Windy with a Chance of Profit - Bid Strategy and Analysis for Wind Integration

Sagar Kurandwad, Chandrasekar Subramanian,
Venkata Ramakrishna P, Arunchandar Vasan,
Venkatesh Sarangan, Vijaysekhar Chellaboina
Innovation Labs,
TATA Consultancy Services
IIT Madras Research Park, Chennai 600113, India
venkataramakrishna.p@tcs.com

Anand Sivasubramaniam
Dept. of Comp. Sci. & Eng.
Pennsylvania State University
University Park, PA 16802, USA
anand@cse.psu.edu

ABSTRACT

Integration of wind power with the grid has become an important problem. For integration, a producer needs to bid in a time-ahead market to deliver an amount of energy at a future point in time. Because wind speed and price are both uncertain, a producer needs to place bids on the basis of expected wind power yield and price. To this end, improving the accuracy of the prediction of wind speed has received much attention. However, the trade-off between expected profit and the prediction errors over a multi-period setting has been less studied.

We fill this gap by quantifying trade-offs between profits and prediction errors. First, we obtain, under idealized conditions on the price and the yield processes, an optimal bid strategy as a closed-form expression. Next, we evaluate the profit-vs-prediction trade-off using this idealized bidding strategy on synthetic datasets which satisfy all the idealistic assumptions. We also consider two baselines - a naive strategy and an oracle strategy that has perfect knowledge over a limited horizon. Finally, we relax our assumptions and evaluate all strategies under real-world datasets. We identify and work around limitations of the idealized bidding strategy when the underlying assumptions are violated.

On synthetic datasets, with no buffering and a (relative) prediction error of 25% , we find that our bidding approach performs significantly better than a naive approach and compares favourably (86%) to an oracle with a look-ahead of two time-slots and infinite buffer. On real-world datasets, with buffer equivalent to 20% of the maximum yield, our approach exceeds the naive approach by 25%, while remaining within 62% of a two-step look-ahead oracle that uses infinite buffering.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Stochastic processes; J.7 [**Computers in Other Systems**]

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
e-Energy'14, June 11–13, 2014, Cambridge, UK.
Copyright 2014 ACM 978-1-4503-2819-7/14/06 ...$15.00.
http://dx.doi.org/10.1145/2602044.2602056.

Keywords

Renewable integration; Wind prediction; Bidding strategy

1. INTRODUCTION

The global demand for energy is expected to grow by 56% from 2010 to 2040 [20]. With increasing awareness of the carbon emissions of fossil-fuel based power generation such as coal, gas, etc., it has become important for utilities to diversify their energy basket across cleaner sources of energy. Renewable and nuclear energy are the fastest growing cleaner sources of energy [20]. While hydroelectricity has traditionally been considered a clean source of energy, there are significant ecological and financial issues in commissioning new hydroelectric power projects. At this juncture, wind energy installations are expected to grow by 20% per annum [23].

Integration of wind power with the grid involves a producer making a bid to supply an agreed quantity of electricity at a future time. This involves placing a bid in a time-ahead market. For instance, a day-ahead bidding process would mean that the producer makes a bid for supplying wind power one day ahead of the actual generation. When a producer makes a bid for supplying energy, they do so under the following sources of uncertainty: 1) the actual amount of energy produced; 2) the price of energy some time ahead. The uncertainty in energy production can result in a producer producing more or less energy than what they bid for. In the former case, the excess energy produced may or may not be buffered. In the latter case, the producer has to pay a (Unscheduled Interchange - UI) penalty to the utility for failing to meet the contract obligation.

In developing economies (e.g., India, South Africa), where the energy market is still evolving evolving, wind integration primarily involves avoidance of penalties of under-delivery. For instance, in India, a shortage in delivered power beyond 30% of the bid placed, will incur a penalty for the shortage at the UI rate. In developed economies (e.g., UK, EU), the wind producer can seek to maximize profits by participating in spot and future energy markets. In either case, it is important to have a bidding strategy that maximizes the operating profits of the producer [9].

For a small producer of electricity, the pricing of the spot market is largely unaffected by its bid. In other words, the bidding process is the only control knob available to them in the face of the uncertainty. The bid of the producer is determined by the trade-offs between the expected price in the electricity market and the wind power yield. This requires a model for wind power prediction whose expected values can drive the bidding process. In this context, two natural questions arise:

1. How should a producer bid under uncertainty?
2. What is the economic impact of prediction errors?

Wind power prediction has been very well studied in the literature (e.g., [18]). Question 1 has also been addressed very extensively under a variety of assumptions about the price market and the wind yields (e.g., [5,17,21], detailed related work is in Section 2). Question 2, on the other hand, has received little attention and has been typically addressed for one-period problems [4]. The economic impact of the prediction accuracy depends on the how the bidding process uses information about the prediction. While several approaches have been proposed to answer question 1, the trade-off between prediction errors and profits has not been widely studied over a multi-period horizon. We fill this gap and complement existing work by focusing on the economic impact of wind power (mis)prediction.

Exact analysis under simplifying assumptions: First, we develop, *under idealized assumptions*, a closed form analytical expression for idealized bidding (IB). Specifically, we assume an 1) arbitrary price process that is independent of our bidding strategy; 2) a Markovian wind power yield process that is conditionally uniform and independent of the price process; and 3) an infinite buffer. Under these stochastic driving processes, we optimize the expected value of wind power over a finite horizon of one year as opposed to several existing works that focus on optimizing the wind power over the next bidding slot alone. We derive the optimal bidding process by solving the Bellman recursion exactly. The expression for the optimal bid in each timestep involves a) the buffer at the beginning of the timestep; b) the expected wind power yield over that timestep; and c) the expected prices over that timestep and the next (details are in Section 3).

Cross-validation with synthetic datasets: We then cross-validate our analytical derivation of the optimal bidding strategy under idealized conditions. To do this, we derive the price and yield from stochastic models that satisfy the requirements of our analysis. Under this synthetic dataset, we evaluate the profit-prediction trade-off. Specifically, we quantify the impact of prediction errors on the profits over the optimization horizon for varying caps on the buffers to store the excess energy produced. We find that a buffer can help mitigate any mispredictions to a certain extent, beyond which the expected profits almost linearly degrade with increasing prediction error. As two extreme baselines, we use a **naive** approach and an **oracle** based approach. The former makes no use of the buffer information to place a bid, while the latter has perfect knowledge of the price and the yield up to a look-ahead window into the future. With no buffering and a (relative) prediction error of 25% , we find that our bidding approach performs significantly better than a naive approach and compares favourably (86%) to an oracle with a look-ahead of *two time-slots and infinite buffer.*

Empirical evaluation with real-world datasets: Next, we relax each of our simplifying assumptions and evaluate our approach with real-world data and a finite buffer. We identify the limitations of our approach when the data (specifically, the price process) does not match our assumptions. Then we modify our approach to work around this limitation (details in Section 5). We compare our approach by comparing with the two baselines mentioned above. With buffer equivalent to 20% of the maximum yield, our approach exceeds the naive approach by 25%, while remaining within 62% of a two-step look-ahead oracle that uses infinite buffering.

Economic impact of prediction errors: We want to explore the profit vs. prediction error trade-off in a manner that is agnostic to the predictor. Specifically, we do this by constructing sample paths of the predicted yields with a parametrized prediction error from the actual wind yield obtained from real-world data. Then we evaluate the expected profits under our idealized bidding algorithm, and the naive and oracle-based baselines. This quantifies the trade-off between the profits from wind power and the errors in misprediction. For a given tolerance of the profit from the maximum possible, we also quantify the trade-off between the buffer sizing required as a function of the prediction error.

Contributions: Our specific findings and contributions include the following:

- We analyze a one-period ahead bidding strategy to obtain a closed-form expression for the optimal bid under idealized conditions. The closed-form expression provides intuition behind the choice of the bid as a trade-off point between the higher and lower limits of the yield distribution and the expected prices.
- We cross-validate the analytical expression on synthetic datasets for the price and yield processes. The price and yield processes match real-world data in mean and deviation. We quantify the profit vs. prediction error trade-off with varying absolute and relative prediction errors. We find that prediction errors do not matter significantly beyond a buffer capacity of 30% of the maximum generation. For a buffer capacity of 20%, we find that the impact of prediction errors on profits to be less than 1.2%.
- We evaluate our approach on real-world datasets by relaxing the assumptions on the price and the yield processes. Specifically, we modify our approach to account for the fact that price and yield processes in reality violate our idealized assumptions. Under our modified bidding strategy, we quantify the profits vs. prediction error trade-off. We find that modified IB remains stable on the real-world data.
- We present two baselines as comparison with IB for both synthetic and real datasets. At the lower extreme, we consider a naive bid strategy; and at the higher extreme, we consider an oracle that gives perfect prediction of yield and price processes over the next few time-slots.
- We quantify the impact of the increasing penalties on the expected net profits. We find that modified idealized bidding is almost agnostic to variations in the penalty.

The rest of the paper is organized as follows. Related work is surveyed in Section 2. Section 3 presents our analysis of optimal bidding using a Bellman recursion and identifies the formula for the optimal bid. Section 4 evaluates our algorithm on synthetic datasets that satisfy all the assumptions of our analysis. Section 5 evaluates our algorithm on real-world datasets for price and yield and explains how to improve our algorithm when the conditions assumed in the analysis are not met. Section 6 studies the impact of penalty on net profits. Section 7 concludes.

2. RELATED WORK

A number of wind yield prediction techniques have been tried. These techniques are based on statistical methods, numeric weather prediction methods, machine learning models and hybrid forecasting approaches [11]. Reference [18] uses alternative models based on the variables involved and combine these models to obtain the final prediction. Numerical Weather Prediction (NWP) models outperform statistical and machine learning techniques over long term prediction horizons. However, the statistical techniques perform better over shorter horizons [19], [2], [12].

A large section of the related literature focuses on optimal bidding as a one period problem, optimizing for just the next period, under various settings. Usaola et al [22] look at the problem of bidding one period ahead in an intraday market, given a position in a forward market. Botterud et al [5] consider the question of how to make the next bid in a day ahead market for different types of objective functions (Conditional Value at Risk (CVAR), expected profit, etc.). Puglia et al [17] also look at minimization of a risk measure, under the presence of both a day ahead and an ancillary market. Liang et al [13] address the problem of bidding optimally for the next period, provided the wind producer has access to both an energy market and a reserve market. Bitar et al [3] solve the problem of optimally bidding for the next time slot, allowing wind output to vary continuously within the slot. They then generalize the problem somewhat by allowing storage to be used within the slot, but it is still with a constant bid over the entire period. Other papers that look at one-period bidding problems include [16], [24], [6], [7], [8], [25] and [21]

There have been a few papers that have looked at N-period horizon bidding problems. Lohndorf et al [14] consider such an N-period problem, with finite buffers. They model the problem as a Markov decision process, and use an approximate dynamic programming approach to provide approximate solutions to it. Apart from the fact that our paper provides detailed empirical analysis of the value of prediction accuracy to bidding (which is not considered in [14]), a few other differences can also be noted between this work and ours. We allow for an arbitrary price process, while this work models price by an AR(1) process and the solution critically depends on this assumption. Further, our work provides an analytical solution that is optimal. Our optimal solution also turns out to be simple and can be implemented at each state of the system in constant time; on the other hand, the ADP methods used in [14] take significant computational resources. Morales et al [15] solve an N-period problem by reducing continuous variables into discrete scenarios and using linear programming. However, they do not allow buffers (any excess is sold off immediately), and therefore it is fundamentally not different from just a one-period problem. Further, the running time can quickly go up with increasing number of scenarios chosen. Giannitrapani et al [10] employ a similar LP technique as [15]; [10] differs from [15] by allowing buffers, but assuming a constant price and penalty throughout.

One of the main objectives of this paper is to analyze how better prediction accuracy can translate to better bidding, and therefore better profits. There has been some recent initial effort in this area, but only for a one-period case. Bitar et al, in [4], consider a problem of choosing an optimal constant bid for the next time slot, allowing wind output to vary continuously within the slot. They also analyze value of information about next period's wind power to the quality of bidding. Specifically, they quantify information using the conditional value-at-risk (CVaR) measure, and show how better information can increase profits. However, in this paper, we seek to evaluate (empirically) the value of prediction accuracy in a more general situation with multiple periods and with storage.

3. IDEALIZED BIDDING (IB)

Consider a wind power producer who needs to bid at time-slot $t-1$ for supplying energy in the slot t. The bid placed k_t needs to be decided by the producer in response to the predictions of the yield Y_t and the price P_t processes. Suppose buffering is allowed and the buffer at the end of a slot t is b_t, then the energy delivered d_t in slot t is given by $\min(Y_t + b_{t-1}, k_t)$.

If buffering is allowed, the buffer b_t stores any energy that is produced in excess of the bid, and can be used to meet any shortfall in the yield below the bid. We ignore buffer related energy losses in this work. Any unit of energy that is actually delivered from the yield or the buffer earns a revenue of P_t and any energy delivered by procuring from the spot market incurs a penalty of L_t. The net profit is difference between the revenue and the penalty paid. In sum, our problem statement is as follows:

Problem statement: Over an optimization horizon of T (say, a year) divided into equal number of (say, six hour) slots, given (probabilistic) knowledge about the yield and the price, we want to identify a bid strategy that maximizes the expected value of the net profit over the entire duration T at time $t = 0$ with future profits discounted by a factor δ.

Symbol	Meaning
T	Optimization horizon
t	Timeslot id; ranges in $[1..T]$
CF_t	Net profit in slot t
Y_t	Yield in t
y_l	Lower bound on $Y_t \vert Y_{t-1}$
y_h	Upper bound on $Y_t \vert Y_{t-1}$
P_t	Unit price in t
b_{t-1}	Buffer at end of $t-1$
k_t	Bid placed at $t-1$ for delivery in t
d_t	Amount actually delivered in t given by $d_t = \min(Y_t + b_{t-1}, k_t)$
L_t	Penalty for each unit of under-delivery
δ	Discount factor to translate cash-flows between slots
p_0	Terminal sale price at last slot T
$\mathbb{E}_t[]$	Expectation given all information up to and inclusive of t
$\mathbf{P}()$	Probability distribution

Table 1: Notation

Notation used is summarized in Table 1. Our goal is to find the optimal strategy $\{k_1^*, ..., k_T^*\}$ to maximize

$$\sum_{t=1}^{T} \delta^t \mathbb{E}_0(CF_t)$$

In other words, we want to identify a strategy that maximises the value function where the expected cash flows of future timeslots are discounted by the factor δ. The cash flow at the terminal time slot T is:

$$CF_T = d_T P_T - [k_T - d_T]L_T + [b_{T-1} + Y_T - d_T]p_0$$

where $d_T = \min(Y_T + b_{T-1}, k_T)$ Specifically, at T, there is no further buffering. Anything in excess of the delivered amount is dumped at the dumping price p_D. The cash flow at slots $t < T$ is:

$$CF_t = d_t P_t - [k_t - d_t]L_t$$

We obtain the optimal strategy using Bellman dynamic programming. The value function is:

$$V_t = CF_t + \delta \cdot \mathbb{E}_t(V_{t+1}^*), t < T$$

with $V_T = CF_T$. We update the buffer at every t as follows:

$$b_t = b_{t-1} + Y_t - d_t$$

The problem at each $t \in \{0, 1, ..., T-1\}$ is:

$$V_t^* = \max_{k_{t+1}} V_t$$

The value under the optimal strategy $\{k_1^*, ..., k_T^*\}$ is V_0^*.

3.1 Assumptions

For the sake of analytical tractability, we make the following simplifying assumptions (that will be relaxed in the experimental evaluation):

- At $t-1$, we bid k_t for delivery during t.
- Y_t is conditionally uniform: $Y_t|Y_{t-1}$ is $\sim U(y_l, y_h)$, where the range is determined by Y_{t-1}.
- Y_t and P_t are independent.
- Buffers are infinite and costless.
- The per-unit price of dumping is low:

$$p_0 < \mathbb{E}_t(P_{t+1}), \forall t \tag{1}$$

- Penalty is sufficiently large to avoid simple hoarding:

$$\mathbb{E}_{t-1}[L_t] > \mathbb{E}_{t-1}[-P_t + \delta P_{t+1}], \forall t \tag{2}$$

  Essentially, this precludes someone from buying in the spot market at time t and hoarding it to sell it at time $t+1$.

3.2 Preliminaries

Our strategy is to calculate the expected delivery $\mathbb{E}_{t-1}[d_t]$ and see how it changes with the bid value k_t. Since the delivered amount is the minimum of the bid placed and the available energy (the sum of the buffer and the current yield), we proceed as follows:

$$\mathbb{E}_{t-1}[d_t] \quad = \quad \mathbb{E}_{t-1}[b_{t-1} + \min\{Y_t, k_t - b_{t-1}\}] \tag{3}$$

If Y is a uniform random variable in $[y_l, y_h]$, then for $a \in [y_l, y_h]$

$$\begin{aligned}\mathbb{E}[\min(Y, a)] &= \int_{-\infty}^{a} x f(x)dx + \int_{a}^{\infty} a f(x)dx \\ &= a\alpha_t - \frac{a^2}{2(y_h - y_l)} + g(y_h, y_l)\end{aligned}$$

where $\alpha_t = \frac{y_h}{y_h - y_l}$ and $g(.)$ is a function of only y_h and y_l and independent of a. Using this result in Equation 3, we have

$$E_{t-1}[d_t] = b_{t-1}(1-\alpha_t) + k_t\alpha_t - \frac{(k_t - b_{t-1})^2}{2(y_h - y_l)} + g(.) \tag{4}$$

Differentiating $E_{t-1}[d_t]$ with respect to k_t yields,

$$\frac{\partial \mathbb{E}_{t-1}(d_t)}{\partial k_t} = \alpha_t - \frac{(k_t - b_{t-1})}{(y_h - y_l)} \tag{5}$$

3.3 Solving the Bellman recursion for $t = T$

We solve for the optimal bid in every timestep using Bellman's recursion working from backwards. Recall that the value function is recursively defined with

$$V_t = d_t P_t - [k_t - d_t]L_t + \delta \mathbb{E}_t(V_{t+1}^*)$$

with a terminal condition at:

$$V_T = d_T P_T - [k_T - d_T]L_T + (b_{T-1} + Y_T - d_T)p_0$$

Now $\dfrac{\partial V_{T-1}}{\partial k_T} = \delta \dfrac{\partial}{\partial k_T}\mathbb{E}_{T-1}[V_T]$

$$\begin{aligned}\frac{\partial \mathbb{E}_{T-1}[V_T]}{\partial k_T} &= \left(\frac{\partial E_{T-1}[d_T]}{\partial k_T}\right)\left(\mathbb{E}_{T-1}[P_T + L_T] - p_0\right) \\ &- \left(\mathbb{E}_{T-1}[L_T]\right)\end{aligned} \tag{6}$$

We use Equation 5 to evaluate $\frac{\partial E_{T-1}[d_T]}{\partial k_T}$. To solve for the optimal k_T we equate Equation 6 to 0. Under the assumption given in Equation 1, the second derivative w.r.t. k_T (that we omit for the sake of brevity) is negative and hence the local optimum is a maximum. Thus, the optimal bid for T works out to:

$$\boxed{k_T^* = y_h + b_{T-1} - \frac{\mathbb{E}_{T-1}[L_T](y_h - y_l)}{\mathbb{E}_{T-1}[P_T + L_T] - p_0}} \tag{7}$$

3.4 Solving the Bellman recursion for $t < T$

Solving the recursion for $t < T$ is slightly more involved, and we give only the highlights of the analysis. Recall that

$$V_{t-1} = d_{t-1}P_{t-1} - [k_{t-1} - d_{t-1}]L_{t-1} + \delta \mathbb{E}_{t-1}(V_t^*)$$

$$\begin{aligned}\frac{\partial V_{t-1}}{\partial k_t} &= \delta \frac{\partial}{\partial k_t}\mathbb{E}_{t-1}(V_t^*) \\ &= \delta\left(\left(\frac{\partial \mathbb{E}_{t-1}[d_t]}{\partial k_t} \cdot \mathbb{E}_{t-1}[P_t + L_t]\right) - \mathbb{E}_{t-1}[L_t]\right) \\ &+ \delta^2 \frac{\partial}{\partial k_t}\mathbb{E}_{t-1}[V_{t+1}^*]\end{aligned} \tag{8}$$

The first term of Equation 8 is similar to the expression for the case of $k = T$ that we saw earlier in Equation 6 (only with p_0 set to zero). Because $t < T$ does not have a terminal condition, we need to evaluate the second term of Equation 8 separately. Using the law of iterated expecations and some algebraic transformations, we have

$$\delta^2 \frac{\partial}{\partial k_t}\mathbb{E}_{t-1}[V_{t+1}^*] = -\delta^2 \frac{\partial}{\partial k_t}\left(\mathbb{E}_{t-1}[d_t] \cdot \mathbb{E}_{t-1}[P_{t+1}]\right) \tag{9}$$

Now, substituting Equation 9 back into Equation 8 and equating it to zero, we have

$$k_t^* = \alpha_t(y_h - y_l) + b_{t-1} - \frac{\mathbb{E}_{t-1}[L_t](y_h - y_l)}{\mathbb{E}_{t-1}[P_t - \delta P_{t+1} + L_t]}$$

which is

$$\boxed{k_t^* = y_h + b_{t-1} - \frac{\mathbb{E}_{t-1}[L_t](y_h - y_l)}{\mathbb{E}_{t-1}[P_t - \delta P_{t+1} + L_t]}} \tag{10}$$

This choice of k_t^* is a maxima because the assumption in Equation 2 ensures that the second derivative is negative.

3.5 Discussion

We make the following observations about the expression for the optimal bid. As $\mathbb{E}_{t-1}[L_t]$ increases, k_t^* decreases, indicating that as the penalty increases, the bidding strategy should be more conservative. Similarly, if $\mathbb{E}_{t-1}[P_t]$ increases, k_t^* increases, but if $\mathbb{E}_{t-1}[P_{t+1}]$ increases, k_t^* decreases. If the expected price in the next slot t is higher, we bid more, if the expected price in the slot $t+1$ is higher, we defer. As y_h increases, k_t^* increases if next period's expected price is relatively more attractive (i.e. if $\mathbb{E}_{t-1}[P_t] > \delta \cdot \mathbb{E}_{t-1}[P_{t+1}]$)

Data	Value	Units
Min Wind Yield	0	KwH
Max Wind Yield	21000	KwH
Average Wind Yield	5358	KwH
Min Unit Price	8.48	£/KwH
Max Unit Price	257.21	£/KwH
Average Unit Price	25.41	£/KwH

Table 2: Digest of real-world data for slots of 6 hours. All yields are normalized to a single turbine. The data is from the UK geography and the prices are in GBP (£)

Parameter	Symbol	Value	Units
Price	μ_P	16.92	£/KwH
	α_P	0.47	
	σ_P	9.41	£/KwH
Yield	μ_Y	1533	KwH
	α_Y	0.73	
	σ_Y	4061	KwH
Risk-free interest rate	r_f	3.5% p.a.	
Discount factor	δ	0.99	
Absolute prediction error σ	Varied		KwH
Relative prediction error ρ	Varied		% of prediction

Table 3: Values of parameters used to generate price and yield processes. The data is from the UK geography and the prices are in GBP (£)

4. EVALUATION ON SYNTHETIC DATASETS

We now evaluate our IB strategy with data that is generated synthetically but matching real-world data in mean and deviation. We use a synthetic dataset to ensure that all the assumptions made for the analysis are satisfied by the dataset. On this, we explore the profit vs. prediction tradeoff for IB.

Time model: We consider a bidding horizon of 1 year. Each bidding slot is 6 hrs, i.e., we consider wind prediction and bid over (non-overlapping) 6 hour blocks. Thus we have $T = \frac{1yr}{6hrs} = 1460$ slots. Under this horizon, we want to maximize the expected profits over all timeslots with the profit at each slot t discounted by δ^t. To get the discount factor δ, we use the annual risk-free rate of interest $r_f = 3.5\%$ and divide it across 6 hour blocks. That is, $\delta = \frac{1}{(1+r_f/T)} = \frac{1}{(1+0.025/1460)} = 0.99$

Price model: For the price process, following [14], we use an AR(1) process. We use real-world datasets obtained from [1] to obtain the mean, variance, and the first-order correlation of the price process. Specifically, the price process is modeled as $P_{t+1} = \alpha_P P_t + \mu_P + \xi_t$, where P_{real} denotes the price from a real-world dataset, $\mu_P = \mathbb{E}[P_{real}](1 - \alpha_P)$ and ξ_t is $\mathcal{N}(0, \sigma)$ with $\sigma^2 = Var[P_{real}](1 - \alpha^2)$.

Under this model the assumption that the bidding knows $\mathbf{P}(P_{t+1}|P_t)$ at time t and hence $\mathbb{E}_t[P_{t+1}]$ is satisfied.

Yield and prediction errors: For the yield process, we similarly use real-world data [1] to identify parameters for an AR(1) process. We first generate a sample path $\hat{Y}(t)$ from an AR(1) process. We introduce a **prediction error** to the actual yield $\hat{Y}_t$ to get a predicted yield Y_t as follows. Specifically, once a sample path $\hat{Y}(t)$ is generated, we assume that the predicted yield is given by

$$\mathbf{P}(Y_{t+1}|\hat{Y}_t) = \hat{Y}_{t+1} + Uniform[-\sigma, \sigma] \quad (11)$$

to model absolute errors in prediction that are same over the entire region of wind values or

$$\mathbf{P}(Y_{t+1}|\hat{Y}_t) = \hat{Y}_{t+1} + Uniform[-\rho, +\rho\,] \times \hat{Y}_{t+1} \quad (12)$$

to model relative errors in prediction whose ranges depend on the actual value being predicted. This satisfies the assumption in the analysis that the conditional yield is uniformly distributed and that the lower and upper limits of the uniform distribution are determined by $\hat{Y}_t$. The parameters used for generating the yield and the price processes are shown in Table 3.

Bidding strategies evaluated: We evaluate the following bid strategies:

1. Our approach (IB) described in Section 3. Here y_h would be $\hat{Y}_t + \sigma$ and y_l would be $\hat{Y}_t - \sigma$ in the case of absolute prediction errors and the corresponding ρ term for relative prediction errors.
2. The first baseline strategy (termed **Naive**) bids exactly the expected value of the yield according to the wind prediction model, i.e., $k_t = \hat{Y}_t$ + prediction noise, while the actual yield is $\hat{Y}_t$ and the noise is given in Equations 11 and 12 respectively.
3. The second baseline strategy (termed **Oracle(l)**) is as follows:
   - The bidder knows perfect information about the price for the next l steps.
   - With an infinite buffer, the yield Y_t is always sold at the maximum price in the slots $[t, t+l-1]$
   - With $l = 1$, slot t's yield is sold off at that slot's price P_t with no penalty at all.

4.1 The profit-prediction tradeoff

Under these generating processes for the price and the yield, we generate an ensemble of 1000 sample paths for each set of parameters. We compute (ensemble average) expected net profit over the entire bidding horizon. Then we plot the prediction error on the X-axis and the expected net profit on the Y-axis for the various bidding strategies. This quantifies the tradeoff between prediction error and the expected net profit.

The results are shown in Figures 1a, 1b, and 1c for IB, naive, and oracle bid strategies when the prediction errors are absolute (i.e., the error magnitude is the same irrespective of the predicted value). The X-axis shows the absolute prediction error, and the Y-axis shows the expected net profits for each of the bidding strategies. Each plot shows a family of curves parametrized by various buffer values (B_{cap}). We make the following high-level observations:

- The expected profits fall with increasing errors in both IB and naive strategies. In the Oracle strategy, with increasing look-ahead knowledge, we have improved profits.
- The rate of fall is higher for the naive strategy than IB despite both using buffers of the same capacities.
- In general, for increasing buffer capacities, the tolerance of the net profit to error increases for both IB and Naive. For IB, beyond a buffer cap of about 6000 KwH (which is 30% of

[1]We do not mention the source of real-world data for the wind power yields due confidentiality requirements.

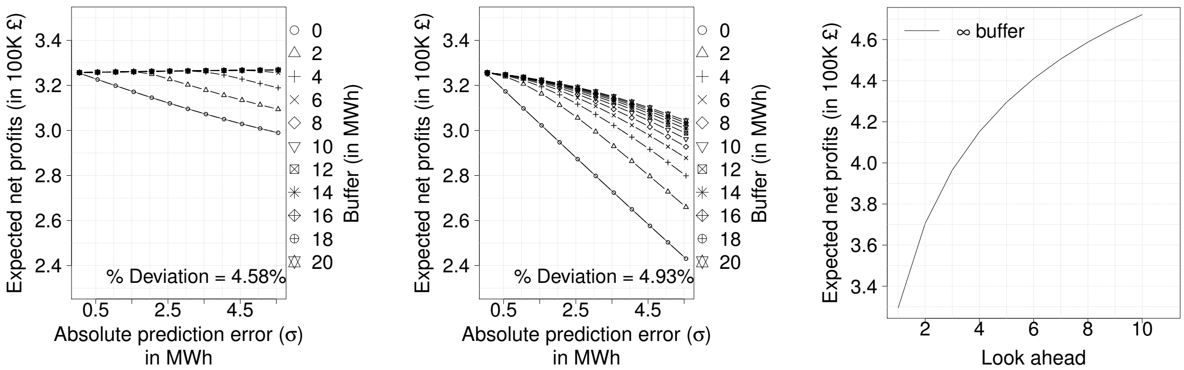

(a) IB (b) Naive (c) Oracle

Figure 1: Profit-vs-prediction tradeoff with varying buffer capacities and absolute prediction errors

the peak yield), prediction errors do not matter significantly. With a buffer capacity of 20% of peak yield, the impact of prediction errors on profits is less than 1.2%.

- The Oracle strategy with a perfect look-ahead of length $l = 1$ matches closely with IB, indicating that IB is nearly optimal when the look-ahead knowledge is perfect only over short-intervals.

Figures 2a and 2b show the profit-vs-prediction tradeoff for relative errors in prediction. We note that the trends are similar to the absolute prediction errors. We now analyze each of these bidding strategies in detail.

4.2 Analysis of IB

Figure 1a shows IB's performance for varying values of prediction noise. Specifically, the X-axis shows the absolute error in prediction of the yield. The Y-axis shows the expected value of the cashflow over the entire bidding period. The curves are parametrized by increasing values of the buffer cap (B_{cap}). We see that as the prediction error increases, the expected cash flow falls as one would expect. For very high values of B_{cap}, there is no significant fall as the buffer can handle the excess and deficit over multiple timeslots. For lower values of B_{cap} the expected net cash flow stays flat for lower values of σ, but show a sharp dip when σ exceeds some threshold. All deviations are within 5% of the ensemble means. Figure 2a shows similar results when the errors are no longer absolute but are a constant fraction (ρ) of the predicted value. As before, the X-axis shows increasing values of ρ and the Y-axis shows the expected cash flow. For the sake of brevity, we restrict our analysis of the IB to the case where the prediction errors are absolute (i.e., for the errors parametrized by σ). Specifically, we explain the behaviour of the expected cash flows with increasing σ. We have assumed that the prediction error is uniform in $[-\sigma, +\sigma]$. For sake of brevity, we omit $(t-1)$ in $E_{t-1}[.]$ and use $\mathbb{E}[.]$. Because $P_{t+1} = \alpha_P P_t + \mu + \xi_t$ is a weak-sense stationary process, we have as $\alpha_P < 1$

$$\lim_{t\to\infty} \mathbb{E}[P_t] = \lim_{t\to\infty} \mathbb{E}[P_{t+1}] = \frac{\mu}{(1-\alpha_P)} \tag{13}$$

Therefore, under the assumption that $\mathbb{E}[L_t] = \lambda \mathbb{E}[P_t]$, the optimal bid k_t^* simplifies to

$$k_t^* = y_h + b_{t-1} - \frac{\lambda}{(1-\delta)+\lambda} \times (y_h - y_l) \tag{14}$$

Further, because $\delta \approx 1$ (explained earlier in Section 4) we have $k_t^* = b_{t-1} + y_l$. Recall that y_l is the lower limit of the uniform

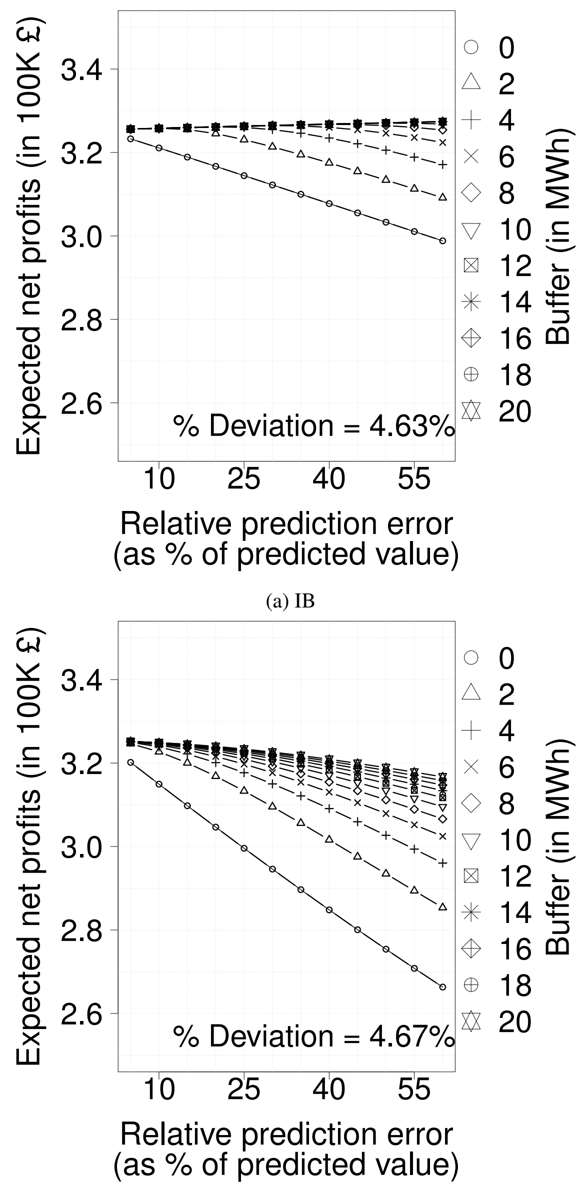

(a) IB

(b) Naive

Figure 2: Profit-vs-prediction tradeoff for varying buffer capacities and relative prediction errors

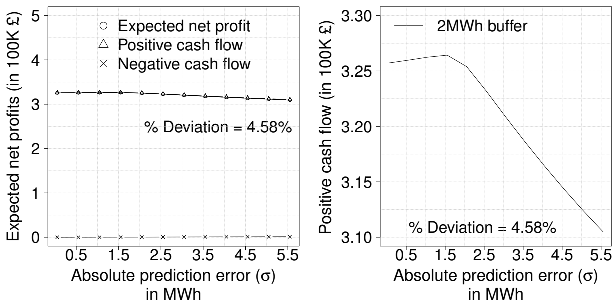

Figure 3: Break-up of the expected net cash flow into the positive and negative cash flows confirming the analytical predictions. The buffer cap used is 2000 KwH.

marginal $P(Y_{t+1}|Y_t)$. On an average, the expected bid at t is given by $\mathbb{E}[b_{t-1}] + \mathbb{E}[y_l]$ Because we bid at the lower limit, the penalty is never paid. And we store $\mathbb{E}[Y_t - y_l]$ in the buffer. For a constant error of prediction parametrized by σ, we have $\mathbb{E}[Y_t - y_l] = \sigma$. Therefore, we expect to discharge σ from the buffer and store σ on an average over every t.

If the buffer cap $\mathbf{B_{cap}} > \sigma$, all of the yield excess over the bid value can be stored over all timesteps. The expected per-step cash flow over all timesteps would be $\lim_{t\to\infty} \mathbb{E}[Y_t]\mathbb{E}[P_t]$. By our assumptions of independence, this would simplify to $\frac{\mu_Y \mu_P}{(1-\alpha_P)(1-\alpha_Y)}$ which is independent of B_{cap} and σ. Therefore, we get a flat line behaviour. The expected steady-state buffer size is given by σ.

If the buffer cap $B_{cap} < \sigma$, we will expect to lose some energy in every timestep when the yield is in excess of the bid value because the buffer cannot store it. Specifically, the expected per-step cash flow is given by $\mathbb{E}[P_t(b_{t-1} + Y_L) + P_D(Y_t - Y_L)]$. The expected buffer size is given by B_{cap} and $\mathbb{E}[Y_L] = \mathbb{E}[Y_t] - \sigma$. Thus the per-step cash flow becomes

$$\frac{\mu_P}{1-\alpha_P} \times \left(B_{cap} + \frac{\mu_Y}{1-\alpha_Y} - \sigma \right) + (\sigma - B_{cap})\, P_{dump}$$

Because $P_{dump} < \dfrac{\mu_P}{1-\alpha_P}$, the expected cash flow is a linearly decreasing function of σ.

Figure 3 shows the expected net cash flow and the expected negative and positive cash flows of the IB algorithm for varying values of σ for $B_{cap} = 2000$. As we can see, the expected negative cash flow stays almost constant at zero, while the expected positive cash flow decreases after $\sigma > B_{cap}$ as predicted by the analysis thus confirming our analysis of the IB algorithm.

Figure 4 shows the time evolution of the expected buffer size for the IB algorithm. As expected, for $\sigma < B_{cap}$, the steady-state buffer value is approximately σ, while for $\sigma > B_{cap}$, the buffer saturates close to the maximum cap available B_{cap}. Thus our experiments confirm our analytical predictions.

4.3 Analysis of naïve bidding

Figure 1b shows the performance of naïve bidding for various values of the prediction noise.

For the naïve bidding strategy, where $k_t = \mathbb{E}[Y_t]$, we see a smooth behaviour in the curve (as opposed to a threshold-based knee like behaviour seen for IB.) When we bid $E[Y_t]$, any excess of the yield over the bid is buffered and any deficit in the yield can be met from the buffer. However, the buffer value itself is not used to make a bidding decision. Therefore, if $\epsilon_t = E[Y_t] - Y_t$ the

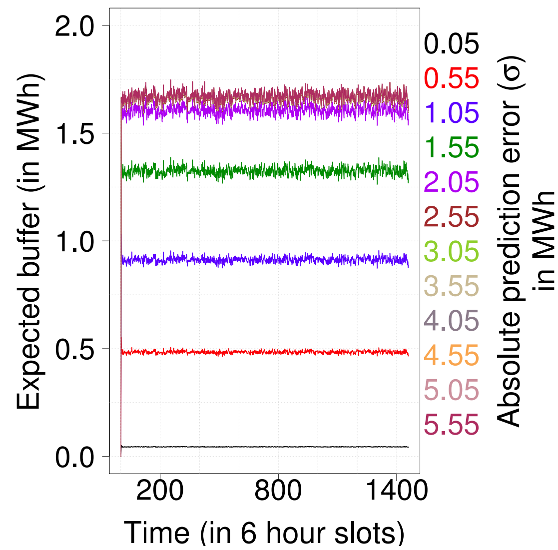

Figure 4: Time evolution of the ensemble average buffer $\mathbb{E}[B_t]$ for $t \in [0, 1460]$ for the bid strategies with prediction errors. IB

buffer evolves according to the following equations:

$$B_{t+1} = \begin{cases} \min(B_t - \epsilon_t, B_{cap}) & \text{if } \epsilon_t < 0 \\ \max(B_t - \epsilon_t, 0) & \text{if } \epsilon_t > 0 \\ B_t & \text{if } \epsilon_t = 0 \end{cases}$$

Because ϵ_t is uniform in range $[-\sigma, +\sigma]$, we can solve the stochastic difference equation for the steady state behaviour of B_t as $t \to \infty$ similar to diffusion processes with barriers. We omit the analysis for sake of brevity and state the following results:

$$\begin{aligned} \lim_{t\to\infty} \mathbb{E}[B_t] &= B_{cap}/2 \\ \lim_{t\to\infty} \mathbf{P}(B_t) &\xrightarrow{D} Unif[0, B_{cap}] \\ &\quad \text{with point masses } p_m \text{ at } \{0, B_{cap}\} \end{aligned}$$

An illustration of the buffer distribution is shown in Figure 5.

At timeslot t, if there is no penalty, then the expected profit is precisely $E[P_t]E[Y_t]$. However, if the yield happened to be less than the bid **and** the buffer cannot meet the deficit, we end up paying a penalty. To estimate the probability of paying a penalty in the steady-state we proceed as follows. Because Y_t is exogenous, it and, therefore, ϵ_t are independent of B_{t-1}. Recall that ϵ_t and B_{t-1} are continuous and mixed (i.e., continuous with point masses) random variables respectively. So we have

$$\mathbf{P}(\text{penalty}) = \int_x \mathbf{P}(\epsilon_t \in [x, x+dx]) \times \mathbf{P}(B_{t-1} < x)$$

For $B_{cap} > \sigma$, we have

$$\begin{aligned} \mathbf{P}(\text{penalty}) &= \int_0^\sigma \frac{dx}{2\sigma} \times \int_0^x \mathbf{P}(B)dy \\ &= \int_0^\sigma \frac{dx}{2\sigma} \times \left(p_m + \frac{x(1-2p_m)}{B_{cap}} \right) \\ &= \frac{p_m}{2} + \frac{\sigma(1-2p_m)}{4B_{cap}} \end{aligned}$$

Because $\sigma < B_{cap}$ this is an increasing function of p_m, and consequently the penalty would increase with increasing p_m. For

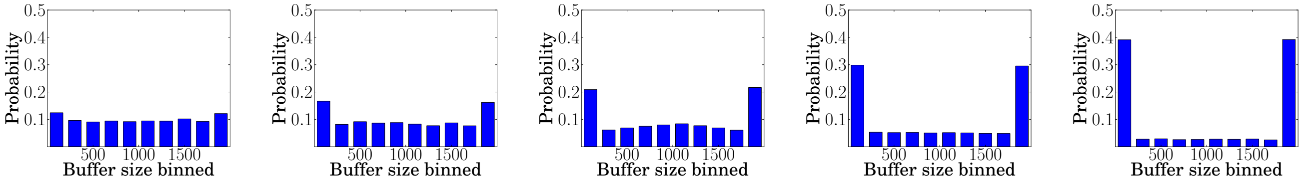

Figure 5: PDF of the steady-state buffer distribution for $B_{cap} = 2000$ units for σ in the list [200, 600, 1000, 2000, 3800] units.

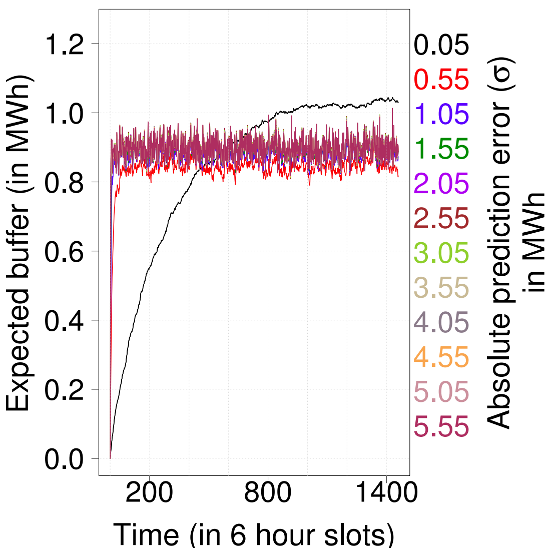

Figure 6: Time evolution of the ensemble average buffer $\mathbb{E}[B_t]$ for $t \in [0, 1460]$ for the bid strategies with prediction errors. Naive approach

$\sigma > B_{cap}$, we have

$$\begin{aligned}\mathbf{P}(\text{penalty}) &= \frac{(\sigma - B_{cap})}{\sigma} + \int_0^{B_{cap}} \frac{dx}{2\sigma} \times \int_0^x \mathbf{P}(B)dy \\ &= \frac{(\sigma - B_{cap})}{\sigma} + \frac{B_{cap}}{4\sigma}\end{aligned}$$

independent of p_m and increases smoothly with σ. For $\sigma = B_{cap}$, the expressions for the latter two cases both simplify to $\frac{1}{4}$ as can be expected.

Summary: Therefore, the expected penalty also increases super-linearly and smoothly with σ for $\sigma < B_{cap}$ as well as $\sigma > B_{cap}$. So the net cash flow falls smoothly. The inflection point occurs at $\sigma = B_{cap}$. Figure 6 shows the time evolution of the expected buffer size averaged across 1000 runs for the range of timeslots for $B_{cap} = 2000$ and various values of σ. As the analysis predicts, it converges to approximately $B_{cap}/2$. The slight deviation from $B_{cap}/2$ is because the empirical evaluation does not allow the yield values to become negative, while the analysis does not explicitly account for the case (instead focusing on the buffer values alone being non-negative).

5. EVALUATION ON REAL-WORLD DATA SETS

As mentioned before, we had access to historical data for wind yield over 2003 and 2005. The price data for the same period was obtained from [1]. Figure 7 shows the price and the yield data for 2003; Figure 8 shows the same for 2005. We evaluated the performance of the IB algorithm and the two benchmarks as follows. The expected value of the wind was uniformly chosen within an interval of the actual wind speed. This makes the prediction error to be uniformly distributed in a pre-determined interval.

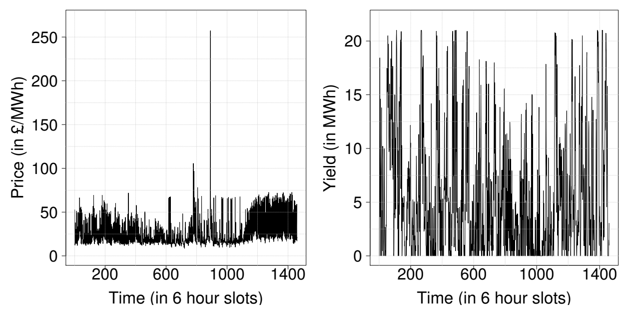

(a) Price (b) Yield

Figure 7: Price and yield for wind in 2003

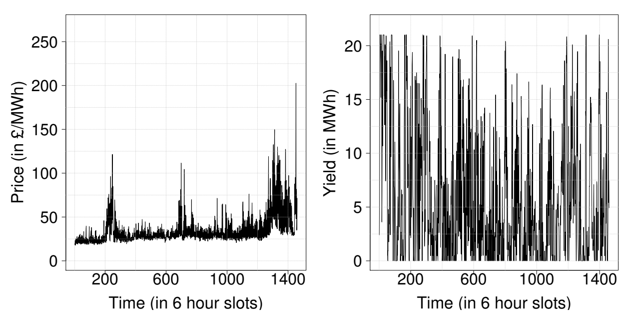

(a) Price (b) Yield

Figure 8: Price and yield for wind in 2005

Under this assumption on the prediction error, we evaluated the IB algorithm with real-world price and yield data (2003). Figures 9a and 9b show the results for absolute and relative errors respectively. IB performs significantly worse than the other naive bidding algorithms that do not even use the buffer or the buffer state.

Explanation of poor performance of IB: To understand this behaviour, we evaluated if the conditions (specified in Equations 1 and 2) for the IB algorithm to be optimal are valid. The assumption that the price process satisfies the condition that $E[L_t] > E[P_t] - \delta E[P_{t+1}]$ is violated in the real-world price data that we used. Intuitively, the IB algorithm assumes that the price process varies smoothly. Indeed, if $L_t = \lambda P_t$, then $\mathbb{E}[P_{t+1}] < \frac{(1+\lambda)\mathbb{E}[P_t]}{\delta}$. The violation of this assumption has two consequences:

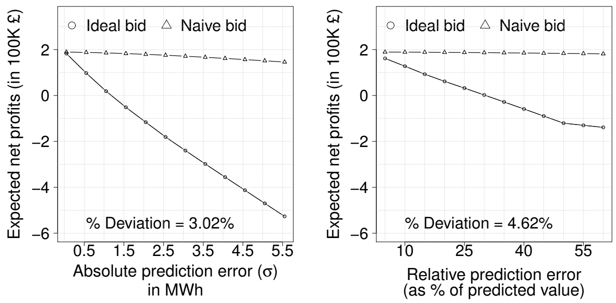

(a) Absolute errors. (b) Relative errors.

Figure 9: Poor performance of the IB algorithm on the real-world data for 2003.

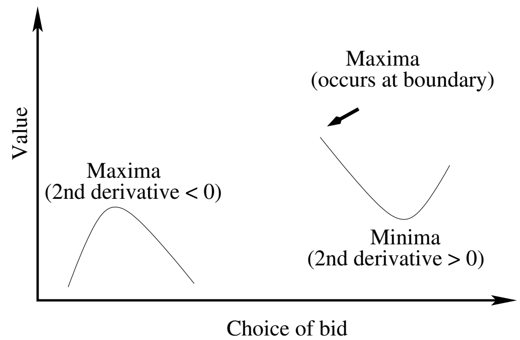

Figure 11: Intuition behind choice of optimal bid when the condition for local maxima is flipped.

- The second derivative used for checking the optimality in the expected per-slot payoff flips signs to become positive and therefore indicative of a minima.
- In the optimal bid formula

$$k_t^* = y_h + b_{t-1} - \frac{\mathbb{E}_{t-1}[L_t](y_h - y_l)}{\mathbb{E}_{t-1}[P_t - \delta P_{t+1} + L_t]}$$

  the denominator $\mathbb{E}_{t-1}[P_t - \delta P_{t+1} + L_t]$ becomes negative, and so the bid placed becomes

$$k_t^* = y_h + b_{t-1} + \text{a positive quantity}$$

  In other words, IB systematically overbids above the range of possible yields in every timestep leading to a very high negative cashflow and, consequently, a poor performance.

Modifying IB for real-world data: We observe that even under the violation of the price assumption, the expression for the optimal k_t^* still identifies a local extrema. Without going into the details of the mathematics, we give the intuition in Figure 11.

When the second derivative changes sign, we expect the bid to be a local minima rather than a maximum. Because there are no other local extrema, the boundary value of the function at the lowest choice of the bid is the local maxima in that interval. This corresponds to a bid value of precisely 0. Therefore, we modify the **IB algorithm to bid zero** when the condition on the expected price is violated.

Performance of modified IB: We now evaluate the modified IB algorithm with the price and yield data (2005) obtained from real-world sources. We also restrict the buffer capacity to 20% of the maximum yield (i.e., a value of 4 MwH) to avoid unrealistic buffer sizes. Figure 10a shows the results of IB when the error involved in the prediction is absolute. In both cases, the IB algorithm is relatively stable in the face of increasing errors, while the naive strategy quickly degrades in performance.

6. IMPACT OF PENALTY

So far, we studied the impact of prediction inaccuracy assuming that the penalty is a constant multiple of the bid price. In other words, $L_t = \lambda P_t$ with λ being held constant. We now study the impact of prediction errors on the profits when the L_t varies significantly. Specifically, for a given constant prediction error σ, we vary L_t as follows:

- Relative to P_t, i.e., increasing the value of λ
- Independent of P_t, with L varying from $\frac{P_{max}}{10}$ to $2P_{max}$ in steps of $\frac{P_{max}}{10}$.

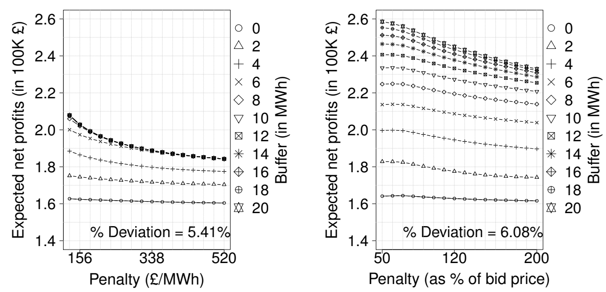

(a) Absolute penalty (b) Relative penalty

Figure 12: Effect of penalty on the expected profit for IB. Somewhat constant, but falls with increasing values of penalty.

Figures 12a and 12b show the effect of increasing penalty values on the expected profits on the real-world data for IB. These can be understood as follows. Because IB in the real-world data incurs penalties, we expect the expected profit to decline with increasing penalties. This trend is seen across the same value of B_{cap}. With increasing B_{cap}, we expect the effect of penalty to be less as the buffer helps smoothen out the bid-yield mismatch. This trend is also observed in both figures. A similar behaviour is seen in Figures 13a and 13b for the naive bidding strategy, except that the penalty increases rapidly with increasing penalty.

For the synthetically generated datasets, as indicated in the analysis in Section 4, the bid formula does not depend on the penalty

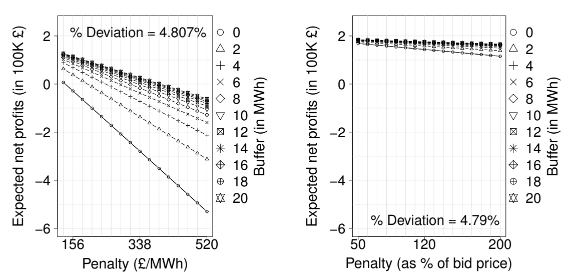

(a) Absolute penalty (b) Relative penalty

Figure 13: Effect of penalty on the expected profit for naive bidding for 2003 real dataset. The expected profits fall off more rapidly than IB.

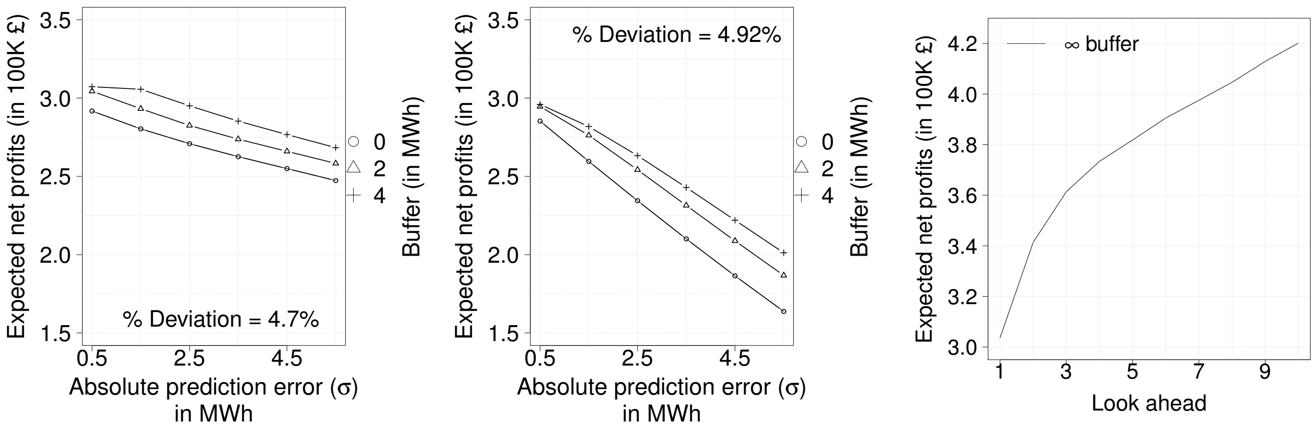

(a) IB (b) Naive (c) Oracle

Figure 10: Profit vs. prediction for realistic buffer capacities and absolute prediction errors for IB, Naive, and Oracle approaches. Data is from year 2005.

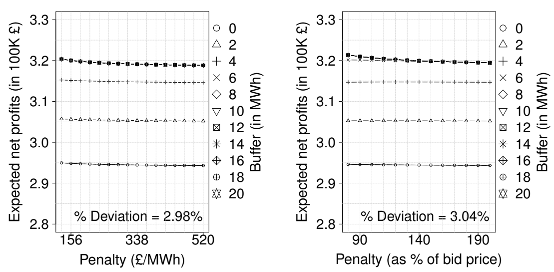

(a) Absolute penalty (b) Relative penalty

Figure 14: Effect of penalty on the expected profit for synthetic datasets and IB. The net profit is almost independent of the penalty as IB hardly pays any penalties.

value, and the penalty paid is close to zero, and thus the actual value of the per-unit penalty price does not affect the net profit. Figures 14a and 14b confirm the expected trend.

7. CONCLUSIONS

Integration of wind energy with the grid is an important problem. Placing a bid in a time-ahead market is a key mechanism for wind power generators to integrate with the grid. A bidding strategy uses predictions about the wind power yield and the price. To do this, we need a bidding strategy and understand the trade-off between profits and the prediction accuracy. We presented a bidding strategy that computes the optimal bid under idealized situations. We cross-validated this approach under synthetic datasets that match the idealized requirements. Next, we identified the modifications required for our idealized bidding strategy to work on real-world datasets. For both synthetic and real-world datasets, we explored the trade-off between the net profits and the prediction errors. On real-world data, we find that buffers can help reduce the impact of prediction errors significantly, but are not required if the errors are within 10% to achieve 83% of baseline profits. Future directions of work include generalizing the strategy for costs associated with buffering and evaluating the profit-prediction trade-off when buffers incur capital expenditure and operational expenses.

8. REFERENCES

[1] APX Power Exchange, UK. UKPX Historical Reference Price Data. http://www.apxgroup.com/market-results/apx-power-uk/ukpx-rpd-historical-data/, 2003.

[2] Melam Bhaskar, Amit Jain, and N Venkata Srinath. Wind speed forecasting: Present status. In *Power System Technology (POWERCON), 2010 International Conference on*, pages 1–6. IEEE, 2010.

[3] E. Bitar, A. Giani, R. Rajagopal, D. Varagnolo, P. Khargonekar, K. Poolla, and P. Varaiya. Optimal contracts for wind power producers in electricity markets. In *Decision and Control (CDC), 2010 49th IEEE Conference on*, pages 1919–1926, 2010.

[4] E.Y. Bitar, R. Rajagopal, P.P. Khargonekar, K. Poolla, and P. Varaiya. Bringing wind energy to market. *Power Systems, IEEE Transactions on*, 27(3):1225–1235, 2012.

[5] A. Botterud, Zhi Zhou, Jianhui Wang, R.J. Bessa, H. Keko, J. Sumaili, and V. Miranda. Wind power trading under uncertainty in LMP markets. *Power Systems, IEEE Transactions on*, 27(2):894–903, 2012.

[6] F. Bourry, J. Juban, L. M. Costa, and Kariniotakis G. Advanced strategies for wind power trading in short-term electricity markets. In *European Wind Energy Conference*, Brussels, Belgium, 2008.

[7] Franck Bourry and George Kariniotakis. Strategies for wind power trading in sequential short-term electricity markets. In *European Wind Energy Conference*, 2009.

[8] D.W.H. Cai, S. Adlakha, and K.M. Chandy. Optimal contract for wind power in day-ahead electricity markets. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 1521–1527, 2011.

[9] USA Department of Energy. Strategies and Decision Support Systems for Integrating Variable Energy Resources in Control Centers for Reliable Grid Operations. 2011.

[10] Antonio Giannitrapani, Simone Paoletti, Antonio Vicino, and Donato Zarrilli. Optimal bidding strategies for wind power producers with meteorological forecasts. Available at: http://knowledgecenter.siam.org/CP2-0022/2, 2013.

[11] Gregor Giebel, Richard Brownsword, George Kariniotakis, Michael Denhard, and Caroline Draxl. The state-of-the-art in short-term prediction of wind power: A literature overview. Technical report, ANEMOS. plus, 2011.

[12] Matthias Lange and Ulrich Focken. State-of-the-art in wind power prediction in germany and international developments. In *Second Workshop of International Feed-In Cooperation*, volume 94, 2005.

[13] Jiaqi Liang, S. Grijalva, and R.G. Harley. Increased wind revenue and system security by trading wind power in energy and regulation reserve markets. *Sustainable Energy, IEEE Transactions on*, 2(3):340–347, 2011.

[14] Nils Lohndorf and Stefan Minner. Optimal day-ahead trading and storage of renewable energies - an approximate dynamic programming approach. *Energy Systems*, 1(1):61–77, 2010.

[15] J.M. Morales, A.J. Conejo, and J. Perez-Ruiz. Short-term trading for a wind power producer. *Power Systems, IEEE Transactions on*, 25(1):554–564, 2010.

[16] P. Pinson, C. Chevallier, and G.N. Kariniotakis. Trading wind generation from short-term probabilistic forecasts of wind power. *Power Systems, IEEE Transactions on*, 22(3):1148–1156, 2007.

[17] L. Puglia, D. Bernardini, and A. Bemporad. A multi-stage stochastic optimization approach to optimal bidding on energy markets. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 1509–1514, 2011.

[18] Ismael Sanchez. Short-term prediction of wind energy production. *International Journal of Forecasting*, 22(1):43–56, 2006.

[19] Saurabh S Soman, Hamidreza Zareipour, Om Malik, and Paras Mandal. A review of wind power and wind speed forecasting methods with different time horizons. In *North American Power Symposium (NAPS), 2010*, pages 1–8. IEEE, 2010.

[20] US Energy Information Administration. International energy outlook. `http://www.eia.gov/forecasts/ieo/pdf/048428201329.pdf`, 2013.

[21] J. Usaola and J. Angarita. Bidding wind energy under uncertainty. In *Clean Electrical Power, 2007. ICCEP '07. International Conference on*, pages 754–759, 2007.

[22] J. Usaola and M.A. Moreno. Optimal bidding of wind energy in intraday markets. In *Energy Market, 2009. EEM 2009. 6th International Conference on the European*, pages 1–7, 2009.

[23] World Wind Energy Association. World wind energy report. `http://wwindea.org/home/index.php`, 2011.

[24] Haifeng Zhang, Feng Gao, Jiang Wu, Kun Liu, and Xiaolin Liu. Optimal bidding strategies for wind power producers in the day-ahead electricity market. *Energies*, 5(11):4804–4823, 2012.

[25] Xiaoxuan Zhang. Optimal wind bidding strategy considering imbalance cost and allowed imbalance band. In *Energytech, 2012 IEEE*, pages 1–5, 2012.