2020.

[19] E. Ela, M. Milligan, A. Bloom, A. Botterud, A. Townsend, and T. Levin. Evolution of wholesale electricity market design with increasing levels of renewable generation. Technical report, National Renewable Energy Laboratory (NREL), Golden, CO., 09 2014.

[20] Brent Eldridge, Richard O'Neill, and Anya Castillo. An improved method for the DCOPF with losses. *IEEE Transactions on Power Systems*, 33(4):3779–3788, 2017.

[21] Electric Reliability Council of Texas. ERCOT, 2024. Accessed: 2024-10-14.

[22] Empyrical Library Documentation. Empyrical: Common financial risk and performance metrics, 2014. Accessed: 2024-09-12.

[23] Philip Ernst, James Thompson, and Yinsen Miao. Portfolio selection: The power of equal weight. *arXiv preprint arXiv:1602.00782*, 2016.

[24] National Grid ESO. National Grid ESO: Electricity Market Operations in the UK, 2025. Accessed: 2025-02-27.

[25] Federal Energy Regulatory Commission. Promoting wholesale competition through open access non-discriminatory transmission services by public utilities; recovery of stranded costs by public utilities and transmitting utilities, order no. 888, 1996.

[26] Federal Energy Regulatory Commission. Regional transmission organizations, order no. 2000, 1999.

[27] Federal Energy Regulatory Commission. RTOs and ISOs. `https://www.ferc.gov/power-sales-and-markets/rtos-and-isos`, 2025. Accessed: 2025-11-08.

[28] Rémi Galarneau-Vincent, Geneviève Gauthier, and Frédéric Godin. Foreseeing the worst: Forecasting electricity DART spikes. *Energy Economics*, 119:106521, 2023.

[29] Matteo Gelmini and Pierpaolo Uberti. The equally weighted portfolio still remains a challenging benchmark. *International Economics*, 179:100525, 2024.

[30] Xuejiao Han and Gabriela Hug. A distributionally robust bidding strategy for a wind-storage aggregator. *Electric Power Systems Research*, 189:106745, 2020.

[31] William W Hogan. Virtual bidding and electricity market design. *The Electricity Journal*, 29(5):33–47, 2016.

[32] ISO New England. ISO-NE, 2024. Accessed: 2024-10-14.

[33] Rajnish Kamat and Shmuel S Oren. Multi-settlement systems for electricity markets: Zonal aggregation under network uncertainty and market power. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences*, pages 739–748. IEEE, 2002.

[34] Jalal Kazempour and Benjamin F Hobbs. Value of flexible resources, virtual bidding, and self-scheduling in two-settlement electricity markets with wind generation—Part I: Principles and competitive model. *IEEE Transactions on Power Systems*, 33(1):749–759, 2017.

[35] Simge Küçükyavuz and Ruiwei Jiang. Chance-constrained optimization under limited distributional information: A review of reformulations based on sampling and distributional robustness. *EURO Journal on Computational Optimization*, 10:100030, 2022.

[36] Daniel Kuhn, Soroosh Shafiee, and Wolfram Wiesemann. Distributionally robust optimization. *Acta Numerica*, 34:579–804, 2025.

[37] Yinglun Li. *Data-driven modeling and algorithmic trading in electricity market*. PhD thesis, UC Riverside, 2024.

[38] Yinglun Li, Nanpeng Yu, and Wei Wang. Machine learning-driven virtual bidding with electricity market efficiency analysis. *IEEE Transactions on Power Systems*, 37(1):354–364, 2021.

[39] Eugene Litvinov. Design and operation of the locational marginal prices-based electricity markets. *IET Generation, Transmission & Distribution*, 4(2):315–323, 2010.

[40] Eugene Litvinov, Tongxin Zheng, Gary Rosenwald, and Payman Shamsollahi. Marginal loss modeling in LMP calculation. *IEEE transactions on Power Systems*, 19(2):880–888, 2004.

[41] Adriana A Londoño and Juan D Velásquez. Risk management in electricity markets: Dominant topics and research trends. *Risks*, 11(7):116, 2023.

[42] Malik Magdon-Ismail and Amir F Atiya. Maximum drawdown. *Risk Magazine*, 17(10):99–102, 2004.

[43] Hossein Mehdipourpicha, Siyuan Wang, and Rui Bo. Developing robust bidding strategy for virtual bidders in day-ahead electricity markets. *IEEE Open Access Journal of Power and Energy*, 8:329–340, 2021.

[44] Midcontinent Independent System Operator. MISO, 2024. Accessed: 2024-10-14.

[45] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.

[46] Francesco Morri, Hélène Le Cadre, Luce Brotcorne, and Pierre Gruet. Learning in Stackelberg Games with Application to Strategic Bidding in the Electricity Market. In *2024 20th International Conference on the European Energy Market (EEM)*, pages 1–7. IEEE, 2024.

[47] Antonio Naimoli and Giuseppe Storti. Forecasting volatility and tail risk in electricity markets. *Journal of Risk and Financial Management*, 14(7):294, 2021.

[48] New York Independent System Operator. NYISO, 2024. Accessed: 2024-10-14.

[49] New York Independent System Operator. NYISO Market Administration and Control Area Services Tariff (MST), 2025. Accessed: 2025-11-20.

[50] New York Independent System Operator. NYISO Open Access Transmission Tariff (OATT), 2025. Accessed: 2025-11-20.

[51] OMI. Iberian Peninsula Electricity Market: OMI Overview, 2025. Accessed: 2025-02-27.

[52] Julien Pallage and Antoine Lesage-Landry. Wasserstein Distributionally Robust Shallow Convex Neural Networks. *INFORMS Journal on Optimization*, 2025.

[53] Dimitri J Papageorgiou. Data-driven distributionally robust optimization for long-term contract vs. spot allocation decisions: Application to electricity markets. *Computers & Chemical Engineering*, 179:108436, 2023.

[54] John E Parsons, Cathleen Colbert, Jeremy Larrieu, Taylor Martin, and Erin Mastrangelo. Financial arbitrage and efficient dispatch in wholesale electricity markets. *MIT Center for Energy and Environmental Policy Research*, (15-002), 2015.

[55] Pierre Pinson. Distributionally robust trading strategies for renewable energy producers. *IEEE Transactions on Energy Markets, Policy and Regulation*, 1(1):37–47, 2023.

[56] PJM Interconnection LLC. PJM, 2024. Accessed: 2024-10-14.

[57] Yuliya Plyakha, Raman Uppal, and Grigory Vilkov. Why does an equal-weighted portfolio outperform value-and price-weighted portfolios? *Available at SSRN 2724535*, 2012.

[58] Bala Kameshwar Poolla, Ashish R Hota, Saverio Bolognani, Duncan S Callaway, and Ashish Cherukuri. Wasserstein distributionally robust look-ahead economic dispatch. *IEEE Transactions on Power Systems*, 36(3):2010–2022, 2020.

[59] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.

[60] R. Tyrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

[61] Ehsan Samani, Mahdi Kohansal, and Hamed Mohsenian-Rad. A data-driven convergence bidding strategy based on reverse engineering of market participants' performance: A case of California ISO. *IEEE Transactions on Power Systems*, 37(3):2122–2136, 2021.

[62] William F Sharpe. The Sharpe Ratio. *Streetwise–the Best of the Journal of Portfolio Management*, 3(3):169–85, 1998.

[63] Southwest Power Pool. SPP, 2024. Accessed: 2024-10-14.

[64] Dirk Tasche. Expected shortfall and beyond. *Journal of Banking & Finance*, 26(7):1519–1533, 2002.

[65] Evangelos G Tsimopoulos, Christos N Dimitriadis, and Michael C Georgiadis. Financial arbitrage in electricity markets via virtual bidding. *Computers & Chemical Engineering*, 181:108550, 2024.

[66] U.S. Energy Information Administration. NYISO load zone map. `https://www.eia.gov/electricity/930-content/NYISO-load-zone-map-WEB-062119.jpg`, 2019. Accessed: 2025-11-25.

[67] Yushen Wang, Weiliang Huang, Haoyong Chen, Zhiwen Yu, Linlin Hu, and Yuxiang Huang. Optimal scheduling strategy for virtual power plants with aggregated user-side distributed energy storage and photovoltaics based on CVaR-distributionally robust optimization. *Journal of Energy Storage*, 86:110770, 2024.

[68] Dongliang Xiao, Hongbo Sun, Daniel Nikovski, Shoichi Kitamura, Kazuyuki Mori, and Hiroyuki Hashimoto. CVaR-constrained stochastic bidding strategy for a virtual power plant with mobile energy storages. In *2020 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, pages 1171–1175. IEEE, 2020.

[69] Weijun Xie and Shabbir Ahmed. Distributionally robust chance constrained optimal power flow with renewables: A conic reformulation. *IEEE Transactions on Power Systems*, 33(2):1860–1867, 2017.

[70] Ogun Yurdakul, Fikret Sivrikaya, and Sahin Albayrak. A distributionally robust optimization approach to unit commitment in microgrids. In *2021 IEEE Power & Energy Society General Meeting (PESGM)*, pages 1–5. IEEE, 2021.

[71] Jianzhe Zhen, Daniel Kuhn, and Wolfram Wiesemann. Mathematical foundations of robust and distributionally robust optimization. *arXiv preprint arXiv:2105.00760*, 2021.

[72] Yinda Zhou, Weiming Liu, and Bin Li. Efficient online hyperparameter adaptation for deep reinforcement learning. In *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*, pages 141–155. Springer, 2019.

[73] Zhongbao Zhou, Zhengyang Song, Tiantian Ren, and Lean Yu. Two-stage portfolio optimization integrating optimal Sharp ratio measure and ensemble learning. *IEEE Access*, 11:1654–1670, 2022.