



PDE methods for stochastic dynamic optimisation

DOI:

[10.1098/rsta.2016.0301](https://doi.org/10.1098/rsta.2016.0301)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Johnson, P., Howell, S., & Duck, P. (2017). PDE methods for stochastic dynamic optimisation: an application to wind power generation with energy storage. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375(2100), Article 20160301. <https://doi.org/10.1098/rsta.2016.0301>

Published in:

Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.





Article submitted to journal

Subject Areas:

Math Finance,

Keywords:

Wind Energy, Investment, Battery
Storage, Arbitrage, Dynamic Optimal
Control

Author for correspondence:

Paul Johnson

e-mail:

paul.johnson-2@manchester.ac.uk

PDE methods for stochastic dynamic optimisation: an application to wind power generation with energy storage

Paul Johnson¹, Sydney Howell² and Peter
Duck¹

¹School of Mathematics, The University of Manchester

²Alliance Manchester Business School, The University
of Manchester

A mixed financial/physical Partial Differential Equation (PDE) can optimise the joint earnings of a single wind power generator (WPG) and a generic energy storage device (ESD). Physically, the PDE includes constraints on the ESD's capacity, efficiency and maximum speeds of charge and discharge. There is a mean-reverting daily stochastic cycle for WPG power output. Physically, energy can only be produced or delivered at finite rates. All suppliers must commit hourly to a finite rate of delivery C , which is a continuous control variable that is changed hourly. Financially, we assume heavy 'system balancing' penalties in continuous time, for deviations of output rate from the commitment C . Also, the electricity spot price follows a mean-reverting stochastic cycle with a strong evening peak, when system balancing penalties also peak. Hence the economic goal of the WPG plus ESD, at each decision point, is to maximise expected Net Present Value (NPV) of all earnings (arbitrage) minus the NPV of all expected system balancing penalties, along all financially/physically feasible future paths through state space. Given the capital costs for the various combinations of the physical parameters, the design and operating rules for a WPG plus ESD in a finite market may be jointly optimisable.

1. Introduction

The large-scale use of wind power will pose unprecedented problems, because the supply wattage of a Wind Power Generator (WPG) over a 24-hour lead time is approximately a mean reverting Geometrical Brownian Motion (GBM). Such processes are noisy at any lead time, inside and outside the range of seconds to 10^5 seconds, and that noise can severely disrupt both system balancing and price arbitrage. As kindly pointed out by an anonymous referee, the average contribution of wind in Ireland has already reached 24% in 2015 and in Germany the contribution peaked at 36% in May 2016, both these (and other) countries therefore face a tremendously difficult task managing large fluctuations in power. One possible way to make it run more efficiently would be the introduction of large-scale energy storage on the grid. The optimum use of an Energy Storage Device (ESD) must hourly reset a committed rate of output, such as to allocate the ESD's resources dynamically between the two conflicting priorities of system balancing, which minimises balancing penalties, and price arbitrage, which maximises income.

Mathematical models that can tackle variable flow rates into and out of a store have emerged in the literature over the last 10 years, with applications appearing in gas storage [1] as well as thermal storage [2]. However many 'real options' models of electrical storage assume an infinite or a fixed rate of physical flow into or out of the ESD [3]. This infinite flow rate imitates financial trading in an inertia-free market. Such models set and solve an optimal stopping problem, of when to empty or fill the ESD instantaneously, when the selling price hits a suitable trigger level. Instantaneous emptying is not dynamically possible in a physical power system, but if ESDs were emptied at the largest physically feasible rate, this would create, rather than solve, system balancing problems for the physical systems.

The present paper addresses the above problem in three innovative ways, by using a 'proof of concept' mathematical model. Its main features are:

- (i) The physical state variables include both a finite and instantaneously variable rate of in/out flow, and its integral. Using both variables makes the model physically realisable, and allows it to balance two competing goals, namely system balancing (avoiding real-time penalties, by holding the rate of supply close to some committed rate C) and arbitrage (which maximises income from selling the maximum of energy at the daily peak of selling price – when balancing penalties also peak). The model assumes a market commitment process in which the joint owner of a WPG and an ESD can change C at hourly intervals, and there are continuous penalties in real time for deviating from the presently committed C . Hence C is a continuous control variable that is reset at discrete times. The time path of C has delayed and highly stochastic effects on future system balancing penalties and sales income.
- (ii) To imitate a future power system, having high wind penetration, the PDE model imposes the full economic cost of any system balancing needed for wind power supplied. This is done by a penalty function, which penalises all divergences between supply delivered and commitment C , continuously in real time. The penalty function can be adjusted to model the actual opportunity cost to the grid of 'buying in' system balancing from any source, which in turn sets the value to WPG/ESD of doing some of this balancing itself. In total the PDE interacts with the market economy through three pricing functions: the cost of capital, the contracted sale value of committed electricity and the system balancing penalty.
- (iii) The model's PDE includes stochastic dynamics for the wind supply and a stochastic electricity selling price (both have daily cycles) and it assumes simple deterministic dynamics for the ESD that are straightforward to generalise.

In the PDE, the stochastic, deterministic, financial and physical variables all interact instantaneously, and in the diffusion solution all these variables, including the hourly changeable commitment decision C , also interact through space and time – an example of (joint) Stochastic

Dynamic Optimisation (SDO). The hourly Bellman optimisation of C maximises the expected financial outcome along the infinity of dynamically realisable, continuous trajectories passing through every point in the model's state space. As later results show, the Bellman solutions make unexpectedly subtle trade-offs between balancing penalties and selling price arbitrage, and the trade-offs vary strongly across the state space.

The structure of the paper is as follows: Section 2: A brief introduction to the problem of stochastic storage for wind energy; Section 3: The dynamics of the ESD are discussed; Section 4: A 'proof of concept' SDO PDE is developed; Section 5: Results; Section 6: Conclusions and prospects for future research.

2. Dynamics of Wind and Electricity Price

The static distribution of wind power, relevant to long-term investments in WPG, is often assumed to be best approximated by a Weibull distribution [4,5]. However this lacks a time-dependent behaviour, and it is not tractable over the time scales that are operationally important, which range from seconds (for system balancing) to multi-hour (for arbitrage). Helpfully [6] report finding a Gaussian time series beneath the daily seasonal patterns in wind data. As [7] notes, there are long-term seasonal variations over months (winter versus summer) and in some situations a weak diurnal cycle (e.g. stronger wind at night along coasts). There is also some ability to forecast the local wind strength rather approximately, within any 24 hour period. This lead time is relevant for balancing within-day arbitrage gains against system-balancing penalties. For 'proof of mathematical concept' purposes we model wind strength by a Stochastic Differential Equation (SDE), in which GBM mean-reverts towards a known weak diurnal cycle. The assumed SDE for wind strength X is

$$dX = \kappa_X \left(\theta_X(t) + \frac{1}{\kappa_X} \frac{d\theta_X}{dt} - X \right) dt + \sigma_X X dW. \quad (2.1)$$

Here κ_X is the constant local speed of mean reversion of X towards its local mean $\theta_X(t)$ while σ_X is a constant volatility and dW is the increment of a Wiener variable. The parameter values assumed here are $\kappa_X = 0.1 \text{ h}^{-1}$, $\sigma_X = 0.2 \text{ h}^{-\frac{1}{2}}$ and the daily seasonal forcing function for wind is

$$\theta_X(t) = \bar{\theta}_X (1 + \alpha_X \sin(\gamma(t + \psi_X))), \quad (2.2)$$

where $\bar{\theta}_X = 8 \text{ m s}^{-1}$, $\alpha_X = 0.375$, $\psi_X = 2 \text{ h}$ and $\gamma = \frac{\pi}{12} \text{ h}^{-1}$. A long-term static distribution of wind strength results from this SDE, which can be found by time averaging over the solution to the associated Fokker-Planck equation. If we analyse the properties of the solution to the Fokker-Planck equation, we can derive the probability that $X < 0$ under natural conditions, and we are able to artificially impose non-negativity if we choose simply by using a reflecting (or no-flux) boundary condition at $X = 0$ [8]. When solving the PDE later on we apply reflecting boundary conditions at $X = 0$ that are outlined in section 4(c), meaning that even if negative values were possible we bound our domain by $X = 0$. This distribution is plotted as the thick solid line in figure 1(a), and the general similarity of this static distribution to the Weibull static distribution can be seen clearly. In figure 1(b), the thick solid line represents the median of the distribution, the three lines below it represent the 1%, 5% and 25% quantiles, and the three lines above represent the 75%, 95% and 99% quantiles. The two distributions differ but as noted above, the Weibull does not offer useful time dynamics for system balancing purposes, and as seen later, these different distributions of wind strength lead to rather similar distributions of electrical power output.

The main driver of arbitrage opportunities is the strong but also stochastic daily cycle of electricity price. We take both the wind strength X and the electricity selling price Y to be different GBMs over the following 24 hours, and we assume that each mean-reverts towards its own daily cycle (strong for price and weak for wind), which is identical every day. In this analysis we ignore the longer term cycles of wind across the seasons in order to not over complicate matters. If a more detailed analysis were to be carried out, such effects could be taken into account most readily by

solving the problem (in isolation) for representative seasonal days, and then taking the average value across them. Using a single day cycle allows us to quickly converge to a periodic steady solution when solving the problem numerically.

We convert the wind strength X to a power output, through a smoothed non-linear function $F(X)$. We do not model changes in wind direction, because this variable is not present in most published data on WPG power output. Hence changes in wind direction, and possible engineering responses to these in WPG design, are not needed in a ‘proof of concept’ model of the optimal storage of stochastically arriving wind power. The output is scaled to a representative maximum power output of 1 MW (easily rescaled to model any desired wind power generation system or WPG), with a cut-in speed of approximately 4 m s^{-1} , a rated wind speed of approximately 13 m s^{-1} , and a cut-out speed of approximately 25 m s^{-1} [9]; we show this power curve in figure 2(a). Given the assumed daily trend and stochastics of wind described in (2.1), we can calculate the expected capacity factor at each time of the day, as shown in 2(b).

Comparing these assumptions to the empirical analysis of UK wind farms in [7], the model’s assumptions give a peak output of around 50% of maximum capacity, which is typical of the empirical maximum seen in UK winter, and a trough output of 10% of maximum capacity, which resembles the empirical minimum in summer. The model’s grand mean output of 29.5% also resembles UK empirical values [9]. We note that the existence of cut-in and cut-out speeds reduces the difference between the GBM and the Weibull, as static distributions of actual power supply rates, in that they differ most strongly at wind speeds which are too low or too high for generation.

We model the instantaneous market selling price of electricity as Y by the following (mean-reverting) SDE

$$dY = \kappa_Y \left(\theta_Y(t) + \frac{1}{\kappa_Y} \frac{d\theta_Y}{dt} - Y \right) dt + \sigma_Y Y dZ, \quad (2.3)$$

where κ_Y is the constant local speed of mean reversion of Y towards its local mean $\theta_Y(t)$, and σ_Y is its volatility. Here dZ is the increment of a Wiener variable, uncorrelated with dW . It should be noted that the linear mean-reversion term means that this is not a standard GBM model. Models similar to this have already been used many times in finance, for example [10] fit the standard GBM mean reverting model to data from the nordic power exchange, whereas [11] use a variation on our SDE with linear mean-reversion to model interest rates. It is important to note again that negative prices are possible in the model since the mean-reversion coefficient could go negative. For our ‘proof of concept’ model we choose to take an arbitrary set of parameters that demonstrate some properties of the UK market, such as an average price of around 40 £/MWh and the daily cycle with a peak in the afternoon at 4pm.

$$\kappa_Y = 0.04 \text{ h}^{-1}, \sigma_Y = 0.075 \text{ h}^{-\frac{1}{2}}, \quad (2.4)$$

and the daily cyclical forcing function for Y is

$$\theta_Y(t) = \bar{\theta}_Y (1 + \alpha_Y \sin(\gamma(t + \psi_Y))). \quad (2.5)$$

We choose the parameters as $\bar{\theta}_Y = 40 \text{ £/MWh}$, $\alpha_Y = 0.375$, $\psi_Y = 14 \text{ h}$ and $\gamma = \frac{\pi}{12} \text{ h}^{-1}$. We choose to bound possible values of $Y \geq 0$ using a reflecting boundary condition at $Y = 0$. Usefully, our treatment reduces the size of the computational task and simplifies the optimal strategy.

3. Dynamics of the Energy Storage Device

The instantaneous energy content of the ESD is Q , where $0 < Q < Q_{\max}$ and Q_{\max} is the designed capacity of the ESD. We define the charge or discharge rate from the ESD to be

$$\mathcal{L}(X, Q) = \frac{dQ}{dt}, \quad (3.1)$$

allowing \mathcal{L} to be both positive and negative. At any moment, the total system of the WPG plus ESD has a committed output rate of C . In the ideal case the total system’s output equals C . However if $Q = Q_{\max}$ or $Q = 0$, charge and discharge respectively are impossible. When Q is

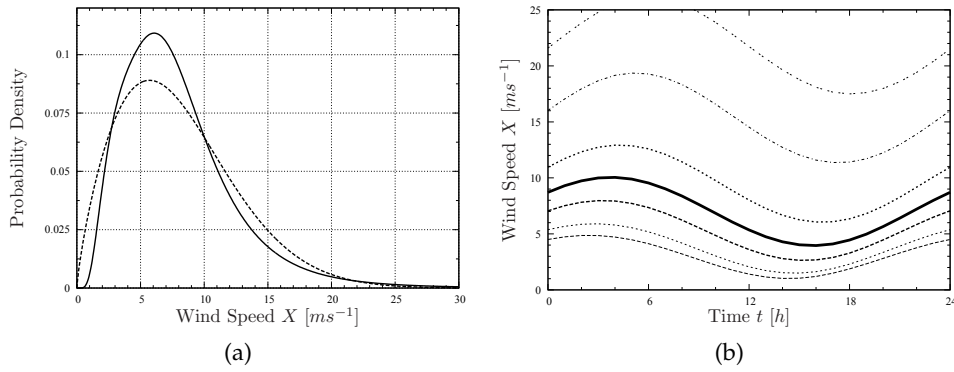


Figure 1. Some summary statistics of the wind speed model used in this paper. On the left in (a) we present the static probability distribution resulting from SDE (2.1) with a solid line, a Weibull distribution with an identical mean and variance is plotted with a dashed line. On the right in (b) we show the quantile plots at different times of the day for SDE (2.1). The solid curve at the centre is the 50%, line, with the three lines below representing the 1%, 5% and 25% quantiles, whilst the three lines above represent the 75%, 95% and 99% quantiles.




Figure 2. On the left in (a) we present a typical wind speed to power output function, as used in this paper. On the right we apply this function to the mean daily cycle of wind strength and plot the mean fluctuation of WPG output as a percentage of maximum output (capacity factor)

not at one of these bounds, the rates of charge and discharge are limited by the design of the ESD to a maximum charge rate L_C and a maximum discharge rate L_D . Richer ESD dynamics can be assumed, and in particular we limit our investigation to cases where L_C and L_D are both equal even though there is no such physical restriction on ESDs. We assume here that there is a fixed percentage loss rate $\kappa = \sqrt{0.7}$, such that a round trip conversion into and out of the ESD will only return 70% of the initial energy. This is actually quite low for an ESD, with many ESD providers claiming up to 95% efficiency. However these figures should be offset by the potential degradation of the ESD through charge/discharge cycles, and since our model cannot currently take account of this we set our efficiency at the lower end. Mathematically the above ESD is described by two scenarios. If the current WPG power output $F(X) > C$ there is a surplus, so the ESD will if charge at up to its maximum rate, hence

$$\mathcal{L}(X, Q) = \min(\kappa(F(X) - C), L_C, \xi_C(Q_{\max} - Q)) \quad (3.2)$$

to ameliorate the discrepancy. If the WPG power output $F(X) < C$ is in deficit, then the ESD will discharge at up to its maximum rate (after efficiency losses), so that

$$\mathcal{L}(X, Q) = \max \left(\frac{1}{\kappa} (F(X) - C), -L_D, -\xi_D Q \right). \quad (3.3)$$

These smoothing parameters, ξ_C and ξ_D , will limit the charging/discharging rate near full/empty to numerically smooth the solutions near the boundary, and they are chosen to be sufficiently large so that increasing them has little effect on the solution. This simplified model of an ESD has been previously used in [12].

Two alternative sets of parameter values are used to model two alternative designs of ESD. We define a ‘small’ ESD with $L_C = L_D = 250$ kW, $Q_{\max} = 1$ MWh and we define a ‘large’ ESD with $L_C = L_D = 1$ MW, $Q_{\max} = 4$ MWh. The ‘small’ ESD can charge or discharge at only a quarter of the WPG’s peak output rate, and it can store only one hour of the WPG’s peak output rate. The small ESD is expected to ‘smooth’ the WPG’s output over an hour or two, but not to ‘arbitrage’ the selling price of the WPG’s output, because this ESD can time shift at most one hour of maximum wind power output towards the evening peak. In contrast the large ESD can charge or discharge at the maximum output rate of the WPG, and it can also move four hours of peak power output away from the early morning peak of physical supply towards the early evening peak, when physical demand and selling price both peak. Hence the large ESD is expected both to smooth the WPG’s power output and to time-shift it.

4. Modelling continuous-time penalties for over- and under-delivery

The smoothing of wind power is becoming an increasingly costly problem for networks around the world. At present there is a low penetration of such power in most of the world, and therefore although each individual WPG produces very ‘noisy’ power, they are too few in aggregate to disturb the overall supply/demand balance – especially in richly connected regions like Germany. Accordingly, the scale of system balancing charges presently in force was historically designed to penalise occasional over- or under-supply by fossil generators (which are fairly stable over short periods). This scheme is often applied unchanged to WPGs, so that a WPG can deliver noisy power throughout an hour without penalty, provided that the cumulative total delivered matches the total committed.

At high levels of penetration, fluctuations of wind power have the serious potential to disrupt the total supply-demand balance on a historically unprecedented scale. Systems will then incur much higher balancing costs, and they are likely to pass these on as penalties to WPGs. This in turn should motivate WPGs towards some mix of under-committing their expected output (in the limit, discarding some power), and storing energy for later smoothing.

This paper’s model assumes a system of balancing charges in a generic form, which can be set sufficiently harshly (as a function of $F(X) - C$) to motivate WPGs to undertake some mix of under-commitment (of expected supply) and storage (of unexpected and/or uncommitted supply). The assumed system balancing charge, penalises both over- or under-supply in real time, so that the only penalty-free outcome is to deliver at exactly the committed wattage rate C . This earns the contracted selling price Y , a stochastic value which is fixed at the moment when C is fixed. Hence the delivery to the grid $D(X, Q, C)$ at any given time is

$$D(X, Q, C) = \begin{cases} F(X) - \frac{1}{\kappa} \mathcal{L} & \text{if } F(X) > C, \quad \mathcal{L} \geq 0 \\ F(X) - \kappa \mathcal{L} & \text{if } F(X) < C, \quad \mathcal{L} \leq 0, \end{cases} \quad (4.1)$$

where the efficiency conversion means that only the quantity $\kappa\mathcal{L}$ is provided from the ESD to the grid. Therefore the instantaneous income function \mathcal{I} is defined by

$$\mathcal{I}(X, Y, Q, C) = \begin{cases} CY + Y(1 - \zeta)(D(X, Q, C) - C) & \text{if } F(X) > C \\ CY & \text{if } F(X) = C \\ CY - Y(1 + \zeta)(C - D(X, Q, C)) & \text{if } F(X) < C \end{cases} \quad (4.2)$$

Here the parameter $0 < \zeta < 1$ defines the rate of price penalty for under- or over-delivery during dt . The rate of extra income earned on over-delivered power, at an instantaneous output rate $D(X, Q, C) - C > 0$, is $Y(1 - \zeta) < Y$. Here the committed wattage rate C has been previously set to the system, therefore the opportunity cost of instantaneous over-delivery (e.g. after under-committing, when unable to store a surplus) is to lose the fraction ζ of the selling price of the wattage over-delivered. Since in our model we choose to bound $Y \geq 0$ and $\zeta < 1$ both hold for all t , the WPG has no incentive to curtail generation. It is often the case in real markets that either $Y < 0$ or $\zeta \geq 1$ causing excess power to be shed from the system. This situation could be included in future models by adjusting parameters and the relevant boundary conditions.

The corresponding penalty for instantaneous under-delivery (in addition to being unable to sell all of the wattage in the commitment C) is a cash ‘fine’ of the fraction ζ of the full spot sales value of whatever wattage fails to be delivered $C - D(X, Q, C)$. A large enough under-delivery causes a negative instantaneous income from the wattage actually delivered.

(a) Derivation of the PDE for system value

Now define the contract $V(X, Y, Q, t)$ in units [£] as the current value of the jointly owned WPG and ESD. This is the expectation under the real-world measure of all future discounted income, which the WPG plus ESD can make on sales in the electricity supply market. Ignoring the optimal commitment for the time being, we can write

$$V(X, Y, Q, t; C) = E \left[\int_t^{\infty} e^{-r(s-t)} \mathcal{I}(X_s, Y_s, Q_s, C) ds \right] \quad (4.3)$$

where $0 \leq t \leq \infty$ and $C = C(X, Y, Q, s)$ is the committed delivery amount at time s which may be a function of all variables in the problem. We define r as a general discount rate, which can be the given risk adjusted or risk free value, as the nature of the problem requires. Changes to r will change the capital value of the WPG/ESD system over its lifetime, but should not change the optimal commitment policy C^* within any one day. The effects of discounting at a typical real rate of return (seldom outside the range 1% to 10% per annum) are far smaller over 24 hours than the expected fluctuation of the real electricity price within that period.

Now let us define

$$V^*(X, Y, Q, t) = \max_C V(X, Y, Q, t; C) \quad (4.4)$$

as the solution to the optimisation problem at time t . Then to introduce our optimal commitment problem, we first assume that electricity supply sales (which commits us to a fixed rate of supply C for a fixed time interval ΔT) happen at discrete time points $t = 0, \Delta T, 2\Delta T, \dots$ which are to be indexed by k . More precisely, at time $t = t_k = k\Delta T$ the owner must choose to commit to deliver at a rate of $C \in [0, C_{\max}]$ units of electricity supply, starting at time t_k and ending at $t_{k+1} = t_k + \Delta T$. Then the value of this contract given a commitment C is

$$V(X, Y, Q, t; C, t_{k+1}) = E \left[\int_t^{t_{k+1}} e^{-r(s-t)} \mathcal{I}(X_s, Y_s, Q_s, C) ds + e^{-r(t_{k+1}-t)} V^*(X_{t_{k+1}}, Y_{t_{k+1}}, Q_{t_{k+1}}, t_{k+1}) \right] \quad (4.5)$$

for $t_k < t < t_{k+1}$; here V^* assumes that all future decisions are made optimally. We can simplify the notation at this stage by dropping variables that do not play an explicit part in the

optimisation, namely X , Y and Q , so that the quantities defined in (4.3), (4.4), and (4.5) become $V(t; C)$, $V^*(t)$ and $V(t; C, t_{k+1})$ respectively.

If we assume that the possible set of commitments is a finite set of values $\{C_i\}$ where $C_i \in [0, C_{\max}]$, then at the decision time t_k , we must choose the optimal C_i . Once the decision is made we must hold that choice of C_i constant over the period $t_k < t < t_{k+1}$, and the value of making that choice we will write as $V_i(t; t_{k+1}) = V(t; C_i, t_{k+1})$. Now the optimisation becomes relatively simple and can be written as

$$V^*(t = t_k) = \max_i [V_i(t_k; t_{k+1})]. \quad (4.6)$$

The optimal condition (4.6) can then be transformed into the terminal condition required to solve for the previous period $[t_{k-1}, t_k]$. The appropriate condition at t_k is then

$$V_i(t = t_k^-; t_k) = \max_j [V_j(t_k^+; t_{k+1})] \quad (4.7)$$

for all i . The positive and negative superscripts take account of the fact that our choice of commitment is in the period $[t_{k-1}, t_k]$ or $[t_k, t_{k+1}]$.

Now if the market includes a delay, so that the owner must precommit to a delivery rate starting in the future, then at time t_k the owner must choose from the set $\{C_i\}$ for the period starting at t_{k+1} and ending at t_{k+2} . We approximate the full set of admissible controls by a finite set of pairs, where each pair includes the commitment for the current period as well as the next. So we denote $C_{i,j} = (C_i, C_j)$ as the scenario in which C_i is the current commitment and C_j is the commitment at the next step, and $V_{i,j}$ to denote the value of implementing that control strategy. We have

$$V_{i,j}(t; t_{k+1}) = E \left[\int_t^{t_{k+1}} e^{-r(s-t)} \mathcal{I}(X_s, Y_s, Q_s, C_i) ds + e^{-r(t_{k+1}-t)} V_j^*(t_{k+1}; t_{k+2}) \right]$$

for $t_k < t < t_{k+1}$. Here $V_j^*(t_{k+1}; t_{k+2})$ indicates the value of committing C_j at t_{k+1} followed by the optimal commitments for $t > t_{k+2}$.

In order to optimally manage the operation of the wind farm, at time $t = t_k$ we lock in the commitment at t_{k+1} by maximising over all possible future commitment decisions j for each of the current possible commitments i (for which we are yet to make a decision on). Mathematically we can write the optimisation as

$$V_i^*(t_k; t_{k+1}) = \max_j [V_{i,j}(t_k; t_{k+1})] \quad \forall i. \quad (4.8)$$

Now as we step back in time across the instant from just after t_k to just before t_k , we must take account of the fact that $C_{i,j}$ refers the current and future commitments, and that current one will become the future commitment for the previous period. So, if we write t_k^- to denote the value the instant before the decision is made, and t_k^+ to denote the instant just after, we obtain the final condition when solving in the period $t \in [t_{k-1}, t_k]$

$$V_{i,j}(t = t_k^-; t_k) = \max_l [V_{j,l}(t_k^+; t_{k+1})]. \quad (4.9)$$

To value the contract in between agreement times $t_k < t < t_{k+1}$, with or without a delay before its commencement, we can use the Feynman Kac formula (according to Itô calculus, see [13] for more details) to derive a partial differential equation for $V_{i,j}$. It follows simply that we must solve the following PDE for all possible controls $C_{i,j}$,

$$\begin{aligned} \frac{\partial V_{i,j}}{\partial t} + \frac{1}{2} \sigma_X^2 X^2 \frac{\partial^2 V_{i,j}}{\partial X^2} + \frac{1}{2} \sigma_Y^2 Y^2 \frac{\partial^2 V_{i,j}}{\partial Y^2} + \mathcal{L} \frac{\partial V_{i,j}}{\partial Q} + \kappa_Y \left(\theta_Y + \frac{1}{\kappa_Y} \frac{d\theta_Y}{dt} - Y \right) \frac{\partial V_{i,j}}{\partial Y} \\ + \kappa_X \left(\theta_X + \frac{1}{\kappa_X} \frac{d\theta_X}{dt} - X \right) \frac{\partial V_{i,j}}{\partial X} - r V_{i,j} + \mathcal{I}(X, Y, Q, C_i) = 0, \end{aligned} \quad (4.10)$$

where r is the discount rate (per hour) and \mathcal{I} is the instantaneous income function as defined in (4.2). Note that the commitment that enters this equation is C_i , the current commitment, whilst the future commitments C_j enter through boundary conditions (4.9). When the problem to be solved does not include a time delay before the chosen commitment becomes effective, the j subscript can be omitted and the final condition (4.7) is used instead.

(b) Boundary Conditions

First we deal with the boundary conditions for X and Y , which are both dealt with in the same way. In order to derive boundary conditions at $X = 0$ or $Y = 0$, we simply set the term to zero in the PDE and solve the resulting equations. For example, on $X = 0$ we solve

$$\begin{aligned} \frac{\partial V_{i,j}}{\partial t} + \frac{1}{2}\sigma_Y^2 Y^2 \frac{\partial^2 V_{i,j}}{\partial Y^2} + \mathcal{L} \frac{\partial V_{i,j}}{\partial Q} + \kappa_Y \left(\theta_Y + \frac{1}{\kappa_Y} \frac{d\theta_Y}{dt} - Y \right) \frac{\partial V_{i,j}}{\partial Y} \\ + \kappa_X \theta_X \frac{\partial V_{i,j}}{\partial X} - rV_{i,j} + \mathcal{I}(0, Y, Q, C_i) = 0, \end{aligned}$$

and a similar equation is derived for $Y = 0$.

For the large X and large Y boundaries we assume that the respective second order differential term is small and can be neglected so that we may solve the resulting equation. So for large X the following equation is solved

$$\begin{aligned} \frac{\partial V_{i,j}}{\partial t} + \frac{1}{2}\sigma_Y^2 Y^2 \frac{\partial^2 V_{i,j}}{\partial Y^2} + \mathcal{L} \frac{\partial V_{i,j}}{\partial Q} + \kappa_Y \left(\theta_Y + \frac{1}{\kappa_Y} \frac{d\theta_Y}{dt} - Y \right) \frac{\partial V_{i,j}}{\partial Y} \\ + \kappa_X \left(\theta_X + \frac{1}{\kappa_X} \frac{d\theta_X}{dt} - X \right) \frac{\partial V_{i,j}}{\partial X} - rV_{i,j} + \mathcal{I}(X, Y, Q, C_i) = 0. \end{aligned}$$

The boundary conditions in Y are derived in a similar way.

For the boundaries in Q , note that we have $\mathcal{L} \leq 0$ on $Q = 0$, and $\mathcal{L} \geq 0$ on $Q = Q_{\max}$, meaning that the characteristics of the solution do not travel outside the domain $[0, Q_{\max}]$ (helped in the discretisation scheme by the terms ξ_C and ξ_D) when moving backwards in time. The result is that we are able to solve the full PDE on the boundaries.

We search for a perpetual solution to the problem which in this case does not mean that $\partial/\partial t \rightarrow 0$ because our solution will be periodic. Observe that (2.1) and (2.3) both have 24-hour periodic functions, then we might expect that any steady solution to the problem will also have a 24-hour period. So we proceed by solving successive 24-hour periods, using the solution from the last 24-hour period as the final condition for the next period. We check on the difference between solutions at the end of each period, and once we have a solution that satisfies

$$\|V(X, Y, Q, t) - V(X, Y, Q, t + 24)\| \leq \epsilon$$

for some small parameter ϵ we stop solving. The convergence towards this periodic perpetual solution is laboriously slow, but we have developed some heuristic methods to speed matters up.

(c) Numerical methods

The wind storage problem has the high dimensionality (for a PDE) of four state variables, including time. A further difficulty is introduced by the control set, meaning that the high dimensional PDE must be solved for each and every possible control strategy. In the case of a trading delay, this means that an extra factor of n^2 will be applied to the computation time, where n is the number of possible commitments at each time. Such calculations are time consuming, but when their purpose is to value an investment proposition a solution time of days, or even weeks, may be acceptable.

We apply standard finite-difference techniques, whose convergence and stability properties for this problem are described in [14]. A typical calculation will divide the domain (X, Y, Q, t, C) into a fixed grid comprising of $251 \times 51 \times 51 \times 401 \times 21$ points to solve for each contract period.

We find that using a fine grid will start to take a substantial time, whereas coarser grids with for example $101 \times 21 \times 21 \times 201 \times 11$ points provide comparable results in a much shorter timescale. Most of the results in this paper are calculated on the coarser grid. In terms of the X and Y grid, we choose $X_{\max} = 100$ and $Y_{\max} = 100$. A much smaller number of points in Y and Q can be used since there are no strong gradients in those directions.

5. Results

In this section we define $C^*(t)$ as the commitment that applies at time t and that has in general been set optimally at some time before t . In calculating the option's value V over the four-dimensional space of X, Y, Q, t , and over all possible strategies for commitment C , the diffusion PDE's numerical solution, the value V in \mathcal{E} , is the result of a probability weighted integral over infinitely many potential future trajectories for the continuous variables, between optimisation at the successive discrete fixings of C .

Clearly the valid C^* , a scalar, contains no time indexed information about the possible future time-domain trajectory of any of the variables X, Y, Q , or of V or C (even though the physically and financially feasible trajectories for all these variables have been integrated over whilst calculating the C^* at each state point). The absence of time-domain output can have slightly different effects on WPG investors and on WPG operators.

Investors, in financial theory, are willing to spend any sum $\leq V$ instantaneously to acquire the instantaneous value V of owning an (optimised) WPG plus ESD system; time-domain trajectories for any of these variables do not interest such investors. In contrast WPG operators live in the time domain, and therefore need to understand how the optimised policy will drive the entire dynamic system in the time domain. They need to trust that it can and will remain economically optimised. A wiser (real-world) investor will also always check that the operators can understand, and will commit to achieving, the predicted optimality of any dynamic trajectories in the time domain which result from implementing $C^*(X, Y, Q, t)$ at any point.

To assist intuition for operators and wiser investors, some sample plots in figure 3 of the joint time trajectories of the following subset of optimised system variables: Q, C^*, D, t , where D is the instantaneous rate of power delivery from the WPG and ESD. Each plot was generated from a single 24 hour realisation of the joint stochastic paths of Q_t, X_t, Y_t and of the resulting $C^*(t)$ as computed from the SDEs of the problem. The system balancing penalty payments \mathcal{I} and the resulting cash flows are omitted from the plots to avoid clutter. At the same time the effects of day-specific stochastic realisations of X and Y , and of the other omitted variables, are partly visible in the behaviour of those variables actually plotted. The following results mostly use the default set of parameters shown in table 1.

(a) Optimised dynamics of the physical variables: sample paths and summary statistics

Simulated paths of the physical dynamics are shown for both a large and a small ESD in Figure 3. Taking first the large ESD in Figure 3(a), the fine dotted line plots the quantity of energy Q stored in the ESD. During the nightly peak of wind supply (hours 5 to 9) the optimal C^* completely fills the ESD. In this state no more energy can be stored, so all wind power delivered (solid line of delivery D) has Brownian disturbances. The coarse dotted line of the optimal commitment C^* balances the opposite risks of under-committing the Brownian output (causing a lower mean price for all amounts over-delivered) and of over-committing it (causing severe cash penalties for amounts under-delivered).

The commitment C^* is the locally horizontal coarse dotted line, whose hourly jumps in level were optimally selected one hour previously. In hours 1 to 5 (from midnight) the ESD fills rapidly, and when the ESD is full (saturated) between hours 5 and 9, deliveries (solid) are free to deviate from commitment (coarse dotted). Here the commitment does rise, but it remains low enough to

| | Parameter | Value | Parameter | Value |
|--------------------------------|-------------------|----------------------------------|---------------------|---|
| Wind Speed | σ_X | $0.2 \text{ h}^{-\frac{1}{2}}$ | $\bar{\theta}_X$ | 8 m s^{-1} |
| | α_X | 0.375 | ψ_X | 2 h |
| | κ_X | 0.1 h^{-1} | Capacity factor | 29.5% |
| Wind Power Conversion Function | Max output Rating | 1 MW 15 m s^{-1} | Cut out Constraints | 25 m s^{-1} $C_{\min} = 0, C_{\max} = 1$ |
| Smoothing ESD | Q_{\max} | 1 MWh | κ | $\sqrt{0.7}$ |
| | L_C | 250 kW | L_D | 250 kW |
| | ξ_C | 5 h^{-1} | ξ_D | 5 h^{-1} |
| Arbitraging ESD | Q_{\max} | 4 MWh | κ | $\sqrt{0.7}$ |
| | L_C | 1 MW | L_D | 1 MW |
| | ξ_C | 1.2 h^{-1} | ξ_D | 1.2 h^{-1} |
| Trading | σ_Y | $0.075 \text{ h}^{-\frac{1}{2}}$ | $\bar{\theta}_Y$ | 40 £/MWh |
| | α_Y | 0.375 | ψ_Y | 14 h |
| | κ_Y | 0.04 h^{-1} | Contract length | 1 h |
| | ζ | 0.5 | r_a | 0.05 pa |
| | C_{\max} | 1MW | r | $r_a / (365 * 24) \text{ h}^{-1}$ |

Table 1. Default parameters used in all calculations unless otherwise stated.




Figure 3. Monte Carlo simulations demonstrating the optimal strategy for (a) a large ESD ($Q_{\max} = 4$) and (b) a small ESD ($Q_{\max} = 1$). We show the % energy level in the ESD Q/Q_{\max} (dotted), the % commitment C^*/C_{\max} (dashed) and the % amount delivered D/C_{\max} (solid). Underneath we show the simulated processes X_t and Y_t (solid lines) used in both scenarios in (c) and (d) respectively. We also plot θ_X and θ_Y (dashed lines) for both processes to indicate if the levels are above/below expected.

ensure that almost all deliveries are over-deliveries. It is not in general optimal to avoid every risk, and close examination of the second half of hour 10 shows a brief threat of under-delivery, as a spike of demand causes a brief discharge from the ESD. Early in hour 11 the ESD is refilled swiftly, due to a low commitment. It is notable that this refilling is due to a prediction made during hour 9, not to error feedback of a stock fall-off during hour 10. This is because the demand spike arises late in hour 10, but the low optimised commitment C^* during hour 11, which swiftly recharges the ESD at the start of hour 11, was itself set at the start of hour 10.

From hour 12 to hour 24, the period of highest price, the commitment line C and the delivery line D coincide almost exactly. Because the optimal policy fills the ESD fully overnight (but not without some jumps in C^*), nothing more could have been done to increase the stock of energy available during the evening price peak.

Comparable results for the small ESD are shown in figure 3(b). The small ESD was expected to do continuous smoothing (only) because its small total capacity (one hour of maximum output) and its small rate of charge/discharge (one quarter of peak output) both seem to forbid price arbitrage. There are interesting similarities and differences between the optimal C^* trajectories for the large and small ESDs. Like the large ESD the small ESD takes the stored energy Q to higher levels during the night hours 5 to 9, but in this case not to saturation. Then both the large ESD and the small ESD run Q down during the evening peak hours, 16 to 19.

The intuition here is precisely that the small ESD is so small. On inspection of figure 3(b) the continuous delivery line D is mostly close to (but above) the stepped commitment line C^* . These facts, along with the high level of both D and Q at night, arise because when using only a small ESD, the WPG must deliver most of its power during the nightly peak of wind strength. To minimise the balancing penalties on these low-priced sales, it is optimal for the small ESD to remain unsaturated all of the night, in order to absorb fluctuations in the majority of its deliveries – which must be made at night. Despite this, at the end of the night the optimal policy retains a fairly large quantity Q in store. This reserve, used up in the peak, allows the small ESD to deliver the small quantity of energy that it does sell at evening peak prices exactly on commitment.

Overall, both the small and large ESDs are ‘arbitraging the precision of control’ towards the evening price peak. During the price peak both ESDs go from near full to near empty, and both deliver their cleanest power of the 24 hours (closest to commitment) at this time, when both selling prices and balancing penalties are at their 24 hour peak. Oddly however, the large ESD does NOT as intuitively expected ‘both arbitrage and smooth’. This again is because of the actual size of this ESD. The large ESD ‘chooses not to smooth’ during the night because the balancing penalties, from selling un-smoothed power during the night, are less than the arbitrage value of selling almost all of the stored power (less any used up for later smoothing) during the evening peak of prices and penalties. This suggests two conjectures: (i) only an even larger store than this large ESD would both arbitrage and smooth during all 24 hours; (ii) there must be diminishing returns to increasing the size of the ESD, because the optimal solution has already maximised the returns of a store of this size. These conjectures are explored in the next subsection, which briefly addresses the effect on value V of continuously improving the store capacity Q_{\max} and the charge/discharge rate \mathcal{L} .

The two time plots for the different ESDs are clearly instructive, and are long enough to show some statistically different variation within/between them. However each is one realised sample path over 24 hours, from an infinity of possible other paths from the same starting conditions, over all of which the diffusion solution has been integrated. It is however possible to compute the expected values (as opposed to sample values) of a wide range of statistical measures, when integrated over all the possible trajectories for a given optimised solution.

As [15] have pointed out, the optimal decision function $C^*(X, Y, Q, t)$ has set the values of all feasible changes in C^* to maximise V . Here V is the integral (from any point in the state space) of the expected time-discounted future value of \mathcal{I} . So far \mathcal{I} has been used as an indicator for each state-point of the actual rate of income earned/paid over dt when the system is at that state

space point. The Bellman-optimised control policy C^* at each state point simply maximises the probability-weighted (diffused) discounted value of \mathcal{I} along all trajectories from that state point.

Therefore we are able to compute the expected and discounted value, under $C^*(X, Y, Q, t)$, of any event or variable at each state point. To do so it is sufficient: (i) to retain the optimised $C^*(X, Y, Q, t)$; (ii) to replace the cash flow indicator function \mathcal{I} at each state point with an indicator function for the event of interest (e.g. to record the qualitative fact of a physical over-delivery at any state point, set an indicator function $1_{C>D}$; to record the size of a physical over-delivery set $C - D$; to record its square, set $(C - D)^2$, and so on); (iii) to run the model to equilibrium backwards in time using the new \mathcal{I} and unchanged C^* . The resulting value in each X, Y, Q, t cell is the expected discounted value under C^* (over the time horizon being evaluated) of the outcome indicator in question. From this discounted value and the interest rate r the undiscounted time-average of the required qualitative or quantitative variable over all trajectories can be recovered.

(b) Effects on annuity returns $AR = rV$ due to varying the ESD design

We return to comparing different ESD designs as in the subsection above, and we use the same system balancing penalty as before, namely $\zeta = 0.5$, but we now examine the effects on V due to continuous variations in some of the important and expensive (design) parameters namely Q_{\max} and L_C, L_D .

Each solution is a hypercube of $V(X, Y, Q, t)$ values, and \pounds year⁻¹ annuity values are given by the simple formula $AR = r_a V$ where r_a is the annual interest rate so as to remove the effect of the interest rate from the valuation of a stochastic stream of cash flows. For intuitive clarity the hypercube of V values is averaged over a truncated state space, so reducing the value function V to a single representative scalar. This seems reasonable as a first approximation, because it is not likely that an investor would base a decision to buy or design an ESD based on its behaviour inside any small region of the state space, or of the design space.

The maximised value of the combined WPG and ESD is found by optimising the commitment C as in (4.9) under the original forcing function \mathcal{I} (4.2) so as to calculate the net present value of operating the ESD in perpetuity. In the case of a known limit on system life this perpetuity value can be easily adjusted to a finite annuity. In general there might be large positive or negative payments at the end of the asset's life, but these are not considered here, as they are a well understood problem in more conventional 'real options' analysis.

The total problem of engineering and economic design for a complete WPG/ESD system has three main steps: (i) calculate the NPV of the system's maximised system earnings (income earned less system balancing penalties) for various combinations of design or operating parameters Q_{\max}, L_C and L_D ; (ii) calculate the capital costs of building an ESD with each set of design parameters, taking into account non-linear economies of scale in building and technology readiness risks; (iii) optimise the investment decision overall, using conventional real-option tools to model any long-term random or non-random walks in the financial variables, also including realistic full costs of entry/exit.

In this paper we briefly consider step (i) only, and only for the ESD, because this is where the present paper is more innovative. Hence we study the effects of changing the ESD's operating parameters ($Q_{\max}, L_C = L_D$) on its income after balancing penalties. In figure 4 we vary two design parameters of the modelled class of ESD: the horizontal axis continuously varies ESD capacity Q_{\max} , and the successively higher lines show the values of successively higher discrete rates of charge/discharge ($L_C = L_D$). All these curves share a near vertical region close to the origin. This suggests that much of the value of any ESD comes from smoothing over a period of 10 to 15 minutes, where this value is near-uniform for all the designs plotted. This is because Brownian variance is linear in time, and hence any losses proportional to the standard deviation are proportional to the square root of the time interval.

Diminishing returns are seen in figure 4 to increase for both the capacity Q_{\max} (continuous, left to right on the horizontal axis), and the charge-discharge rate $L_C = L_D$ (successive vertical




Figure 4. The annuity returns added by the ESD, $AR_{ESD} = r_a V - r_a V(Q_{max} = 0)$. The value is shown with L_C and L_D ranging from 250 kW up to 1 MW (both are kept equal) and for a market with system balancing penalty parameter $\zeta = 0.5$. Annuity returns without an ESD are $AR = rV(Q_{max} = 0) = 70583 \text{ £pa}$ with $\zeta = 0.5$.

jumps in the vertical axis direction). Taking $Q_{max} = 1$, $AR_{ESD} = 11500 \text{ £ year}^{-1}$, $L_C = L_D = 250 \text{ kW}$ as an arbitrary datum point for comparison, more value is added by quadrupling the charge/discharge rate $L_C = L_D$ to 1 MW than by quadrupling the capacity $Q_{max}=4$. The non-linear response of gross system value V to changes in Q_{max} and $L_C = L_D$ permits quite complex economic trade-offs between Q_{max} and $L_C = L_D$. The optimum ESD design maximises the (positive) NPV of ESD net earnings minus the capital cost of installing the ESD. The fact that there are sharply diminishing returns to increases in ESD capacity and ESD charge/discharge rate means that there can be (at most) one finite optimum size of ESD for a given WPG, if this pair is to be optimised in isolation from others, as seen easily in the unrealistically simple case where the capital cost of the ESD is linear in all of Q_{max} , L_C , L_D .

Looking next along the lowest value curve (for the lowest discharge rate) in figure 4, $\frac{\partial AR_{ESD}}{\partial Q_{max}}$ is approximately zero after $Q_{max} = 2$, therefore there is little gain from storage capacity exceeding 8 hours of the WPG’s peak output rate – at such a low discharge rate there is too little time for a larger ESD ever to empty or fill completely. In contrast for the higher discharge rates in figure 4, close to $L_D = L_C = 1 \text{ MW}$, the value is still increasing since $\frac{\partial AR_{ESD}}{\partial Q_{max}}$ is significantly larger than zero at the largest Q_{max} plotted, namely $Q_{max} = 4$. This was foreseen in the previous subsection, in our discussion of figure 3. However while there is still always a slight gain from a larger ESD (before its capital cost) the important feature of the top lines is that they crowd closer together. This must be due to the fixed size assumed for the WPG and to the fixed volatilities assumed for X and Y : eventually larger L_* rates add almost nothing to smoothing or arbitrage, and larger Q_{max} adds fewer and fewer opportunities for price arbitrage (due to moving deeper into the tail of the demand distribution, even before allowing for the capital costs of doing so). Within this logic the engineering design and the actual investment decision could be optimised simultaneously.

(c) Effects of the one-hour time delay between fixing C and the start of its effect

In Figure 5 the size of the system balancing penalty ζ varies left to right on the horizontal axis, and the effects of ζ on annuity return AR (vertical axis) are plotted for two states of ESD (present and absent) and for two rules of price setting (with and without an hour’s delay between setting the

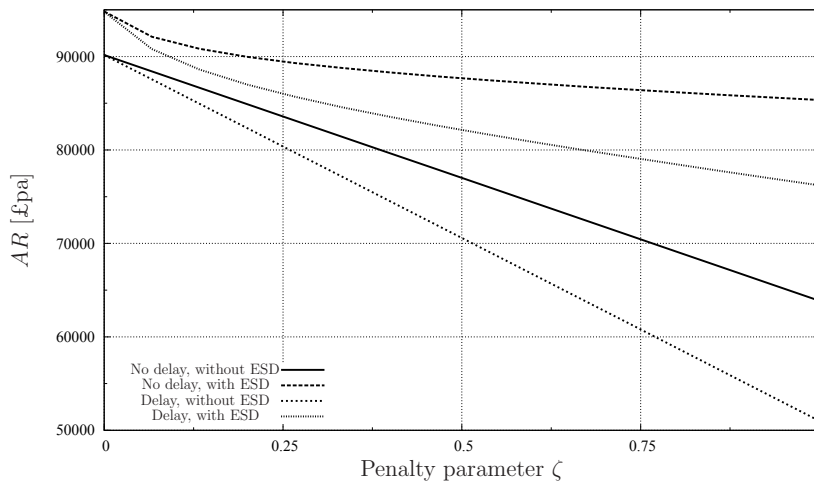


Figure 5. The annuity return $AR = rV$ from the wind farm combined with an ESD. The value is shown with and without a smoothing ESD for an immediate commitment market and for a delayed market, varying the system balancing penalty parameter ζ .

commitment C and the start of the hour committed). In all plots an increase in ζ reduces AR . The plotted falls in AR are all initially non-linear in ζ . Over the range $0 < \zeta < 0.12$ value falls steeply, but at a diminishing rate; after this the falls in AR in all plots tend towards linearity in ζ . This suggests some kind of saturation in the effects of higher ζ on the optimum policy itself: over the range $0 < \zeta < 0.12$, one can reduce penalties by making smaller commitments C . Losses tend to be linear in the remaining variable, the penalty itself. The limiting case is when the ESD's capacity is zero, and when the commitment C must be set one hour before it becomes operative; both apply to the lowest plot on figure 5. This plot shows the lowest level of AR , and also the earliest approach to linearity in ζ . Conversely the highest of the value plots in this figure represents the case with the highest physical and pricing flexibility. These plots show how changes in the market arrangements (a delay in implementing C , and an increase in the system balancing penalty ζ) can change the optimum size of an ESD, as well as changing the optimum rules for operating it.

6. Conclusions and future work

In contrast to the classic Black Scholes equation, which is purely financial and has no inertia or dynamic lags, the PDE of this paper incorporates three modifications to enable it to model a system which is partly physical and/or has 'lagged' deterministic dynamics:

- there is a forcing term, \mathcal{I} [$\mathcal{L} \text{ h}^{-1}$], which defines a rate of income or cost at every point in the state space (this ensures that the system's value is decided not at the instant of its death, as in most financial options, but by the integration of income and expenditure over its complete lifetime);
- there is a control variable C which is continuous, but is reset only at discrete time intervals; and
- there is an 'integration term', that can model deterministic physical and/or financial dynamics.

This integration term has the PDE's units of [$\mathcal{L} \text{ h}^{-1}$], but it is the product of a purely physical coefficient $\frac{dQ}{dt}$ [MW] (which describes the deterministic dynamics of the storage system) and a purely financial coefficient $\frac{\partial V}{\partial Q}$ [\mathcal{L} / MWh] which is calculated by standard backward Black Scholes diffusion methods. This type of PDE is suited to the stochastic dynamic optimisation of

systems in which some of the physical and/or financial variables have deterministic dynamics. Results in the present ‘proof of mathematical concept’ model show that in a future power system which has high wind penetration, and therefore high system balancing penalties, a complex dynamic balance must be, and can be, maintained in real time, between (i) present arbitrage profits; (ii) all future arbitrage profits; (iii) the present system-balancing penalties; and (iv) all future system-balancing penalties. Because this model sets absolute physical bounds and/or absolute physical values, for both the WPG parameters and the ESD parameters, the model can in principle jointly optimise the engineering design parameters, the operating rules and the economical performance – with real options methods available to deal with longer-term stochastic price drifts.

Future work could include energy management within the supply grid and storage problems in general, not restricted to energy – e.g. financial system design/operation. Taking the energy opportunities would involve: better estimates of the parameters for weather, prices, energy storage and trading, plus tests for the model’s robustness to errors in these estimates; wider opportunities within grid storage which would need more realistic dynamic models of energy storage (e.g. electrical circuits and/or mechanical devices for energy storage); and more general ‘new modelling topics’ including the optimal design and/or operation of energy storage in aircraft, ships and cars.

Data Accessibility. All software used was written in c++ code, parameter inputs are provided in the text. The code is available to download from <https://github.com/pjohn/WindFarmESDValuation>.

Authors’ Contributions. PJ participated in the design of the model, wrote the numerical software to solve the problem, analysed the results and drafted the manuscript; SH conceived of the model, analysed the results and helped draft the manuscript. PD helped design the numerical software and drafted the manuscript. All authors gave final approval for publication.

Competing Interests. We have no competing interests.

Funding. PJ on The Engineering and Physical Sciences Research Council (EP/F027842/1)

Acknowledgements. We would like to thank two anonymous referees for providing extensive and constructive feedback that helped improve this work.

References

1. Thompson M, Davison M, Rasmussen H. 2009 Natural gas storage valuation and optimization: A real options application. *Naval Research Logistics (NRL)* **56**, 226–238.
2. Kitapbayev Y, Moriarty J, Mancarella P. 2015 Stochastic control and real options valuation of thermal storage-enabled demand response from flexible district energy systems. *Applied Energy* **137**, 823–831.
3. Moriarty J, Palczewski J. 2017 Real option valuation for reserve capacity. *European Journal of Operational Research* **257**, 251–260.
4. Earl N, Dorling S, Hewston R, Von Glasow R. 2013 1980–2010 variability in UK surface wind climate. *Journal of Climate* **26**, 1172–1191.
5. Correia P, de Jesus JF. 2010 Simulation of correlated wind speed and power variates in wind parks. *Electric Power Systems Research* **80**, 592–598.
6. Hill DC, McMillan D, Bell KR, Infield D. 2012 Application of auto-regressive models to UK wind speed data for power system impact studies. *IEEE Transactions on Sustainable Energy* **3**, 134–141.
7. Sinden G. 2007 Characteristics of the UK wind resource: Long-term patterns and relationship to electricity demand. *Energy Policy* **35**, 112–127.
8. Park BT, Petrosian V. 1995 Fokker-planck equations of stochastic acceleration: Green’s functions and boundary conditions.

- The Astrophysical Journal* **446**, 699.
9. Früh WG. 2013 Long-term wind resource and uncertainty estimation using wind records from Scotland as example.
Renewable Energy **50**, 1014–1026.
 10. Lucia JJ, Schwartz ES. 2002 Electricity prices and power derivatives: Evidence from the Nordic power exchange.
Review of Derivatives Research **5**, 5–50.
 11. Brennan MJ, Schwartz ES. 1982 An Equilibrium Model of Bond Pricing and a Test of Market Efficiency.
Journal of Financial and Quantitative Analysis **17**, 301–329.
 12. Howell SD, Duck PW, Hazel A, Johnson PV, Pinto H, Strbac G, Proudlove N, Black M. 2011 A partial differential equation system for modelling stochastic storage in physical systems with applications to wind power generation.
IMA Journal of Management Mathematics **22**, 231.
 13. Øksendal B. 2003 *Stochastic Differential Equations*. Springer.
 14. Avalos JH, Johnson PV, Duck PW. 2015 SLADI: a semi-Lagrangian alternating-direction implicit method for the numerical solution of advection–diffusion problems with application to electricity storage valuations.
Journal of Computational Finance **19**, 69–108.
 15. Howell SD, Johnson PV, Duck PW. 2011 A rapid PDE-based optimization methodology for temperature control and other mixed stochastic and deterministic systems.
Energy and Buildings **43**, 1523–1530.