A Machine Learning Framework for Algorithmic Trading with Virtual Bids in Electricity Markets

Wei Wang
Department of Electrical and Computer Engineering
University of California, Riverside
Riverside, CA 92521
Email: wwang031@ucr.edu

Nanpeng Yu
Department of Electrical and Computer Engineering
University of California, Riverside
Riverside, CA 92521
Email: nyu@ece.ucr.edu

***Abstract*—This paper develops a machine learning framework for algorithmic trading with virtual bids in electricity markets. In the proposed algorithmic trading strategy, a budget and risk constrained portfolio optimization problem is solved, which selects the virtual transactions to be executed. In order to maximize the expected net earnings of the virtual bid portfolio, a mixture density network model is developed to provide robust and accurate forecasts for electricity price spread between day-ahead and real-time market. By leveraging a coherent risk measure and historical price samples, the risk-constrained portfolio optimization problem is solved efficiently. Backcasting results based on market data from ISO New England show that our proposed mixture density network based trading strategy consistently outperforms the benchmark online learning approach.**

***Index Terms*—Algorithmic trading, electricity market, machine learning, portfolio optimization, virtual transaction**

I. Introduction

The regional wholesale electricity markets in the United States all adopt the two-settlement system, which includes the day-ahead (DA) market and the real-time (RT) market. The DA market is a forward market that determines the hourly DA locational marginal prices (LMPs), the unit commitment plans, and the DA dispatch schedules for generations and dispatchable loads. The decisions are based on the supply offers and the demand bids submitted by load serving entities (LSEs), independent power producers (IPPs), and proprietary trading firms. In the RT market, the RT LMPs and the unit/resource dispatch schedules are calculated based on the updated supply offers and the actual operating conditions described by the state estimator.

Market participants can buy or sell energy in the DA market with an explicit requirement to sell or buy it back in the RT market using virtual bids. Note that there is no requirement for such virtual bids to be backed by physical assets. There are two types of virtual bids, increment offer (INC) and decrement bid (DEC) [1]. INC (DEC) is a financial instrument that enables virtual bidders to sell (buy) energy in the DA market and buy (sell) the same amount of energy back in the RT market at the same pricing node [2].

Virtual transactions were introduced into the two-settlement electricity market to improve market efficiency, promote price convergence [3], provide hedging instruments [4], and enhance market liquidity. IPPs and LSEs could leverage DECs to hedge risks associated with generator forced outage, higher RT electric load and volatile RT LMPs. Proprietary trading firms use virtual bids to arbitrage in the electricity market when the expected RT LMPs differ from the expected DA LMPs.

To further drive price convergence in the electricity market, improve market efficiency and increase profitability, it is imperative for proprietary trading firms to design virtual bids portfolio trading strategies that maximize the expected earnings and minimize risks. Virtual traders in proprietary trading firms regularly place speculative virtual bids to arbitrage the differences between DA and RT LMPs based on their knowledge of the electricity market and the forecasts for the key factors that influence electricity prices. In fact, the cleared virtual transactions in the five major electricity markets in the U.S. was 13% [1] of the total load. In this paper, we aim to develop a machine learning framework for algorithmic trading with virtual bids. We are interested in exploring if the machine learning based approach could find profitable virtual trading strategies in electricity markets.

Only a few researchers have studied portfolio trading strategies with virtual transactions from the perspective of proprietary trading firms. In [5], the portfolio optimization problem with virtual bids is formulated as a multi-armed bandit problem and solved by the algorithm referred to as dynamic programming on discrete set (DPDS). It was shown that DPDS consistently outperforms benchmark heuristic methods [6] when only the historical LMPs are considered. A risk-constrained virtual bids portfolio trading strategy is developed to empirically test for the efficiency of the California Independent System Operator market [2]. The existence of a profitable trading strategy with virtual transactions is explored via hypothesis testing in [3].

This paper pushes the research frontier of algorithmic trading with virtual bids in the following ways. First, instead of relying only on historical LMPs to model and estimate the DA and RT price spreads and payoff of virtual bids, this paper develops a mixture density network (MDN) to infer the conditional distribution of nodal price spreads given the fundamental inputs such as electric load, generation outage, and transmission outage. Second, a risk-constrained portfolio optimization problem for virtual bids is formulated and efficiently solved with a finite number of scenarios. Third, a machine learning framework for algorithmic trading with virtual bids is established by synergistically combining the risk-constrained portfolio optimization framework and the MDN model. The profitability of our proposed trading strategy outperforms the state-of-the-art online learning (OL) approach to virtual trading [5], [6].

The remainder of the paper is organized as follows. Section II lays out the overall framework for portfolio optimization with virtual bids. The technical methods for estimating the price spreads and expected earnings of virtual transactions is presented in Section III. The solution methodology for risk-constrained optimization is described in Section IV. Section V shows the results of numerical study with the data set from ISO New England (ISO-NE). Section IV concludes the paper.

II. Portfolio Optimization with Virtual Bids

A. Modeling of Virtual Bids

In this subsection, we will model two types of virtual transactions, incremental offers (INCs) and decremental bids (DECs) in detail.

Let $\lambda_{i,h}^{DA}$ and $\lambda_{i,h}^{RT}$ denote the DA LMP and the RT LMP for node i and hour h. Define $\lambda_{i,h}^{dif} = \lambda_{i,h}^{RT} - \lambda_{i,h}^{DA}$ as the price spread between the DA LMP and the RT LMP. Let $\lambda_{i,h}^{bid,I}$ and $\lambda_{i,h}^{bid,D}$ represent the bid price ($/MWh) for INCs and DECs for node i and hour h.

INCs are also called virtual supply offers, which sell energy in the DA market and must buy back the same amount of energy in the RT market [7]. A virtual supply offer will be cleared if the bid price is lower than the DA LMP. The expected earning of an INC $E[r_{i,h}^I]$ is

$$E[r_{i,h}^I] = E\big[-\lambda_{i,h}^{dif}\mathbb{1}(\lambda_{i,h}^{bid,I} \le \lambda_{i,h}^{DA})\big] \tag{1}$$

DECs are also called virtual demand bids, which buy energy in the DA market and must sell the same amount of energy in the RT market as a price-taker. A virtual demand bid will be cleared if the bid price is higher than the DA LMP. The expected earning of a DEC $E[r_{i,h}^D]$ is

$$E[r_{i,h}^D] = E\big[\lambda_{i,h}^{dif}\mathbb{1}(\lambda_{i,h}^{bid,D} \ge \lambda_{i,h}^{DA})\big] \tag{2}$$

B. Budget and Risk Constrained Portfolio Optimization

Before the DA market closes, each day a trader determines a portfolio of virtual bids to be submitted to the DA market which maximizes the expected portfolio earnings subject to budget and risk constraints as follows:

$$\max_{\boldsymbol{z}} \sum_{i=1}^{N}\sum_{h=1}^{24} E[r_{i,h}^I z_{i,h}^I + r_{i,h}^D z_{i,h}^D] \tag{3}$$

s.t.

$$\sum_{i=1}^{N}\sum_{h=1}^{24}[Prox_{(i,h)}^I z_{i,h}^I + Prox_{(i,h)}^D z_{i,h}^D] \le B \tag{4}$$

$$\sum_{h=1}^{24} CVaR_\alpha(f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}})) \le C \tag{5}$$

$$f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}}) = -\sum_{i=1}^{N}[r_{i,h}^I z_{i,h}^I + r_{i,h}^D z_{i,h}^D] \tag{6}$$

$$z_{i,h}^I, z_{i,h}^D \in \{0,1\} \tag{7}$$

$z_{i,h}^I$ and $z_{i,h}^D$ are the binary variables indicating whether or not the corresponding virtual bids with a quantity of 1MWh are selected for submission. The portfolio loss function f_h can be calculated as the summation of the loss of individual virtual bids (6), where $\boldsymbol{z_h}$ is the binary decision vector for the submission of virtual bids for all nodes at hour h and $\boldsymbol{\lambda_h^{dif}}$ is the vector of random price spreads for all nodes at hour h. Equation (4) is the budget constraint, where $Prox_{(i,h)}^I$ and $Prox_{(i,h)}^D$ denotes the DA virtual bid financial assurance proxy for INC and DEC at node i and hour h respectively, and B is the total portfolio budget limit. The virtual proxy is utilized by ISO for calculating the financial assurance requirements for virtual transactions [8]. Equation (5) enforces the portfolio risk constraint, where the risk metric is selected as the conditional value-at-risk (CVaR) with confidence level α [9]. The details about estimating the expected earnings of virtual transactions and risk-constrained optimization are described in Section III and IV respectively.

III. Estimation of Expected Earnings of Virtual Transactions

A. Benchmark Algorithm: An Online Learning Approach

The Online Learning (OL) approach proposed in [5] and [6] is used as the benchmark algorithm in this paper. The OL approach assumes that traders only have access to historical DA LMPs, RT LMPs, and bid prices when determining the optimal virtual bids. In addition, the payoff of a virtual bid with a given bidding price has the same distribution as that of the historical bids. To make the OL approach consistent with the portfolio optimization framework presented in (3)-(7), we add the portfolio risk constraint and modify the portfolio budget constraint. Note that most of the ISOs in the U.S. use the virtual proxy for the budget constraint as in equation (4) rather than the bid price as shown in [5].

In this subsection, we follow the convention of notations in [5] and [6] to make the payoff function for INCs and DECs consistent. The DA, RT, and bid prices of INCs are transformed as $\overline{\lambda} - \lambda_{i,h}^{DA}$, $\overline{\lambda} - \lambda_{i,h}^{RT}$, and $\overline{\lambda} - \lambda_{i,h}^{bid,I}$ respectively, where $\overline{\lambda}$ is the upper bound of DA LMP. Similarly, the DA, RT and bid prices for DECs are transformed as $\lambda_{i,h}^{DA} - \underline{\lambda}$, $\lambda_{i,h}^{RT} - \underline{\lambda}$, and $\lambda_{i,h}^{bid,D} - \underline{\lambda}$ respectively, where $\underline{\lambda}$ is the lower bound of DA LMP. This way, the functional form of payoffs for the INCs is the same as that for the DECs.

In the OL approach, the expected payoff of a virtual bid with a given bid price is calculated as the average empirical payoff based on the historical DA and RT LMPs. The expected payoff is calculated sequentially and adaptively based on the new market information. The virtual transaction at each node and each hour is treated as a single product. Let $\boldsymbol{\lambda}_t^{(k)}$ denote the bidding price vector for the kth product from day 1 to day t with the bidding prices sorted in ascending order. Let $\boldsymbol{r}_t^{(k)}$

denote the vector of expected payoff for a virtual transaction corresponding to the bid price in $\boldsymbol{\lambda}_t^{(k)}$.

After receiving the new information about bidding price on day t for product k, $\boldsymbol{\lambda}_t^{(k)}$ can be updated based on $\boldsymbol{\lambda}_{t-1}^{(k)}$ as:

$$\boldsymbol{\lambda}_t^{(k)} = [\boldsymbol{\lambda}_{t-1}^{(k)}(1:i_k), \lambda_{t,k}, \boldsymbol{\lambda}_{t-1}^{(k)}(i_k+1:t-1)] \quad (8)$$

where $\lambda_{t,k}$, is inserted at the i_k+1th index with $i_k = \max_i \boldsymbol{\lambda}_{t-1}^{(k)}(i) < \lambda_{t,k}$. The symbol $(\cdot)$ denotes the operation of fetching elements by indexes.

The vector for corresponding expected payoff $\boldsymbol{r}_t^{(k)}$ can be updated based on $\boldsymbol{r}_{t-1}^{(k)}$ as:

$$\begin{aligned}\boldsymbol{r}_t^{(k)} =& [\frac{t-1}{t}\boldsymbol{r}_{t-1}^{(k)}(1:i_k),\\ &\frac{t-1}{t}\boldsymbol{r}_{t-1}(i_k:t-1) + \frac{1}{t}(\lambda_{t,k}^{RT} - \lambda_{t,k}^{DA})]\end{aligned} \quad (9)$$

where $\lambda_{t,k}^{RT}$ and $\lambda_{t,k}^{DA}$ are the DA LMP and the RT LMP for product k at day t respectively.

For a trader following the OL approach, the bid price needs to be selected for each product k. Obviously, the bid price with the highest expected payoff will be selected. The expected payoff of bidding product k can then be updated as:

$$E[r^k] = max(\boldsymbol{r}_t^{(k)}) \quad (10)$$

B. A Machine Learning Framework

Instead of relying only on historical DA LMPs, RT LMPs, and bid prices to estimate the expected price spreads and the payoffs of virtual transactions, the key features that influence the formation of DA and RT prices should also be included in the prediction model. For example, load forecast, generation capacity, transmission outages, fuel prices, and meteorological forecast are all crucial in determining the distribution of price spreads and payoffs of virtual transactions. A machine learning framework based on MDN is proposed in this subsection to infer the distribution of the price spreads and the expected payoffs for virtual bids based on these key features.

Note that under our proposed machine learning framework for algorithmic trading, the virtual bids are treated as self-schedules, i.e., the bidding price is not a decision variable. In other words, the bid prices of INCs are selected to be the price floor of the DA LMP and the bid prices for DECs are selected to be the price cap of the DA LMP.

1) Input Features: The DA and RT LMPs are determined through the market clearing processes which involve solving security-constrained unit commitment problems and security-constrained economic dispatch problems in the DA and RT markets. The key features that influence the price spreads between DA and RT LMPs can be categorized into three groups. The first group includes all the meteorological variables such as system-wide/zonal temperature, dew point, cloud cover, and wind speed. The second group includes all the relevant fuel prices for natural gas, coal, and diesel. The third group includes the system variables such as the forecast for system/zonal demand, the available generation capacity by fuel type, and transmission outages. Note that only the forecast for these variables will be used for choosing the virtual bids to be submitted to the DA market. All of these data are accessible from the Independent System Operator (ISO)'s website and other online repositories.

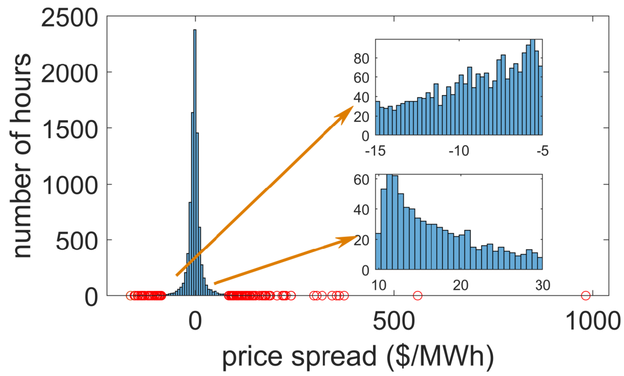

Fig. 1. Histogram of price spreads of a sample node in year 2015

2) Modeling Price Spread with Mixture Density Networks: The price spreads between DA and RT LMPs have extremely high volatility and spikiness [10]. For illustrative purposes, the histogram of the price spreads on a sample node in ISO New England (ISO-NE) is shown in Fig. 1. The frequency of occurrence for price spikes is not visible on the histogram. The price spikes are labeled by the red circles. The two zoomed-in subplots for price ranges $(10, 30)$ and $(-15, -5)$ show that the price spread has a multimodal distribution. The typical neural network model assumes that the output variable has a Gaussian distribution with a mean dependent on the input variables. This Gaussian assumption can lead to very poor price spread predictions.

To deal with multimodality, we adopt mixture density network (MDN) to model the conditional probability distributions of the price spreads. In MDN, a Gaussian mixture distribution is assumed for the conditional distribution of the price spread as follows:

$$p(\lambda^{dif}|X) = \sum_{c=1}^{N_c} \pi_c(X)\mathcal{N}(\lambda^{dif}|\mu_c(X), \sigma_c^2(X)) \quad (11)$$

where X and λ^{dif} denote the input features and the price spread. N_c is the total number of components. $\pi_c(X)$, $\mu_c(X)$, and $\sigma_c(X)$ are the feature dependent weight, mean value and standard deviation for component c respectively. The loss function of MDN is the negative logarithm of the likelihood function:

$$\begin{aligned}L(W) = -\sum_{n=1}^{N_s} \ln \sum_{c=1}^{N_c} &[\pi_c(x_n, W)\\ &\mathcal{N}(\lambda_n^{dif}|\mu_c(x_n, W), \sigma_c^2(x_n, W))]\end{aligned} \quad (12)$$

where N_s is the total number of data samples. x_n and λ_n^{dif} are the features and the price spread target of sample n, while W denotes the network parameters. In order to make the summation of weights equal to 1, the softmax activation function is applied to the output layers for π_c. To make the estimate for standard deviation positive, the exponential activation function is adopted for the output layer of σ_c.

3) Data Preprocessing: Data preprocessing is crucial to minimize noise, capture nonlinear relationships, and flatten the distribution of variables which makes the learning process of neural networks more efficient and robust. Given that the spreads between the DA and RT LMPs are extremely volatile and spiky, the sigmoid function $f(x) = 1/(1 + exp(-x/\theta))$ is leveraged to scale the output price spread data and flatten the price spread distribution. The sigmoidal normalization is especially appropriate to preprocess price spread data, which has many large spikes. The normalization reduces the number of components needed in MDN to avoid over-fitting. Batch normalization is applied to the inputs.

IV. Risk-Constrained Portfolio Optimization

The risk-constrained portfolio optimization problem (3)-(7) formulated in Section II is a mixed integer convex optimization problem. This is because the our chosen coherent risk measure [11], CVaR of the portfolio, is a convex function with respect to the positions in virtual bids. In this section, we will convert the mixed integer convex optimization problem into a mixed integer linear programming (MILP) problem.

First, we briefly review the definition of two risk metrics. The probability of the portfolio loss $f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}})$ at hour h not exceeding ζ_h is defined as

$$\Psi(\boldsymbol{z_h}, \zeta_h) = \int_{f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}}) \leq \zeta_h} p(\boldsymbol{\lambda_h^{dif}}) d\boldsymbol{\lambda_h^{dif}} \tag{13}$$

The value at risk (VaR), i.e., the probability of loss not exceeding ζ_h with confidence level α is defined as

$$VaR_\alpha(\boldsymbol{z_h}) = \min\{\zeta_h : \Psi(\boldsymbol{z_h}, \zeta_h) \geq \alpha\} \tag{14}$$

The $CVaR_\alpha$ of the portfolio loss can then be defined as

$$CVaR_\alpha(f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}})) = \frac{1}{1-\alpha} \int_{f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}}) \geq VaR_\alpha(\boldsymbol{z_h})} f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}}) p(\boldsymbol{\lambda_h^{dif}}) d\boldsymbol{\lambda_h^{dif}} \tag{15}$$

It has been proven that function $F_\alpha(\boldsymbol{z_h}, \zeta_h)$ is an upper bound of $CVaR_\alpha$ [12].

$$F_\alpha(\boldsymbol{z_h}, \zeta_h) = \frac{1}{1-\alpha} \int_{\lambda_h} [f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}}) - \zeta_h]^+ p(\boldsymbol{\lambda_h^{dif}}) d\boldsymbol{\lambda_h^{dif}} \tag{16}$$

By using the historical LMP samples, $F_\alpha(\boldsymbol{z_h}, \zeta_h)$ can be further simplified [13] as

$$F_\alpha(\boldsymbol{z_h}, \zeta_h) = \zeta_h + \frac{1}{(1-\alpha)N_s} \sum_{j=1}^{N_s} [f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_{h,j}^{dif}}) - \zeta_h]^+ \tag{17}$$

where N_s is the total number of samples. By introducing dummy variable u_h^j for sample j at hour h, equation (17) can be transformed as:

$$F_\alpha(\boldsymbol{u_h}, \zeta_h) = \zeta_h + \frac{1}{(1-\alpha)N_s} \sum_{j=1}^{N_s} u_h^j \tag{18}$$

$$u_h^j \geq f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_{h,j}^{dif}}) - \zeta_h \tag{19}$$

$$u_h^j \geq 0 \tag{20}$$

Hence, it can be shown [12] that CVaR of the portfolio loss can be determined from the formula below:

$$CVaR_\alpha(f_h(\boldsymbol{z_h}, \boldsymbol{\lambda_h^{dif}})) = \min_{\zeta_h} F_\alpha(\boldsymbol{z_h}, \zeta_h) \tag{21}$$

Therefore, the original risk-constrained virtual transactions portfolio optimization problem can be reformulated as

$$\max_{\boldsymbol{z}, \boldsymbol{\zeta}, \boldsymbol{u}} \sum_{i=1}^{N} \sum_{h=1}^{24} E[r_{i,h}^I z_{i,h}^I + r_{i,h}^D z_{i,h}^D] \tag{22}$$

s.t.

$$\sum_{h=1}^{24} F_\alpha(\boldsymbol{u_h}, \zeta_h) \leq C \tag{23}$$

$$(4), (6), (7), (18) - (20)$$

The optimization problem is now a MILP which can be tackled by commercial solvers such as Gurobi and CPLEX.

V. Numerical Study

The performances of the proposed virtual transaction bidding strategy based on MDN and the benchmark OL approach are evaluated with the electricity market managed by ISO-NE. The historical LMPs and input variables such as oil and gas prices forecasts, total system demand forecast, wind generation forecast, and weather forecasts, are taken from online data repositories [14], [15]. The DA and RT LMPs from the 994 pricing nodes in ISO-NE are collected. Three years of historical data from year 2015 to 2017 are gathered. The first year's data is used for training and validation purpose. The rest of the data are used for out-of-sample testing. In the testing process, the proposed MDN model is updated on a weekly basis. The costs associated with virtual bids in ISO-NE include transaction fee and net commitment period compensation (NCPC). NCPC is a payment to generators, dispatchable-asset-related demands, demand response resources or the external transactions that did not recover their effective offer costs from the energy market. The transaction fee for virtual bids is \$0.065/MWh. The NCPCs are \$1.25/MWh in 2016 and \$0.77/MWh in 2017 for ISO-NE.

A. Profitability of Algorithmic Trading Strategies

Both the OL approach and the proposed MDN based trading strategy are used to select portfolios of virtual transactions in ISO-NE on a daily basis between January 1, 2016 and December 31, 2017. Two sets of daily portfolio budget and risk constraints of \$50k and \$100k are tested. The confidence level α of CVaR is set as 0.95. The gain, the trading costs, and the net gains are reported in Table I. It can be seen from the table that our proposed MDN based trading strategy consistently outperforms the OL approach in both years and under both sets of budget and risk constraints. When transaction costs and NCPC are taken into consideration, the net gains of our proposed trading strategy are \$1.17 million and \$2.38 million in 2016 and 2017 with a \$100k risk and budget limit. The trading performance results suggest that the MDN provides more accurate prediction for price spreads than historical averages calculated by the OL approach. Finally, note that when the daily budget and risk limit increase, the net gains of the proposed MDN based trading strategy also increase.

TABLE I
PERFORMANCE COMPARISON OF ALGORITHMIC TRADING STRATEGIES

Strategy	Year	Budget / Risk (\$)	Gain (million \$)	Fees & Uplift (million \$)	Net Gain (million \$)
OA	2016	50k / 50k	1.3375	1.2211	0.1164
		100k / 100k	1.8406	2.2876	−0.4469
	2017	50k / 50k	1.0506	0.7149	0.3356
		100k / 100k	1.6967	1.3496	0.3472
MDN	2016	50k / 50k	2.2359	1.4087	0.8271
		100k / 100k	3.9504	2.7804	1.1701
	2017	50k / 50k	2.4221	0.8806	1.5415
		100k / 100k	4.1088	1.7319	2.3770

B. Influence of Rare Events and Changes in Market Rules

The cumulative net gains of our proposed and benchmark virtual bids trading strategies with \$100k daily budget and risk limit for year 2016 and 2017 are depicted in Figure 2. As shown in the figure, the performance of the proposed and benchmark trading strategy can be underwhelming during certain events/periods. It indicates that rare market events and changes in market operation rules not captured by the machine learning model have a significant impact on the performance of the trading algorithms.

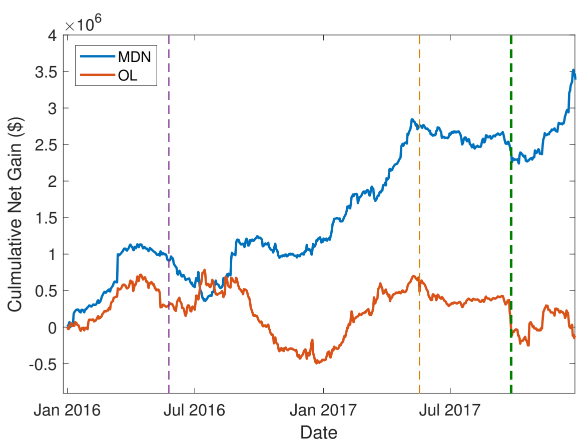

Fig. 2. Cumulative net gains of the proposed and the benchmark trading strategies from 2016 to 2017 with \$100k daily budget and risk limit

The first slump of MDN algorithm in trading performance is due to the implementation of Do-Not-Exceed (DNE) dispatch rules by ISO-NE on May 25, 2016 (labeled by the purple dash line). Before the change in market rules, the practice of manual curtailment for renewable generation resource was adopted by system operators. Under the DNE dispatch rules, renewable generation units must submit supply offers into the DA market. These offers are allowed to settle the DA LMPs on the corresponding pricing nodes when congestion happens. It can be seen that the cumulative net gain of the proposed trading strategy takes a dip right after DNE's implementation. The sharp dip in net gain of the benchmark algorithm on May 18, 2017 (labeled by the orange dash line) and the big drops of MDN algorithm between September 24 and September 27, 2017 (labeled by the green dash lines) are caused by price spikes in RT. These price spikes result from forced outages in RT, which are difficult to predict on a DA basis.

VI. CONCLUSION

This paper develops a data-driven algorithmic trading strategy for virtual transactions in electricity markets. A budget and risk constrained portfolio optimization problem is formulated to select the virtual bids to be submitted. A mixture density network model is developed to forecast the price spreads between DA and RT LMPs. Backcasting results with ISO-NE's market data show that the proposed MDN based algorithmic trading strategy is much more profitable than the benchmark OL approach. In the future, we will consider the price sensitivity in the portfolio optimization framework and develop a machine learning framework to jointly model the price spreads of all pricing nodes.

REFERENCES

[1] J. E. Parsons, C. Colbert, J. Larrieu, T. Martin, and E. Mastrangelo, "Financial arbitrage and efficient dispatch in wholesale electricity markets," *MIT Center for Energy and Environmental Policy Research*, no. 15-002, 2015. [Online]. Available: https://ssrn.com/abstract=2574397

[2] R. Li, A. J. Svoboda, and S. S. Oren, "Efficiency impact of convergence bidding in the California electricity market," *Journal of Regulatory Economics*, vol. 48, no. 3, pp. 245–284, Dec. 2015.

[3] F. A. W. Akshaya Jha, "Testing for market efficiency with transactions costs: An application to convergence bidding in wholesale electricity markets," *Industrial Organization Seminar,Yale University*, May 2013.

[4] W. W. Hogan, "Virtual bidding and electricity market design," *The Electricity Journal*, vol. 29, pp. 33–47, Jun. 2016.

[5] S. Baltaoglu, L. Tong, and Q. Zhao, "Online learning of optimal bidding strategy in repeated multi-commodity auctions," in *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017, vol. 30, pp. 4507–4517.

[6] ——, "Algorithmic bidding for virtual trading in electricity markets," *IEEE Transactions on Power Systems*, pp. 1–1, 2018.

[7] ISO New England Inc. Overview of new englands wholesale electricity markets and market oversight. [Online]. Available: https://www.iso-ne.com/pubs/spcl_rpts/2014/2014_market_overview_050614.pdf

[8] ——. ISO New England financial assurance policy. [Online]. Available: https://www.iso-ne.com/static-assets/documents/2017/09/sect_i_ex_ia.pdf

[9] N. Yu, A. Somani, and L. Tesfatsion, "Financial risk management in restructured wholesale power markets: Concepts and tools," in *IEEE PES General Meeting*, July 2010, pp. 1–8.

[10] N. Yu, H. Sheng, and R. Johnson, "Economic valuation of wind curtailment rights," in *IEEE Power and Energy Society General Meeting (PES)*, 2013, pp. 1–5.

[11] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Mathematical Finance*, vol. 9, no. 3, pp. 203–228, 1999.

[12] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–41, 2000.

[13] P. Krokhmal, J. Palmquist, and S. Uryasev, "Portfolio optimization with conditional value-at-risk objective and constraints," *Journal of Risk*, vol. 4, May 2003.

[14] U.S. Energy Information Administration. [Online]. Available: https://www.eia.gov/

[15] ISO New England Inc. ISO express. [Online]. Available: https://www.iso-ne.com/markets-operations/iso-express/