Optimal Generation and Storage Scheduling in the Presence of Renewable Forecast Uncertainties

Nicolas Gast, Dan-Cristian Tomozei, and Jean-Yves Le Boudec

Abstract—Renewable energy sources, such as wind, are characterized by non-dispatchability, high volatility, and non-perfect forecasts. These undesirable features can lead to energy loss and/or can necessitate a large reserve in the form of fast-ramping fuel-based generators. Energy storage can be used to mitigate these effects. In this paper, we are interested in the tradeoff between the use of the reserves and the energy loss. Energy loss includes energy that is either wasted, due to the inefficiency of the storage cycle and the inevitable forecasting errors, or lost when the storage capacity is insufficient. We base our analysis on an initial model proposed by Bejan, Gibbens, and Kelly. We first provide theoretical bounds on the trade-off between energy loss and the use of reserves. For a large storage capacity, we show that this bound is tight, and we develop an algorithm that computes the optimal schedule. Second, we develop a scheduling strategy that is efficient for small or moderate storage. We evaluate these policies on real data from the UK grid and show that they outperform existing heuristics. In addition, we provide guidelines for computing the optimal storage characteristics and the reserve size for a given penetration of wind in the energy mix.

Index Terms—Energy Storage, Forecasting, Scheduling, Stochastic optimization.

NOMENCLATURE

η	Efficiency of a charging/discharging cycle
$B(t)$	Storage level at time t
$B_{\max}$	Energy capacity of the storage system
B^{opt}	Optimal storage capacity
$\text{BGK}(\lambda)$	Policy of [2]: it targets the storage level λ
$C_{\max}/D_{\max}$	Maximal charging/discharging power
C^{opt}	Optimal storage power
$D(t)$	Demand of energy for time slot t
$\text{DO}(\gamma)$	Dynamic offset policy (weight γ)
$\text{FO}(u)$	Fixed offset policy (offset u)
$\bar{G}^{\pi}$	Actual use of reserves for a policy π
$\bar{L}^{\pi}$	Average energy loss for a scheduling policy π
$M(t)$	Missmatch between production and demand
$P_{t-n}^{f}(t)$	Base load production for time t, set at $t-n$
$W(t)$	Wind production at time t
$W_t^{f}(t+k)$	Forecast for time $t+k$, issued at time t
$W_t^{\text{PF}}(t+k)$	Persistence forecast, equal to $W(t)$
$W_t^{\text{WP}}(t+k)$	Forecast obtained by weather prediction
$W_t^{\text{LC}}(t+k)$	Weather forecast with linear correction

All power and energy are expressed in average wind power (AWP) or average wind energy generated during 1h (AWPh). Losses and use of reserves are expressed in percentage of total wind energy.

I. INTRODUCTION

Renewable energy sources, such as wind and solar, are highly volatile and difficult to predict (current forecast techniques for wind power 12 hours in advance have normalized mean absolute errors to the order of 20% [3]). When *renewables* fall short of providing the required power, a large reserve needs to be available in order to avoid blackouts. From a social planner's perspective, it is shown in [4]–[7] that the undesired effects of the volatility of renewables can be mitigated via the use of energy storage, with a manageable increase in energy costs. Another approach for compensating forecast uncertainties is to increase reserves. Algorithms for computing the optimal reserves needed to insure reliability in the presence of forecast errors are developed in [8]–[10]. They update reserve capacities as new forecasts become available.

The economic aspects have also been analyzed. Fabbri et al. show that the cost associated with prediction errors can be up to 10% of the total generation costs [11]. Hence, several authors investigate profit maximization for wind producers. Constantinescu et al. study the optimal scheduling of thermal generators with uncertain wind forecasts [12]. They use a stochastic optimization approach to minimize fuel generation costs. Storage can be used to deal with forecast uncertainties and therefore maximize revenue. The authors of [13]–[16] consider the day-ahead market with time-varying prices and study the sizing and operation of an electrical storage system owned by a wind-producer. Wind producers can also reduce penalties due to deviations from forecasts by aggregating their production to exploit the spatial smoothing of errors. Incentives and optimal redistribution strategies to encourage this behavior are studied in [17].

Storage management can also be done by an independent storage operator via price arbitrage. When the actors on the market are price-takers (*i.e.*, there is a large number of actors), it is shown in [18]–[20] that decoupling generation control and storage management is socially optimal. However, in the case where the price-taking hypothesis does not hold (small number of actors), the authors of [20], [21] show that storage can be overused or underused if it is owned by independent users or by consumers. The question of the viability of independent storage operation is raised in [22]–[24]. Although installing more storage reduces the negative effect of variability, storage owners might have an incentive to undersize their storage system [19]. Moreover, installing solar panels threatens the revenues of pumped-storage hydroplant operators by generating energy during the noon peak [25].

In this paper we are interested in the global impact of the volatility of renewables and how it can be mitigated by

an optimal use of storage. Optimality is defined by adopting the viewpoint of a social planner. We follow the approach of Bejan et al. [2] and address two performance metrics: (1) the energy loss and (2) the actual use of reserves. The energy loss accounts for the inefficiency of the storage cycle, as well as renewable energy that has to be curtailed. The actual use of reserves is considered as a social performance metric as it often consists of fast-ramping fuel-based generation. Our focus is on wind-energy sources, but a similar analysis can be conducted for other types of renewables. We can think of our system model as a national or regional transmission operator, equipped with a storage system (such as pump-hydro) and a large set of wind sources. The operator schedules a certain amount of production one day ahead, and, according to updated forecasts of wind and demand, updates the value of this scheduled production before the actual consumption (of the order of several hours before). In real time, the mismatch between scheduled production, wind production, and demand is compensated using the storage system and the secondary reserve provided by last-minute fast-ramping generators.

The authors of [8]–[10] consider reserve dimensioning in terms of power. Our approach is complementary: in our analysis, we consider that there is always enough reserve power available and evaluate its actual use. Our results can be used to find optimal scheduling policies. In particular, they show that, when the storage capacity is large enough, it is optimal to use a deterministic policy that schedules a certain fixed surplus of generation, that we are able to compute numerically. If a larger surplus is scheduled, a dramatic increase in energy loss occurs, whereas a smaller scheduled surplus increases the use of reserves without significantly reducing energy loss. Thus, although our approach does not explicitly consider market aspects, it can be used by a social planner to determine various market parameters, such as the required secondary reserve.

Bejan et al. [2] propose a production scheduling policy that aims to maintain the storage system at a constant level (e.g., half the storage capacity). This policy can be regarded as a flavor of model predictive control (MPC): in order to make a scheduling decision, it uses the model of the storage to predict its future level, along with updated wind forecast data and past control decisions. It is a quasi-deterministic policy, as forecast errors are implicitly reduced by using the present storage level as a starting point for predicting the future one. The intuition behind this scheduling policy is that a balanced storage system can better cope with mismatches. We question this intuition in our analysis. Like [2], we do not consider network constraints. We discuss the implications of this simplification in Section VI.

Contributions: Under the statistical assumption that the wind forecasts cannot be improved, we show a fundamental bound on the best performance of any production scheduling policy. This bound is determined by the distribution of forecast errors and by the storage system parameters. We show that the quasi-deterministic policy of [2] does not reach this bound. We exhibit a simple deterministic policy ("fixed offset", FO) that achieves this bound when the storage capacity is large. We demonstrate the existence of an optimal operation point, and we develop an algorithm that uses the distribution of forecast errors to compute this point.

For medium and small storage capacity, we use the statistical properties of the wind-forecast errors to build a heuristic production scheduling policy (called "dynamic offset", DO), obtained via stochastic dynamic programming. We perform numerical evaluations on real measures of aggregated demand and wind production, measured in the UK over a period of almost three years. We find that, for various parameters of the storage system, the policy that maintains the storage in a balanced state is suboptimal and is outperformed by DO. In addition, we provide guidelines for computing optimal storage characteristics and the size of the reserve for a given penetration of wind in the energy mix.

A preliminary version of the ideas in this paper can be found in the extended abstract of our talk at the Greenmetrics Workshop [1].

Outline: We begin by describing our storage model and the dataset that we use in Section II. We provide a lower bound on the performance and we show that the fixed offset policy is optimal for large storage capacity in Section III. We study scheduling policies for small and medium storage capacity in Section IV. We introduce the dynamic offset policy and show that it outperforms the other considered heuristics in all scenarios. In Section V, we provide heuristics for computing the optimal storage characteristics and the size of the reserve using the statistics of the forecast errors. Finally, we discuss our results and conclude in Section VI.

Units. All power and energy are expressed in average wind power (AWP) or average wind energy generated during 1h (AWPh). Losses and use of reserves are expressed in percentage of total wind energy over the considered period.

II. SYSTEM MODEL

A. Storage Model

We consider a slotted time model similar to [2]. We assume that at each time step t we observe a certain power demand $D(t)$ that needs to be satisfied during time period t. We use both non-dispatchable and dispatchable energy sources to satisfy the demand. Non-dispatchable energy sources provide an imposed power at each time; a fraction of this power can be discarded. Dispatchable energy sources can be modulated at the beginning of the time slot to exactly match the demand. Specifically, we rely on

1) scheduled base load power production $P_{t-n}^f(t)$, computed n steps ahead (i.e., at time $t-n$),
2) non-dispatchable wind power $W(t)$,
3) energy storage system (e.g., pumped-storage hydro), and
4) fast-ramping generation $G(t)$ (secondary reserve).

The demand is always satisfied by using the above energy sources. However, we might be confronted with energy losses if the aggregate scheduled power production and available wind power are larger than the demand.

In this paper, we analyze various power scheduling policies for computing $P_t^f(t+n)$ at time t, given the predictions of future wind power. Due to inaccurate predictions, the demand $D(t)$ might not match the scheduled and non-dispatchable
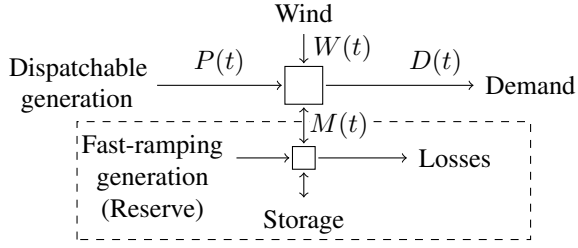
Fig. 1.   The total energy balance: $P + W + M = D$.

generation $P_{t-n}^f(t) + W(t)$. We denote *the mismatch*, i.e., the additional power required to satisfy the demand, by

$$M(t) := D(t) - W(t) - P_{t-n}^f(t). \qquad (1)$$

Note that $M(t)$ can take negative values in the case of over-production, in which case we want to store as much of the surplus power as possible. We also want to use the storage system to compensate a positive mismatch.

We consider that the charging process of the storage has an efficiency η_1 and the discharging process an efficiency η_2. Thus, the storage system has a cycle efficiency $\eta = \eta_1 \eta_2$: only a fraction η of the injected energy can be restored. The factor η accounts for the losses due to charging or discharging, such as losses in transformers or water turbines. The energy capacity, $B_{\max}$, is the quantity of energy that can be retrieved from a fully charged system: the storage can store up to $B_{\max}/\eta_2$ units of energy but can only produce $B_{\max}$ unit of energy. We denote by $B(t) \in [0, B_{\max}]$ the quantity of energy stored at the beginning of time slot t divided by η_2. $B(t)$ represents the net amount of energy that can be extracted and is referred to as the storage level. Due to power constraints, no more than $C_{\max}$ power can be injected in the storage system during a single time-slot and no more than $D_{\max}$ can be generated. We neglect the ramping constraints of the storage system, as they are usually much larger than the ones of conventional generators.

The storage is used to compensate the mismatch $|M(t)|$. Taking into account the constraints described above, we describe the evolution of the storage level by a function[1] of the current storage level and mismatch:

$$B(t+1) = \phi(B(t), M(t)), \qquad (2)$$

$$\phi(B, M) = \begin{cases} (B - \min(M^+, D_{\max}))^+ & \text{if } M \geq 0, \\ \min(B + \eta \min(M^-, C_{\max}), B_{\max}) & \text{if } M < 0. \end{cases}$$

B. Cost Model and Performance Metrics

If $M(t) \geq 0$, then the scheduled power is insufficient. We compensate by using the storage (2) and, only if necessary, the reserve for the remaining unmatched demand:

$$G(t) = M^+(t) - (B(t+1) - B(t))^-. \qquad (3)$$

In this case, the energy loss $L(t)$ during time slot t equals 0.

[1]Throughout the paper, we use the standard notations: $(a)^+ = \max(a, 0)$ and $(a)^- = \max(-a, 0)$.

If $M(t) < 0$, then the scheduled production exceeds current demand, and we store the surplus. There is no use of reserves, hence $G(t) = 0$. The total energy loss $L(t)$ equals

$$L(t) = M^-(t) - (B(t+1) - B(t))^+. \qquad (4)$$

Note that as opposed to [2], the energy loss is composed of both the energy that cannot be stored and the losses due to inefficiency of the storage.

Our cost model assumes that the wind energy is free (wind turbines are already in place), that one unit of energy scheduled in advance costs c_P and that one unit of energy produced by reserve generators costs $c_G > c_P$. As we ignore market aspects, these costs do not vary with time. Over a long time-horizon, the total energy produced is equal to the demand plus the losses: $P + W + G = D + L$. Therefore, the total cost for the operator is equal to

$$c_P P + c_G G = c_P(D - W) + c_P\left(L + \frac{c_G - c_P}{c_P} G\right). \qquad (5)$$

Note that $L = W - D + G + P$. Thus, even if the wind production exceeds the demand, the first term of this equation $D - W$ will be negative but the total cost remains positive.

The demand or the wind cannot be controlled. Hence, to minimize the costs, the operator must minimize a weighted sum of the energy loss and the reserves used. We do not focus on a particular value of $(c_G - c_P)/c_P$ in this paper. We consider the multi-objective problem of finding a policy π that minimizes both the **losses**, $\bar{L}^\pi(T)$, and the actual **use of reserves**, $\bar{G}^\pi(T)$, over time-horizon T. We define

$$\bar{G}^\pi(T) := \frac{\sum_{t=1}^T G(t)}{\sum_{t=1}^T W(t)}; \qquad \bar{L}^\pi(T) := \frac{\sum_{t=1}^T L(t)}{\sum_{t=1}^T W(t)}. \qquad (6)$$

C. Scheduling Policies and Forecast

At each time t, we assume that we are given a forecast $\{W_t^f(t+i)\}_{i=1...n}$ for the wind production at times $t+1$ to $t+n$. We follow [2] and assume that the demand is completely predictable, and that the demand always exceeds the wind production. Hence, mismatches are caused only by wind forecast errors. Our methodology does not rely on a particular forecast method and we do not claim to present a new forecast methodology. Of course, better forecasts will lead to better performance. Hence, our goal is to answer the question: Given a particular forecast, what performance can we achieve and how?

The operator uses the forecasted values of wind production to schedule the power production P. A first approach is to schedule a production at time t equal to the difference in demand and forecast wind production. A specific *scheduling policy* defines a method of computing an *offset* $u_t^f(t+n)$ in the scheduled power. Specifically, the scheduled power production $P_t^f(t+n)$ at time $t+n$ takes the following value:

$$P_t^f(t+n) := (D_t(t+n) - W_t^f(t+n) + u_t^f(t+n))^+. \qquad (7)$$

A negative value of u^f essentially entails dispatching generation at time $t+n$ from the storage or, in the worst, case from the fast ramping generation. A positive value of u^f means that we plan to store energy: we increase the chance of losing energy, but we reduce the use of reserves.

D. Numerical Data Set

To evaluate numerically the performance of the policies, we use data obtained from the BMRA data archive, available at elexonportal.co.uk. This archive is composed of daily reports that contain, among other things, values of aggregated electricity production and consumption in the UK, and weather forecast data. These values are averaged over 30min intervals. Hence, in the simulation, we use a time-step of 30min.

We consider a demand $D(t)$ equal to the aggregated demand and use wind production and day-ahead wind production forecast in the time interval from June 2009 to April 2012. In the considered time frame, the maximum capacity for wind-power generation increased due to the deployment of additional wind farms. In our analysis we normalize the values of production to maintain a constant wind capacity. Hence we use normalized values:

$$W(t) := \frac{\text{production at time } t}{\text{total wind capacity at time } t \times \text{load factor}}. \quad (8)$$

The load factor is the ratio between the average power to the peak power. In our data, this factor is around 28%, which is typical for a good site with modern turbines [26, Chapter 4].

Using this normalization, all units will be expressed in **average wind power** (**AWP**). 1 AWP is equal to the average over time of $W(t)$. Similarly, the unit **AWPh** corresponds to the **average wind energy** generated **during** 1h.

According to the DECC calculator [27], UK produced 16TWh of wind energy in 2010. This means that currently in the UK, one AWP corresponds 2GW. Depending on the scenario, these values should grow to 10GW to 20GW in 2020, 30GW to 30GW in 2035 and 30GW to 120GW in 2050 [27]. As a comparison, the most powerful pumped-storage hydro power station is currently the *Bath County Pumped Storage Station*. It has a power capacity of 3GW, an energy capacity of 30GWh and an cycle-efficiency of about $\eta = 0.77$ [28]. In the UK, the most powerful station is the Dinorwig Power Station. It has a power capacity of 1.7GW, an energy capacity of 10GWh and an cycle-efficiency of $\eta = 0.75$ [29]. This power-plant corresponds to 0.85 AWP and 5 AWPh for today's UK power production or 0.085 AWP to 0.17 AWP and 0.5 AWPh to 1 AWPh for the 2020 scenario. The maximum charging and discharging power is similar for both power-plants. In the numerical evaluation, we will use two values of efficiency: η=0.8, which is close to today's value for pumped-storage, and η=1, which represents an idealistic scenario.

We use three forecast methods in our numerical analysis:

1) W^{PF}: persistence forecast: The first method is to use the average value of wind power during the previous hour to predict any value in the future: $W_t^{\mathrm{PF}}(t+i) := W(t)$. This forecast is the one used by Bejan et al. [2]. It is widely studied in the literature [30], [31] and performs well on short time-scales (see [3] and Figure 3).

2) W^{WP}: numerical weather prediction: The second method is to use the forecast data available in our data set. This data set contains the forecasts of wind-power production, made available each day at 3p.m. for a 24h interval starting at 9p.m. the same day. Thus, the forecast $\hat{W}(t)$ was available between 6h and 30h in advance.
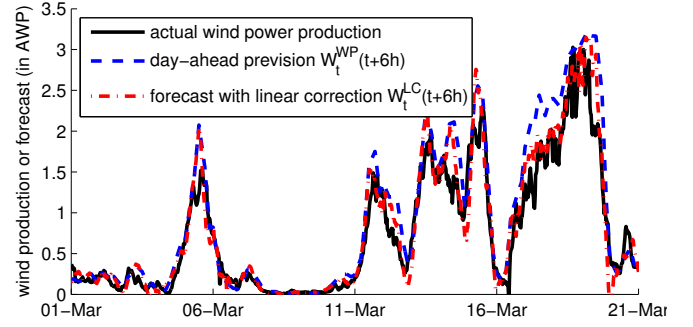

Fig. 2. Typical sample of day-ahead forecast W^{WP} and forecast with linear correction W^{LC} versus actual wind power production (March 2010).

To remove seasonality, we subtract the average forecast error over the previous week and we define

$$W_t^{\mathrm{WP}}(t+i) := \hat{W}(t+i) - \frac{1}{7\text{days}} \sum_{t-7\text{days}<s\leq t} (\hat{W}(s)-W(s))$$

These values are well defined for $i \leq 30$h.

3) W^{LC}: weather prediction with a linear correction: This forecast method is a combination of the two preceding ones. It exploits the fact that weather prediction errors are positively correlated. The forecast for time $t+i$ is obtained by subtracting a fraction α_i of the current weather forecast error:

$$W_t^{\mathrm{LC}}(t + i) = W_t^{\mathrm{WP}}(t + i) - \alpha_i(W_{t-i}^{\mathrm{WP}}(t) - W(t)). \quad (9)$$

The coefficient α_i depends on the time horizon i. It is chosen in order to minimize the mean absolute error, defined in Equation (10). Its value is computed using data from the 6 months preceding the considered evaluation interval.

Accuracy of the forecasts: We plot a typical sample of both the wind production and the forecasts W^{WP} and W^{LC} in Figure 2. The forecast W^{LC} is often better than W^{WP}, but sometimes it over-corrects (*e.g.*, on March, 21st). We measure the accuracy of a forecast as the normalized mean absolute error:

$$\sum_{t\in 3\text{ years}} \left|W^f(t+i) - W(t+i)\right| \Big/ \sum_{t\in 3\text{ years}} W(t). \quad (10)$$

The mean absolute errors of the three forecasts are reported in Figure 3. As expected [3], the persistence forecast outperforms the weather prediction for short time-horizon ($i \leq$ 6h). W^{LC} is always the best forecast. The performance of W^{WP} does not depend on the time horizon considered because it uses forecasts from the day-ahead. We also evaluated more complicated statistical corrections (non-linear or with a larger time-horizon) than W^{LC}, but they were not significantly better.

The parameters values used in the simulation are summarized in Table I.

TABLE I
VALUES OF THE PARAMETERS USED IN THE SIMULATION

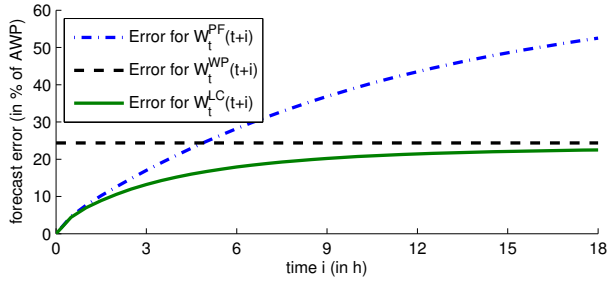| Parameter	Range
Wind production and forecast data	June 2009 to April 2012
Storage energy capacity $B_{\max}$	$\in [1\,\mathrm{AWPh}; 20\,\mathrm{AWPh}]$
Storage maximum power $C_{\max} = D_{\max}$	$\in [0.1\,\mathrm{AWP}; 1.5\,\mathrm{AWP}]$
Cycle efficiency η	$\eta = 0.8$ or $\eta = 1$.

Fig. 3. Mean absolute error of the three wind forecasts, W^{PF} (persistence forecast), W^{WP} (day-ahead weather forecast) and W^{LC} (weather forecast with linear correction) as a function of the horizon of the prediction i.

III. OPTIMALITY OF THE FIXED OFFSET POLICY FOR LARGE STORAGE CAPACITY

At time t, the fixed offset policy $\mathrm{FO}(u)$ schedules a production at time $t+n$ equals to the forecasted mismatch between demand and wind power plus a constant power u:

$$P_t^{\mathrm{FO}}(t+n) := D(t+n) - W_t^f(t+n) + u. \quad (11)$$

A higher offset u results in higher losses but a lower use of reserves. As we will see next, this policy is optimal for large storage capacity.

A. A Lower Bound on the Energy Cost

The quantity $P_t^f(t+n)$ of energy production scheduled for time $t+n$ is set at time t, with the information[2] available at time t. A scheduling policy π defines a method of computing the additional scheduled power $u_t^f(t+n)$ based on all the information available at time t. For such a scheduling policy π, we denote by $\bar{G}^\pi(T)$ and $\bar{L}^\pi(T)$ the average values of the use of reserves and energy loss over time horizon $[1;T]$.

Let $e_t^f(t+n) = W(t+n) - W_t^f(t+n)$ denote the forecast error. The next theorem provides a bound on the performance of any scheduling policy, assuming that the prediction cannot be improved by using the knowledge up to time t and that errors are identically distributed. By this, we mean that there exists a random variable ε such that the forecast error $e_t^f(t+n)$ is independent of the passed values of predictions or generation and $e_t^f(t+n)$ is distributed as ε. Note that we do not assume $e_t^f(t+n)$ to be independent of $e_t^f(t+1)$ or $e_t^f(t+n-1)$.

For a fixed offset u, let $\ell(u)$ and $g(u)$ be defined by

$$\ell(u) := \mathbb{E}\left[(\varepsilon+u)^+\right] - f(u) \quad (12)$$

$$g(u) := \mathbb{E}\left[(\varepsilon+u)^-\right] - f(u) \quad (13)$$

where f is equal to $f(u) = \min(\eta\mathbb{E}\left[\min\left((\varepsilon+u)^+, C_{\max}\right)\right],$ $\mathbb{E}\left[\min\left((\varepsilon+u)^-, D_{\max}\right)\right])$.

Theorem 1 (Lower Bound on the Energy Cost): Let us assume that for all t, the forecast error $e_t^f(t+n)$ is independent of $\mathcal{F}_t$ and distributed as a random variable ε. Then, for any control policy π, there exists u such that

$$\bar{G}^\pi(T) \geq g(u) - \frac{B_{\max}}{T} \text{ and } \bar{L}^\pi(T) \geq \ell(u) - \frac{B_{\max}}{T}. \quad (14)$$

[2]The information available at time t is denoted $\mathcal{F}_t$. $\mathcal{F}_t$ is also called the filtration associated with the process. $\mathcal{F}_t$ represents the knowledge of wind production and prediction up to time t as well as production scheduled for time t to $t+n-1$ and prediction for time t to $t+n$.
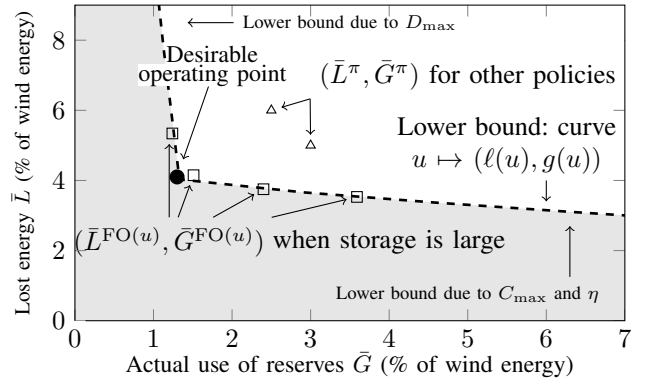

Fig. 4. Illustration of Theorems 1 and 2. The performance of each policy is represented by a point $(\bar{L}, \bar{G})$ where $\bar{L}$ is the average lost energy and $\bar{G}$ the average use of reserves. The dashed curve is the lower bound of Theorem 1. By Theorem 1, the performance of a policy cannot be in the gray zone without improving the forecast. By Theorem 2, the fixed offset policy performs close to optimally for large storage capacity: its performance is on the dashed curve.

Moreover, the function $g(u) \mapsto \ell(u)$ is well defined, decreasing and convex.

Proof: The proof is given in Appendix A. ∎

When the time horizon T is large, the term $B_{\max}/T$ becomes negligible. Thus, this theorem shows that the performance of any policy will be above the curve $u \mapsto (\ell(u), g(u))$.

Theorem 1 provides a lower bound on the energy cost. The next theorem shows that this bound is attained by the fixed offset policy when the storage capacity is large.

Let $\bar{L}^{\mathrm{FO}(u)}$ be the average lost energy when using the fixed offset policy $\mathrm{FO}(u)$: $\bar{L}^{\mathrm{FO}(u)} = \limsup_{T \to \infty} \bar{L}^{\mathrm{FO}(u)}(T)$ and $\bar{G}^{\mathrm{FO}(u)} = \limsup_{T \to \infty} \bar{G}^{\mathrm{FO}(u)}(T)$ the average use of reserves.

Theorem 2 (The fixed policy is optimal for large storage): For all $\delta > 0$, there exists a storage capacity B_0 such that for all u, if $B_{\max} \geq B_0$, we have

$$\bar{G}^{\mathrm{FO}(u)} \leq g(u) + \delta, \quad \text{and} \quad \bar{L}^{\mathrm{FO}(u)} \leq \ell(u) + \delta \quad (15)$$

Proof: The proof is given in Appendix B. ∎

This theorem, combined with Theorem 1, shows that when the storage capacity is large, the FO policies perform at least as well as any other policy. This fact is illustrated in Figure 4: the performance of the fixed offset policy is close to the lower bound. These policies are Pareto-optimal: the performance of any other policies must be above this curve.

B. Desirable Operating Point

All the points of the curve corresponding to the lower bound of Theorem 1 are Pareto-optimal points: for point on this curve and any storage capacity, there exists no policy that simultaneously uses less reserves and wastes less energy. Moreover, as shown in Figure 4, this curve exhibits a sharp knee: it is composed of two branches that intersect in a single point, that we denote the *desirable operating point*.

This point is desirable in the sense that for a large range of weight factors γ, it minimizes the weighted sum $L + \gamma G$. Because of the particular shape of the curve, any policy that decreases the use of reserves by a few percents will at least double its energy loss. This indicates that any reasonable

policy should operate close to this point. The value of this point depends on the power capacity and cycle-efficiency of the storage.

C. Optimal Fixed Offset

When the storage capacity is large, the fixed offset policy achieves the lower bound. FO can operate close to the desirable operating point when the offset u is such that on average, we charge $\eta \mathbb{E}\left[\min\left((\varepsilon+u)^+, C_{\max}\right)\right]$ as much as we discharge $\mathbb{E}\left[\min\left((\varepsilon+u)^-, D_{\max}\right)\right]$. This shows that the optimal fixed offset is largely independent of the cost ratio between energy produced using reserve generators and scheduled energy. Algorithm 1 uses this fact to provide an efficient way of computing a fixed offset u^* that is optimal for most values of γ.

Algorithm 1 Computation of the optimal fixed offset

Require: Samples of $W^f - W$, $C_{\max}$, $D_{\max}$, η.
  Let $\varepsilon_t := W_t^f(t+n) - W(t+n)$
  Use a binary search to find u such that:
$$\eta \sum_t \min((\varepsilon_t+u)^+, C_{\max}) = \sum_t \min((\varepsilon_t+u)^-, D_{\max}).$$
  Return u.

IV. HEURISTICS FOR LOW AND MODERATE STORAGE CAPACITIES

When the storage capacity is low, the fixed offset policy performs far from the lower bound as it does not take into account the level of the storage. In this section, we describe two scheduling policies (BGK and DO) that choose $u_t^f(t+n)$ as a function of the estimate $B_t^f(t+n)$ of the storage level at time $t+n$. This estimate is computed by using all the previously computed decisions $P_{t-n}^f(t), \ldots, P_{t-1}^f(t+n-1)$ and wind forecasts available at time t: $B_t^f(t+1) := B(t+1)$ and

$$B_t^f(t+i+1) := \phi(B_t^f(t+i), D(t+i) - P_{t+i-n}^f(t+i) - W_t^f(t+i)).$$

A. The Bejan-Gibbens-Kelly Policy (BGK)

The Bejan-Gibbens-Kelly policy $\text{BGK}(\lambda)$, introduced in [2], aims at maintaining a constant level of the storage $\bar{B} = \lambda B_{\max}$ (typically, $\lambda = 0.5$). To compute $P_t^{\text{BGK}}(t+n)$ at time t, this policy first forecasts the storage level at time $t+n$. Subsequently, $u_t^{\text{BGK}}(t+n)$ is computed as the required energy to bring the storage level at time $t+n+1$ closest to $\bar{B}$ under the operating constraints:

$$P_t^{\text{BGK}}(t+n) := D(t+n) - W_t^f(t+n) + f_\lambda(B_t^f(t+n)),$$
$$\text{where} \quad f_\lambda(B) = \min(\tfrac{1}{\eta}(\lambda B_{\max} - B)^+, C_{\max})$$
$$- \min((\lambda B_{\max} - B)^-, D_{\max}).$$

B. The Dynamic Offset Policy (DO)

The dynamic offset policy chooses decision $u_t^{\text{DO}}(t+n)$ as a deterministic precomputed function $\delta_\gamma : [0, B_{\max}] \to \mathbb{R}$ of the forecast storage level $B_t^f(t+n)$:

$$P_t^{\text{DO}}(t+n) := D(t+n) - W_t^f(t+n) + \delta_\gamma(B_t^f(t+n)), \quad (16)$$

The control law δ_γ is derived from the optimal policy of a Markov decision process (MDP) that represents a simplified storage model. It takes into account the statistics of forecast errors made on the wind $W^f(t+n)$ and on the storage level $B_t^f(t+n)$. In the remainder of the paper, we show that it performs at least as good as the two previous heuristics in all scenarios studied and outperforms them in most of them.

1) Description of the MDP: We consider a MDP, whose state is the storage level $[0, B_{\max}]$, and whose action space is $\mathbb{R}$. Let ε be an *i.i.d.* sequence of variables on $\mathbb{R}$, with the same distribution as the errors of the corrected forecast $W(t+n) - W_t^f(t+n)$; and let ξ be an *i.i.d.* sequence of variables on $\mathbb{R}$, with a similar distribution as the error made on the storage level $\sum_{i<n}(W(t+i) - W_t^f(t+i))$. If $B(t)$ is the state at time t, and $u(t)$ the offset chosen for the next time slot, the state at time $t+1$ becomes

$$B(t+1) = \phi\Big(\big(B(t) + \xi(t+1)\big)_0^{B_{\max}}, -\varepsilon(t+1) - u(t)\Big), \quad (17)$$

The function ϕ represents the evolution of the storage in one time step and is defined by Equation (2). The notation $(x)_0^{B_{\max}}$ denotes a truncation of x: $(x)_0^{B_{\max}} = \max(0, \min(B_{\max}, x))$.

Given a storage level $b := x + \xi$ and a mismatch $m := -u - \varepsilon$, we consider an instantaneous cost that is a weighted sum of the lost energy and the use of reserves:

$$c(b, m) = m^-(1-\eta) + \left(m + \eta \min\left(\frac{B_{\max} - b}{\eta}, C_{\max}\right)\right)^-$$
$$+ \gamma\left(m - \max(-b, -D_{\max})\right)^+ \quad (18)$$

The first line corresponds (for its first term) to the losses due to inefficiencies in the storage system, and (for its second term) to the losses due to ramp-up constraints or overflows. The second line corresponds to the use of reserves (derived via Equation (3)). The weight γ characterizes the trade-off and indexes the family of policies.

We consider a controller that seeks to minimize the average cost over an infinite time horizon:

$$\min_{u \in \text{policy}} \left\{ \lim_{T \to \infty} \frac{1}{T} \mathbb{E}\left[\sum_{t=1}^T c((B(t)+\xi)_0^{B_{\max}}, u(t)+\varepsilon)\right] \right\}. \quad (19)$$

The choice of $u(t)$ can depend on the value of $B(t)$, as well as all values of $B(s)$ and $u(s)$ until time $s < t$.

2) Numerical computation of the control law δ_γ: We consider a discretized version of the MDP, by replacing the state space $[0, B_{\max}]$ by $1 + (B_{\max}/h)$ values: $\{0, h, 2h, \ldots, B_{\max}\}$. By [32, Theorem 8.1.2], there exists an optimal control law δ_γ that only depends on $B(t)$. Moreover, this MDP belongs to the class of communicating MDPs with average cost. This guarantees that the relative value iteration algorithm (Algorithm 2) converges to v (see [32, Chapter 9]). It provides an efficient way to numerically compute the control law δ_γ.

In our numerical evaluations, we apply Algorithm 2 with a precision threshold of 10^{-7}. To accelerate the computation, we restrict the search to control laws that decrease in u. An illustration of the optimal control law δ_γ as computed by Algorithm 2 is shown in Figure 5. We see that the control law δ_γ is less aggressive than BGK and reacts more smoothly.

Algorithm 2 DO: computation of the control law δ_γ.

Require: Trade-off γ, step-size h, statistics of errors ξ and ε.

$k \leftarrow 0$; $V_0(b) \leftarrow 0$ for all $b \in \{0, h, 2h, \ldots, B_{\max}\}$

while $\sup_b |V_k(b) - V_{k-1}(b)| \geq$ precision threshold **do**

for all $b \in \{0, h, 2h, \ldots, B_{\max}\}$, **do**

$$\bar{V}_{k+1}(b) \leftarrow \inf_{u \in \mathbb{R}} \mathbb{E}\Big[c((b+\xi)_0^{B_{\max}}, u+\varepsilon)$$
$$+ V_k(\phi((b+\xi)_0^{B_{\max}}, u+\varepsilon))\Big]$$

$$V_{k+1}(b) \leftarrow \bar{V}_{k+1}(b) - \inf_b \bar{V}_{k+1}(b)$$

end for

$k \leftarrow k+1$.

end while

$$\delta_\gamma(b) := \arg\min_{u \in \mathbb{R}} \Big\{\mathbb{E}\big[c((b+\xi)_0^{B_{\max}}, u+\varepsilon) + V_k(\phi(b+e, u+\varepsilon))\big]\Big\}$$
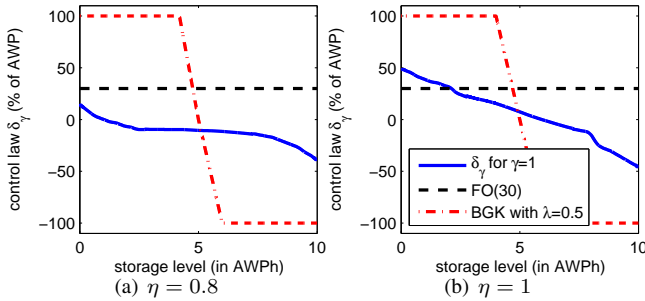

Fig. 5. Offset set by the different policies as a function of the forecasted storage level $B_t(t+n)$. We plot the three policies FO(30AWP), BGK(0.5) and DR(1). The offset set by FO does not depend of the forecasted storage level. The offset of BGK varies abruptly around $0.5 B_{\max}$. DO reacts more smoothly. The parameters of the storage system are $B_{\max}$=10 AWPh and $C_{\max}=D_{\max}$=1 AWP. The offset are expressed in percentage of average wind production (AWP).

The policy DO differs from the dynamic reserve (DR) policy described in our previous work [1] through Equation (17). This equation includes the term $\xi(t+1)$, *i.e.*, the error on the forecasted storage level. DR can be computed using Algorithm 2 by setting $\xi(t+1) = 0$. Although we do not show the comparison, DO performs at least as well as DR in all our examples.

C. Performance for Various Storage Characteristics

To evaluate the performance, we simulate the fixed offset policy (FO), dynamic offset policy (DO), and the BGK policy. Figure 6 reports the results for a small storage capacity $B_{\max}$=3 AWPh and a large storage capacity $B_{\max}$=20 AWPh. Each time, we set $C_{\max}$=0.30 AWP and we simulate η=0.8 and η=1. We apply the BGK policy with the three wind forecasts W^{PF}, W^{WP}, and W^{LC}. For the dynamic offset policy, we use the corrected forecast W^{LC} and we vary its parameter γ from 0.01 to 100.

We make three important observations. First, the plots show that, for all tested values of storage characteristics, the BGK policies, which try to maintain the storage at a fixed level, are outperformed by the dynamic offset policies with corrected forecast. Even when $\eta = 1$, the DO policy reduces both the energy loss and the use of reserves. This suggests that taking

$$\eta = 0.8 \qquad\qquad \eta = 1.$$
(a) $B_{\max}$=3 AWPh and $C_{\max}=D_{\max} = 0.30$ AWP

$$\eta = 0.8 \qquad\qquad \eta = 1$$
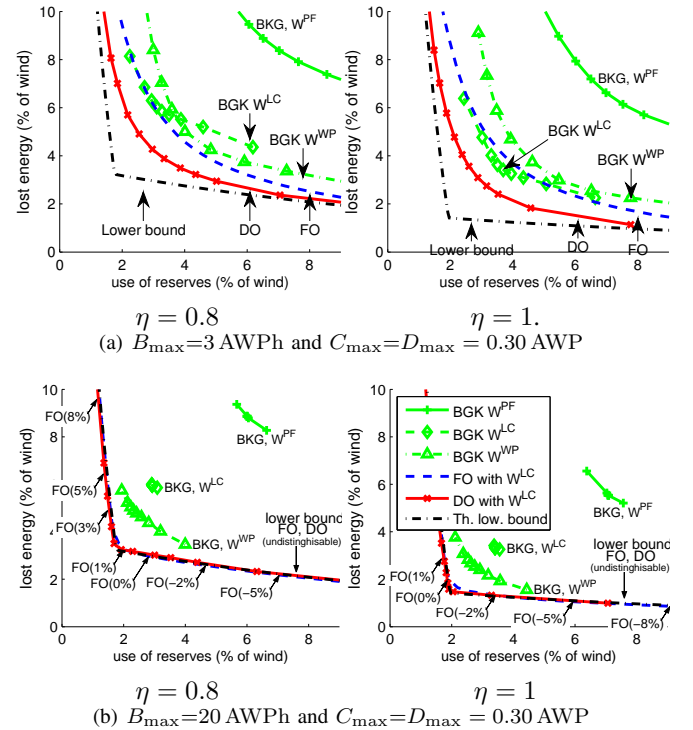(b) $B_{\max}$=20 AWPh and $C_{\max}=D_{\max} = 0.30$ AWP

Fig. 6. Performance of the fixed offset policy, the BGK policy and the dynamic offset policy for various storage characteristics. The left plots are for an efficiency of $\eta = 0.8$ while the right plots are for $\eta = 1$. Values are expressed in percentage of average wind production (AWP).

into account the statistical nature of the prediction error leads to performance gains higher than trying to maintain the storage in a balanced state. For these storage capacities, the fixed offset policy is not optimal.

Second, when the storage capacity is large (20 AWPh here), both the fixed offset policy and the dynamic offset achieve the lower bound. The BGK policy is not optimal for both η=0.8 and η=1 and loses more energy than the best fixed offset policy. In this case, the performance of BGK is best when run with W^{WP}, but it is still away from the lower bound.

Finally, as we see in the different graphs, BGK can perform better either with W^{WP} or W^{LC}, depending on the storage size or the target storage level λ. For example, when $n = 6h$, $B_{\max} = 3$ AWPh and $C_{\max} = D_{\max} = 0.30$ AWP, high threshold favors BGK run with W^{LC}, whereas low thresholds favor BGK run with W^{WP}.

D. Consistency of the Dynamic Offset Policy

As shown in Figure 6, the performance of BGK does not necessarily improve as the forecast improves: BGK with W^{WP} performs better than BGK with W^{LC} that, in turn, performs better than BGK with W^{PF}, although the relative error exhibits the opposite order for $n \leq 6h$ (see Figure 3). We attribute this performance mismatch to the fact that BGK performs better when the prediction does not depend on the time at which the prediction was issued. As the control law of BGK is sharper than the one of DO (see Figure 5), a time-varying prediction causes BGK to overreact to certain situations.

To verify this assumption, we construct two artificial scenarios A and B. For the two scenarios, the wind generation

is equal to its real value $W(t)$. We define the two forecasts W^A, W^B as follows. Let $(L(t))_t$ be a sequence of *i.i.d.* Laplace random variables of mean 0. The forecasted wind production issued at time t for time $t+i$ is $W_t^A(t+i) = W(t+i) + \sum_{k=0}^i L(t+k)$ and $W_t^B(t+i) = W(t+i) + \sum_{k=i-n}^i L(t+k)$.

The prediction error of W^A is $\sum_{k=0}^i L(t+k)$, which is smaller than the one corresponding to W^B. For a given time s, the forecast $W_t^B(s)$ does not depend on time t at which it has been issued, whereas the forecast $W_t^A(s)$ does.
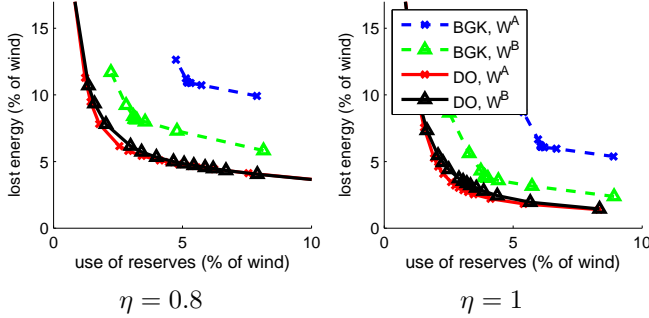

Fig. 7. Performance of BGK and DO for the two forecasts W^A and W^B. Although the forecast W^A is better than the forecast W^B, BGK performs better with W^B. The performance of DO is slightly better with W^A. All values are expressed in percentage of average wind production (AWP). The storage capacity is $B_{\max} = 10$ AWPh and the maximum power is $C_{\max} = D_{\max} = 0.5$ AWP.

Figure 7 reports the values of loss energy and fast ramping generation used for a capacity $B_{\max} = 10$ AWPh, $C_{\max} = 0.5$ AWP and two efficiencies: $\eta = 0.8$ or $\eta = 1$. We observe that in each case, the performance of BGK degrades when we use the better prediction W^A. The dynamic offset policy always outperform BGK and does not exhibit this problem: it performs (slightly) better with the better forecast W^A. This is why we say that DO is consistent: a smaller error results in a better performance.

V. OPTIMAL STORAGE CHARACTERISTICS

A. Optimal Charging and Discharging power ($C_{\max}$, $D_{\max}$)

Let us define the global performance of a policy π by the sum of the energy loss plus the total use of reserve:

$$\text{Global performance of } \pi := \bar{L}^\pi + \bar{G}^\pi. \quad (20)$$

Figure 8 depicts the global performance of the dynamic offset policy DO run with W^{LC} and the BGK policy run with W^{WP}. For BGK, we chose W^{WP} as it often outperforms BGK with W^{LC} (see Figure 6). For each point, we choose the parameter (γ or λ) that minimizes $\bar{L} + \bar{G}$. The time horizon is $n = 6h$.

As observed in the previous sections, DO performs always at least as well as BGK. It significantly outperforms BGK for low storage capacities. Moreover, the performance improves dramatically for small power capacities and does not change much when $C_{\max} > 0.80$ AWP. This indicates that 0.80 AWP is enough to mitigate forecast error for $6h$. Note that this optimal ramping constraints does not depend on η.

Increasing the ramping constraints $C_{\max}$ does not improve the performance when the forecast error has a small probability
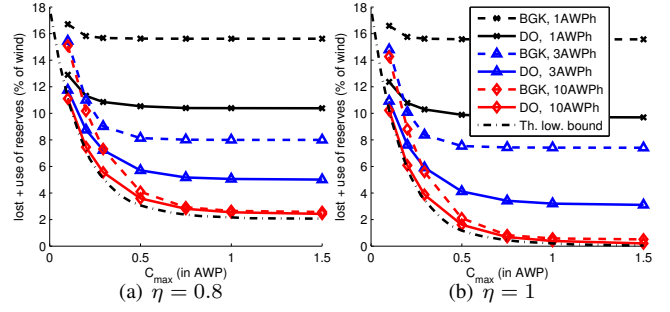

Fig. 8. Global performance (Eq.(20)) of the dynamic offset policy and of BGK as a function of $C_{\max} = D_{\max}$. We see that the knee of the performance is close 0.80 AWP. All values are expressed in percentage of average wind production (AWP).

of exceeding the ramping constraints. Therefore, we define the optimal ramping constraints of the storage as

$$C^{\mathrm{opt}}(n) = \min\{c \text{ s.t } \mathbb{P}\left(|\varepsilon| > c\right) \le 1\%\}. \quad (21)$$

In this equation, ε is defined as in §IV-B1: it is distributed as the forecast error $W(t+n) - W_t^{LC}(t+n)$ and depends on n.

Figure 9(a) indicates that this optimal power capacity converges to 1 AWP as n goes to infinity. It quickly reaches 0.80 AWP for $n = 6h$ and then saturates. This is consistent with the relative error of W^{LC} shown on Figure 3.

Using statistical analysis, the authors of [31] argue that the power capacity needed to compensate for persistence forecast errors is between 50% and 100% of AWP for $n=1h$ and 150% to 250% for $n=12h$. The values reported in Figure 9 are significantly smaller because W^{LC} is significantly better than W^{PF}. Applying our methodology to W^{PF} leads to 50% for $n = 1h$ and 190% for $n = 12h$, close to the number of [31].

B. Optimal Storage Capacity

The storage capacity influences the lost energy or the reserve use when there is an overflow or an underflow of the storage system. As the storage capacity increases, the probability of such an event occurring decreases. As we observe on Figure 10, this decrease is significant for low values of $B_{\max}$ but saturates for high capacity. For example, when $C_{\max} = 0.30$ AWP, $\eta = 0.8$ and $n = 6h$, increasing the storage capacity to more than 10 AWPh has a negligible effect.
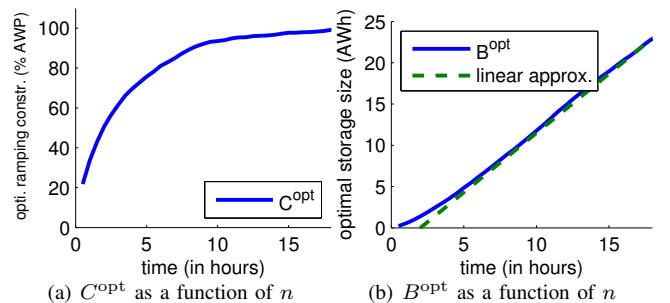

Fig. 9. Optimal storage characteristics C^{opt} and B^{opt} as a function of n, given by Equation (21) and Equation (23).
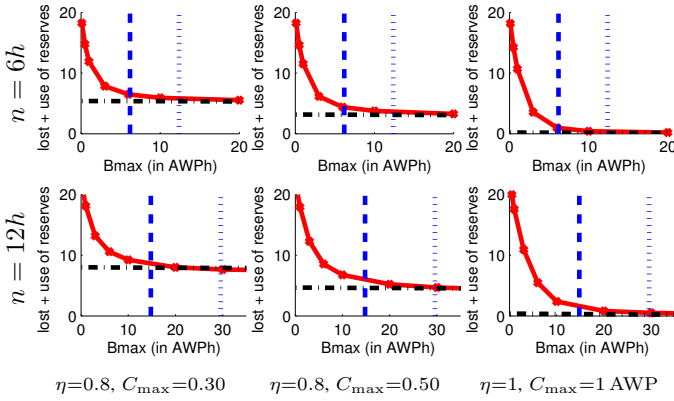
Fig. 10. Global performance (Eq.(20)) of the policy DO as a function of $B_{\max}$. A different value of n is associated to each row and a different value of $C_{\max}=D_{\max}$ and η to each column. In each case, the knee of the performance is between the two vertical lines which represents $B^{\text{opt}}(n)$ and $2B^{\text{opt}}(n)$, defined in Eq.(23). The knee does not depend on $C_{\max}$ or η.

If we remove power constraints and efficiency, the control problem (17) becomes

$$B(t+1) = \max(\min(B(t) + \xi + \varepsilon + u, B_{\max}), 0). \quad (22)$$

By setting $u := B_{\max}/2 - B(t)$, the probability of having an overflow or an underflow is $\mathbb{P}(2|\xi + \varepsilon| > B_{\max})$. This suggests that if $B_{\max}$ is such that this probability is small, there is no need to increase the storage capacity.

Therefore, we define the optimal storage capacity as

$$B^{\text{opt}}(n) = \min\{b \text{ s.t } \mathbb{P}(2|\xi + \varepsilon| > b) \leq 1\%\}. \quad (23)$$

where ξ and ε are defined as in §IV-B1, and depend on n.

Figure 10 shows a good match between this heuristic and the observed value of the optimal storage: the observed optimal storage capacity lies between B^{opt} and $2B^{\text{opt}}$.

A numerical evaluation of B^{opt} given by Equation (23) is depicted in Figure 9(b). This indicates that the optimal storage capacity grows sub-linearly for $n < 6h$ and then increases by about 1.5 AWPh per hour of time horizon for $n \geq 6h$. This corresponds to 150% of the average wind energy produced in one hour, or 40% of the peak production.

VI. CONCLUSION AND FUTURE WORK

In this paper, we adopt the viewpoint of a social planner and investigate efficient ways of integrating renewables in the power mix with the aid of fast-reacting energy storage systems characterized by a cycle efficiency $\eta < 1$. The cost is decomposed into two metrics: the amount of used reserves versus the wasted renewable energy. In Theorem 1, we derive a fundamental lower bound on the cost of any production scheduling policy, given a certain irreducible stochasticity in the forecast errors of the renewable production. We define FO, a deterministic policy that schedules an amount of production equal to the production mismatch, i.e., the difference between the forecast demand and the forecast renewable output, offset by a fixed value. We use the statistics of forecast errors to construct a second policy (DO) through stochastic optimization. The policy DO schedules an amount of production equal to the production mismatch, offset by a dynamic value that

depends of the forecasted storage level. When the available energy capacity of storage is sufficiently large, we prove that the policy FO matches the theoretical lower bound. In all cases, DO outperforms all tested policies, in particular the policy that aims to keep the storage level balanced (at half the capacity) [2]. We show that there exists a desirable operating point for the system, which achieves a trade-off between the two metrics of interest. This point can be determined using the heuristic provided in Algorithm 1. The proposed policies are evaluated on real data collected in the UK made available by ELEXON. This study can serve for designing an efficient energy market that accounts for the available storage, more specifically, incentive and pricing mechanisms that match the socially-optimal performance.

Our dynamical storage management policy, DO, could also be adapted for profit maximization in a market environment with variable or uncertain prices. The only difference in the construction of the the stochastic optimization problem is that the instantaneous cost (18) needs to be replaced by a time-varying cost function that depends on instant prices. When storage is large, we conjecture that the cost of the policy DO attains the lower bound (like the deterministic FO policy). This conjecture is left for further study.

In order to keep the model tractable and to develop a generic methodology, we do not consider the network physical constraints. They intervene in a situation where the transmission network is unable to balance a local overproduction of renewable sources and an underproduction in a remote location. In this case, the placement of storage becomes crucial. Our methodology can be used in a multiple-stage optimization problem: first by using local storage to solve local imbalance, and finally by aggregating these results.

APPENDIX A
PROOF OF THEOREM 1

Proof: Equations (12) and (13) come from two terms: one is due to the charging constraint $C_{\max}$ and storage inefficiencies and the other to discharging constraint $D_{\max}$.

By definition (4), if $M(t) := -e^f_{t-n}(t) - u(t)$ is negative, we store energy and the energy loss between t and $t+1$ is $M(t)$ minus what is stored in the dam[3]. If we neglect the overflows due to $B_{\max}$, we obtain the lower bound (24):

$$L(t) = (M(t))^- - (B(t+1) - B(t))^+$$
$$\geq (M(t))^- - \eta \min((M(t))^-, C_{\max}). \quad (24)$$

Similarly, the use of reserves is

$$G(t) \geq (M(t))^+ - \min((M(t))^+, D_{\max}) \quad (25)$$

For any $x \in \mathbb{R}$, let us define $C(x)$ and $D(x)$ by:

$$C(x) = \eta\mathbb{E}\left[\min((\varepsilon + x)^+, C_{\max})\right] \quad (26)$$
$$D(x) = \mathbb{E}\left[\min((\varepsilon + x)^-, D_{\max})\right]. \quad (27)$$

To ease the notation, we denote by $u(t)$ the additional power $u^f_{t-n}(t)$ that was scheduled at time $t-n$. As $u(t)$ was

---

[3]Recall that we denote by $(M(t))^-$ the negative part of $M(t)$, which is equal to 0 if $M(t) \geq 0$ and to $-M(t)$ otherwise.

scheduled using information available at time $t-n$, it is $\mathcal{F}_{t-n}$-measurable. Therefore, using that $M(t) = -e_{t-n}^f(t) - u(t)$ and that $e_{t-n}^f(t)$ is independent of $\mathcal{F}_{t-n}$, Equation (24) implies

$$
\begin{aligned}
\mathbb{E}\left[L(t)\right] &= \mathbb{E}\left[\mathbb{E}\left[L(t)|\mathcal{F}_{t-n}\right]\right] \\
&\geq \mathbb{E}\Big[\mathbb{E}\big[(e_{t-n}^f(t)+u(t))^+ \\
&\qquad - \eta\min((e_{t-n}^f(t)+u(t))^+, C_{\max})|\mathcal{F}_{t-n}\big]\Big] \\
&= \mathbb{E}\left[(\varepsilon + u(t))^+ - \eta\min((\varepsilon + u(t))^+, C_{\max})\right] \\
&= \mathbb{E}\left[(\varepsilon + u(t))^+\right] - D(u(t)).
\end{aligned}
\tag{28}
$$

For a fixed value ε, $x \mapsto (\varepsilon+x)^+ - \eta\min((\varepsilon+x)^+, C_{\max})$ is convex. Therefore, the function $x \mapsto \mathbb{E}\left[(\varepsilon+x)^+\right] - C(x)$ is convex. Combined with Equation (28), this shows that

$$
\begin{aligned}
\bar{L}(T) = \frac{1}{T}\sum_{t=1}^T \mathbb{E}\left[L(t)\right] &= \frac{1}{T}\sum_{t=1}^T \left[\mathbb{E}\left[(\varepsilon + u(t))^+\right] - D(u(t))\right] \\
&\geq \mathbb{E}\left[(\varepsilon+\bar{u})^+\right] - C(\bar{u})
\end{aligned}
\tag{29}
$$

where the last inequality comes from Jensen's inequality and $\bar{u} = \sum_{t=1}^T u_t^f(t+n)/T$.

Starting from (25) and using similar ideas for, we also have

$$
\bar{G}(T) \geq \mathbb{E}\left[(\varepsilon+\bar{u})^-\right] - D(\bar{u}).
\tag{30}
$$

Using the energy balance equations (3) and (4), we have

$$
B(t+1) - B(t) = G(t) - L(t) - M(t).
\tag{31}
$$

Summing this equation from $t=1$ to T implies that

$$
\bar{G}(T) = \bar{L}(T) + \frac{1}{T}\sum_{t=1}^T \mathbb{E}\left[M(t)\right] + \frac{1}{T}(B(T) - B(0)).
\tag{32}
$$

As $|B(0) - B(T)| \leq B_{\max}$, combining this with (29) and Equation (30) shows that

$$
\begin{aligned}
\bar{G}(T) &\geq \bar{L}(t) - \mathbb{E}\left[\varepsilon+\bar{u}\right] - B_{\max}/T \tag{33} \\
&= \mathbb{E}\left[(\varepsilon+\bar{u})^-\right] - C(\bar{u}) - B_{\max}/T \tag{34} \\
\bar{L}(T) &\geq \bar{G}(t) + \mathbb{E}\left[\varepsilon+\bar{u}\right] - B_{\max}/T \tag{35} \\
&= \mathbb{E}\left[(\varepsilon+\bar{u})^+\right] - D(\bar{u}) - B_{\max}/T. \tag{36}
\end{aligned}
$$

Since the function f is defined by $f(u) = \min(C(u), D(u))$, this concludes the proof of Equations (13) and (12).

It should be clear that $g(u)$ is decreasing in u and $\ell(u)$ is increasing in u. This implies that the function $g(u) \mapsto \ell(u)$ is well-defined and decreasing. Moreover, by definition of ℓ and g, we have $\ell(u) = \mathbb{E}\left[(\varepsilon+u)^+\right] - f(u) = \mathbb{E}\left[\varepsilon\right] + u + g(u)$. Thus, the derivative of $\ell(u)$ w.r.t $g(u)$ equals

$$
\frac{d\ell}{dg} = \frac{d\ell/du}{dg/du} = 1 + \frac{1}{dg(u)/du}.
\tag{37}
$$

Since g is convex (it is the maximum of two convex functions), the function $u \mapsto dg/du$ is increasing. As g is a decreasing function of u, the dg/du decreases with g. This implies that (37) increases with g, which shows that the parametric curve $u \mapsto (g(u), \ell(u))$ is convex. ∎

APPENDIX B
PROOF OF THEOREM 2

We first start by a technical lemma.

We say that a sequence of variable $(X(t))_{t=0,1,\ldots}$ is uniformly integrable if for every $\nu > 0$, there exists $R > 0$ such that $\mathbb{E}\left[|X(t)|\, \mathbf{1}_{|X(t)>R|}\right] \leq \nu$. In particular, this is true if the variables $(X(t))_t$ are identically distributed.

Lemma 1: $(X(t))_t$ be a uniformly integrable sequence of random variables adapted to a filtration $\mathcal{F}_t$. Assume that $\mathbb{E}\left[X(t)|\mathcal{F}_{t-n}\right] = 0$. Then, for all $\alpha, \delta > 0$, there exists $T_0 > 0$ such that for all $T \geq T_0$, we have:

$$
\mathbb{P}\left(\sup_{\tau < T}\left|\sum_{t=1}^\tau X(t)\right| > \alpha T\right) \leq \delta.
\tag{38}
$$

Proof: For all $1 \leq k \leq n$ and $t \geq 0$, we define the variable $Y^k(t) := X(nt+k)$ and we define $\mathcal{G}_t^k := \mathcal{F}_{nt+k}$. It should be clear that $(Y^k(t))_{t \geq 0}$ is a $\mathcal{G}_t^k$-adapted process. Moreover, the sequence of variable $(Y^k(t))_{t \geq 0}$ is uniformly integrable and satisfies $\mathbb{E}\left[Y^k(t+1)|\mathcal{G}_t^l\right] = 0$. Therefore, by Lemma 16 of [33], there exists $T^k > 0$ such that for all $T \geq T^k$:

$$
\mathbb{P}\left(\sup_{\tau < T}\left|\sum_{t=1}^\tau Y^k(t)\right| > \alpha T\right) \leq \frac{\delta}{k}.
\tag{39}
$$

Let $T_0 := n\max_{1 \leq k \leq n} T^k$. Applying the union bound, we have that for all $T \geq T_0$,

$$
\mathbb{P}\left(\sup_{\tau < T}\left|\sum_{t=1}^\tau X(t)\right| > \alpha T\right)
\tag{40}
$$

$$
\leq \sum_{k=1}^n \mathbb{P}\left(\sup_{\tau < \lceil T/n\rceil}\left|\sum_{t=1}^\tau Y^k(t)\right| > \alpha\frac{T}{n}\right) \leq \delta.
\tag{41}
$$
∎

We are now ready to prove the main theorem.

Proof: Let us now consider that the dispatcher applies the fixed offset policy u and let us first assume that $D(u) > C(u)$.

The only neglected terms in Equation (24) are the losses due to overflows (*i.e.* the storage is full and cannot be charged):

$$
L(t) = M^-(t) - \eta\min(M^-(t), C_{\max}) + \text{overflow}(t).
\tag{42}
$$

The mismatch at time t is $M(t) = -(e(t) + u)$, where $e(t)$ denotes the forecast error at time t. Let $c(t) = \eta\min((e(t) + u)^+, C_{\max})$ be the quantity that is charged at time t.

An overflow occurs at time $t + 1$ if $c(t)$ is greater than $B_{\max} - B(t)$. Let us denote $O(t+1) = (c(t) + B(t) - B_{\max})^+$ the quantity of overflow at time $t + 1$. Then, for any $b > 0$,

$$
\begin{aligned}
\mathbb{E}\left[O(t+1)\right] &= \mathbb{E}\left[O(t+1)\mathbf{1}_{B(t) \leq B_{\max}-b}\right] \tag{43} \\
&\quad + \mathbb{E}\left[O(t+1)\mathbf{1}_{B(t) > B_{\max}-b}\right] \tag{44}
\end{aligned}
$$

Let $\delta > 0$. The first term, Equation (43), is equal to

$$
\begin{aligned}
\mathbb{E}&\left[(c(t) + B(t) - B_{\max})^+ \mathbf{1}_{\{B(t) \leq B_{\max}-b\}}\right] \tag{45} \\
&\leq \mathbb{E}\left[(c(t) - b)^+ \mathbf{1}_{\{B(t) \leq B_{\max}-b\}}\right] \leq \mathbb{E}\left[(c(t) - b)^+\right] \tag{46}
\end{aligned}
$$

Since the errors are identically distributed, there exists b_0 such that if $b \geq b_0$, then Equation (43) is less than $\delta/2$.

As for the second term, (44), we have

$$\mathbb{E}\left[O(t+1)\mathbf{1}_{B(t)>B_{\max}-b}\right] \le \mathbb{E}\left[c(t)\right]\mathbb{P}\left(B(t)>B_{\max}-b\right). \tag{47}$$

Let us define $X(t)$ by:

$$X(t) = \eta\min((e^f_{t-n}(t)+u)^+, C_{\max}) - C(u) \tag{48}$$
$$- \min((e^f_{t-n}(t)+u)^-, D_{\max}) + D(u). \tag{49}$$

Let $\alpha = D(u) - C(u)$. The storage $B(t)$ satisfies

$$B(t+1) = (B(t)+X(t)-\alpha)_0^{B_{\max}}. \tag{50}$$

By the assumptions on the noise $e(t)$ and the definition of C and D, $X(t)$ satisfies the assumption of Lemma 1. Let $\nu = \delta/(2\mathbb{E}\left[c(t)\right])$ (this is independent on t as we use the fixed offset policy and the noise is identically distributed). By Lemma 1, there exists T_0 such that if $T \ge T_0$:

$$\mathbb{P}\left(\sup_{k\le T}\left|\sum_{t=1}^{k}X(t)\right| > \frac{\alpha}{4}T\right) \le \nu. \tag{51}$$

Let $B = T_0\alpha$ and assume that $B_{\max} \ge \min(B, 2b_0)$. Let $T = B_{\max}/\alpha$. For all $t \ge T$ the probability $\mathbb{P}(B(t) \ge B_{\max}/2)$ satisfies

- If $B(s) \ne 0$ for all $t \le s < T$, then

$$B(T) \le B(0) - \alpha T + \sum_{t=1}^{T}X(t) \le \sum_{t=1}^{T}X(t) \tag{52}$$

- Otherwise, let t' be such that $B(t') = 0$ and $B(s) \ne 0$ for all s such that $t' \le s < T$. We have

$$B(T) \le \sum_{s=t'+1}^{T}X(s) = \sum_{s=1}^{T}X(s) - \sum_{s=1}^{t'}X(s) \tag{53}$$

$$\le \left|\sum_{s=1}^{T}X(s)\right| + \left|\sum_{s=1}^{t'}X(s)\right| \tag{54}$$

This shows that for $t \ge T$,

$$B(T) \le 2\sum_{k<T}\left|\sum_{t=1}^{k}X(t)\right|. \tag{55}$$

By Equation (51), this is lower than $B_{\max}/2$ with probability ν. Let $B_0 := \max(B, 2b_0)$. Equations (43-44) imply that for $B \ge B_0$, $\mathbb{E}\left[O(t+1)\right] \le \delta$ for all $t \ge T+1$.

This shows that for all $B^{\max} \ge B_0$, we have:

$$\limsup_{k\to\infty}\frac{1}{k}\sum_{t=1}^{k}\mathbb{E}\left[O(t)\right] \le \delta, \tag{56}$$

which by definition of the losses, Equation (42), implies that

$$\bar{L} = \limsup_{k\to\infty}\frac{1}{k}\sum_{t=1}^{k}\mathbb{E}\left[L(t)\right] \le \mathbb{E}\left[(\varepsilon+u)^+\right] - C(u) + \delta. \tag{57}$$

This proof applies mutatis mutandis to the case $D(u) < C(u)$. In that case, this leads to $\bar{G} \le \mathbb{E}\left[(\varepsilon+u)^+\right] - D(u) + \delta$.

Using Equation (32), this implies that for $D(u) \ne C(u)$, $\bar{L} \le \ell(u) + \delta$ and $\bar{G} \le g(u) + \delta$. Since $\bar{L}$ is an increasing function of u (and $\bar{G}$ a decreasing function of u) and $\ell(u)$ and $g(u)$ are continuous in u, this also holds for $D(u) = C(u)$. ∎

REFERENCES

[1] N. Gast, D. Tomozei, and J. Le Boudec, "Optimal storage policies with wind forecast uncertainties," *Greemetrics 2012*, 2012.

[2] A. Bejan, R. Gibbens, and F. Kelly, "Statistical aspects of storage systems modelling in energy networks," *46th Annual Conference on Information Sciences and Systems*, March 2012.

[3] A. Costa, A. Crespo, J. Navarro, G. Lizcano, H. Madsen, and E. Feitosa, "A review on the young history of the wind power short-term prediction," *Renewable and Sustainable Energy Reviews*, vol. 12, no. 6, pp. 1725–1744, 2008.

[4] H. Holttinen, P. Meibom, A. Orths, B. Lange, M. O'Malley, J. O. Tande, A. Estanqueiro, E. Gomez, L. Söder, G. Strbac, *et al.*, "Impacts of large amounts of wind power on design and operation of power systems, results of iea collaboration," *Wind Energy*, vol. 14, no. 2, pp. 179–192, 2011.

[5] K. Heussen, S. Koch, A. Ulbig, and G. Andersson, "Unified system-level modeling of intermittent renewable energy sources and energy storage for power system operation," *Systems Journal, IEEE*, vol. 6, pp. 140 –151, march 2012.

[6] M. Arnold and G. Andersson, "Model predictive control of energy storage including uncertain forecasts," in *Power Systems Computation Conference (PSCC), Stockholm, Sweden*, 2011.

[7] A. Ulbig and G. Andersson, "On operational flexibility in power systems," in *IEEE PES General Meeting, San Diego, USA*, 2012.

[8] R. Doherty and M. O'Malley, "A new approach to quantify reserve demand in systems with significant installed wind capacity," *Power Systems, IEEE Transactions on*, vol. 20, no. 2, pp. 587–595, 2005.

[9] M. A. Ortega-Vazquez and D. S. Kirschen, "Estimating the spinning reserve requirements in systems with significant wind power generation penetration," *Power Systems, IEEE Transactions on*, vol. 24, no. 1, pp. 114–124, 2009.

[10] J. Warrington, P. Goulart, S. Mariéthoz, and M. Morari, "Robust reserve operation in power systems using affine policies," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 1111–1117, IEEE, 2012.

[11] A. Fabbri, T. Gomez San Roman, J. Rivier Abbad, and V. Mendez Quezada, "Assessment of the cost associated with wind generation prediction errors in a liberalized electricity market," *Power Systems, IEEE Transactions on*, vol. 20, no. 3, pp. 1440–1446, 2005.

[12] E. M. Constantinescu, V. M. Zavala, M. Rocklin, S. Lee, and M. Anitescu, "A computational framework for uncertainty quantification and stochastic optimization in unit commitment with wind power generation," *Power Systems, IEEE Transactions on*, vol. 26, no. 1, pp. 431–441, 2011.

[13] T. K. Brekken, A. Yokochi, A. von Jouanne, Z. Z. Yen, H. M. Hapke, and D. A. Halamay, "Optimal energy storage sizing and control for wind power applications," *Sustainable Energy, IEEE Transactions on*, vol. 2, no. 1, pp. 69–77, 2011.

[14] J. Garcia-Gonzalez, R. de la Muela, L. Santos, and A. González, "Stochastic joint optimization of wind generation and pumped-storage units in an electricity market," *IEEE Transactions on Power Systems*, vol. 23, no. 2, pp. 460–468, 2008.

[15] E. D. Castronuovo and J. A. P. Lopes, "Optimal operation and hydro storage sizing of a wind-hydro power plant," *International Journal of Electrical Power & Energy Systems*, vol. 26, no. 10, pp. 771 – 778, 2004.

[16] M. Korpaas, A. T. Holen, and R. Hildrum, "Operation and sizing of energy storage for wind power plants in a market system," *International Journal of Electrical Power & Energy Systems*, vol. 25, no. 8, pp. 599 – 606, 2003.

[17] E. Y. Bitar, E. Baeyens, P. P. Khargonekar, K. Poolla, and P. Varaiya, "Optimal sharing of quantity risk for a coalition of wind power producers facing nodal prices," in *American Control Conference (ACC), 2012*, pp. 4438–4445, IEEE, 2012.

[18] S. Chatzivasileiadis, M. Bucher, M. Arnold, T. Krause, and G. Andersson, "Incentives for optimal integration of fluctuating power generation," in *17th Power Systems Computation Conference*, 2011.

[19] N. Gast, J. Le Boudec, A. Proutière, and D. Tomozei, "Impact of storage on the efficiency and prices in real-time electricity markets," in *ACM E-Energy 2013*, 2013.

[20] R. Sioshansi, "Welfare impacts of electricity storage and the implications of ownership structure," *Energy Journal*, vol. 31, no. 2, p. 173, 2010.

[21] X. HE, W. Delarue, E. abd D'Haeseleer, and J.-M. Glachant, "Coupling electricity storage with electricity markets: a welfare analysis in the french market," *TME working paper - Energy and Environment*, 2012.

[22] R. Walawalkar, J. Apt, and R. Mancini, "Economics of electric energy storage for energy arbitrage and regulation in new york," *Energy Policy*, vol. 35, no. 4, pp. 2558–2568, 2007.

[23] R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss, "Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects," *Energy Economics*, vol. 31, no. 2, pp. 269–277, 2009.

[24] A. Tuohy and M. O'Malley, "Impact of pumped storage on power systems with increasing wind penetration," in *Power & Energy Society General Meeting.*, pp. 1–8, IEEE, 2009.

[25] M. Hildmann, A. Ulbig, and G. Andersson, "Electricity grid in-feed from renewable sources: A risk for pumped-storage hydro plants?," in *Energy Market (EEM), 2011 8th International Conference on the European*, pp. 185 –190, may 2011.

[26] D. MacKay, *Sustainable Energy-without the hot air*. UIT Cambridge, 2008.

[27] "UK DECC pathway calculator," *http://2050-calculator-tool.decc.gov. uk*, accessed 26/06/2013.

[28] https://www.dom.com/about/stations/hydro/ bath-county-pumped-storage-station.jsp, accessed 26/06/2013.

[29] http://www.fhc.co.uk/dinorwig.htm, accessed 26/06/2013.

[30] B. Hodge and M. Milligan, "Wind power forecasting error distributions over multiple timescales," in *Power and Energy Society General Meeting, 2011 IEEE*, pp. 1–8, IEEE, 2011.

[31] H. Bludszuweit, J. A. Domínguez-Navarro, and A. Llombart, "Statistical analysis of wind power forecast error," *Power Systems, IEEE Transactions on*, vol. 23, no. 3, pp. 983–991, 2008.

[32] M. Puterman, *Markov decision processes: Discrete stochastic dynamic programming*. J. W. & S., 1994.

[33] N. Gast and B. Gaujal, "Markov chains with discontinuous drifts have differential inclusion limits," *Performance Evaluation*, vol. 69, no. 12, pp. 623–642, 2012.