Strathprints Institutional Repository

Howell, S. and Duck, P.W. and Pinto, Helena and Strback, G. and Hazel, A. and Proudlove, N. and Black, M. (2006) *A PDE and option based approach to valuing and designing stochastic storage for wind-generated electricity.* In: Portuguese Finance Network 4th Finance Conference, 2006-07-06 - 2006-07-08, Porto. (Unpublished)

http://strathprints.strath.ac.uk/

A PDE and option-based approach to valuing and designing stochastic storage for wind-generated electricity

S. D. Howell, P. W. Duck, H. Pinto, G. Strbac, *Member, IEEE*, A. Hazel, N. Proudlove and M. Black

Abstract—Significant penetration of wind generation will inevitably impose additional requirements on the remaining large conventional plant to deliver both the flexibility and reserve necessary to deal with variability and unpredictability of wind power, which will inevitably have cost implications. Energy storage systems appear to be an obvious solution for dealing with the unpredictability of renewable sources: during periods when intermittent generation exceeds the demand, when the surplus could be stored and then used to cover periods when the load is greater than the generation. However, views on the role of bulk storage remain highly controversial and, somewhat surprisingly, there has been very little work carried out to demonstrate the necessity (or otherwise) and economics of storage based applications in systems with high penetration of wind energy. In this paper we assess the economic value of such storage by an innovative real-options method. Solution of the resulting PDE (partial differential equation) requires novel numerical techniques that are developed and successfully applied in this paper (and which yield results in excellent agreement, but at considerably less cost, than full simulations). This PDE approach may have applications in many physical, engineering and economic systems.

Index Terms—wind energy, storage value, stochastic storage

I. INTRODUCTION

It seems likely that renewable and other low carbon energy sources will become a major part of the future electricity generation system for various reasons. Large penetration of wind power and other forms of renewable generation by 2020 and beyond may displace significant amounts of energy produced by large conventional plant. However as variability and non-controllability are inherent characteristics of wind energy, the ability to maintain the balance between demand and supply in systems with large penetration of wind generation is one of the main concerns in system operation.

Increased dependence on wind generation will impose additional requirements on the remaining large conventional plant to deliver both the flexibility and reserve necessary to maintain the continuous balance between load and generation, which will inevitably have cost implications [1]. Recently completed studies (e.g. SCAR in the UK [2], DENA in Germany [3]) investigated a number of possible scenarios showing that extending variable renewable generation would increase system costs associated with the integration of this generation into the operation and development of the power system. The

S. D. Howell and N. Proudlove are at Manchester Business School, P. W. Duck and A. Hazel are at the School of Mathematics, University of Manchester, H. Pinto is at the Accounting and Finance Department, University of Strathclyde, G. Strbac is at the Department of Electrical and Electronic Engineering, Imperial College, London and M. Black is at CE Electric UK.

analysis of total costs between (i) balancing and capacity, (ii) transmission, and (iii) distribution networks, demonstrated that balancing and capacity costs dominate all other costs.

Managing the control of risk in energy supply has of late become an important topic in the commercial sector. Various aspects of the management of supply have been undertaken in [4], [5], [6], [7] and [8].

The area of focus of this paper is balancing demand and supply in real time, in systems with significant penetration of wind power. In order to deal with the corresponding increased uncertainty, the system will require increased amounts of reserve. When analysing the need for additional reserve requirements time horizons of up to several hours are normally considered [2]. Traditionally required reserves could be provided by a combination of synchronised and standing reserve. In order for synchronised conventional plant to provide reserve it must run part loaded. Thermal units operate less efficiently when part loaded, with an efficiency loss of between 10% and 20%. Whenever this reserve is called on, other plant with higher marginal cost must be run, and this is another source of cost.

In addition to synchronised reserve, provided by part-loaded plant, the balancing task will also be supported by standing reserve, supplied by higher fuel cost plant, such as OCGTs (Open Cycle Gas Turbines) and pump-hydro storage facilities. The allocation of reserve between synchronised and standing plant is a trade-off between the cost of the efficiency losses of part-loaded synchronised plant (plant with relatively low marginal cost) and the cost of running standing plant with relatively high marginal cost. The balance between synchronised and standing reserve could be optimised to achieve a minimum overall reserve cost of balancing [9]. Another possibility is energy storage at some point(s) in the supply-demand chain.

In this paper we develop a novel technique to assess the value of a generic, fully flexible, energy storage technology when used to absorb random fluctuation of wind power in real time. The paper is not concerned with the cost of particular storage technology, but rather with quantifying the benefits that storage technology could bring to the system by increasing the efficiency of its operation.

The wind-power output is assumed to be subject to Brownian motion (BM) disturbances at short lead times of up to several hours. In a simplified representation of the operation of the power system, we assume that the BM disturbance is absorbed entirely by energy storage. We demonstrate that the economic value of such a storage system can be evaluated by an innovative real-options method. Solution of the resulting

PDE (partial differential equation) requires novel numerical techniques that are developed and successfully applied in this paper and which provide results in very close agreement with a simulation approach, but at a fraction of the cost. The overall approach developed is then used to analyse the impact of storage design parameters, such as capacity and power ratings, on the value of storage.

The results of our present paper are limited to determining and solving the PDE when the stochastic storage system is left in an unregulated state, which means that it is left to fill and empty randomly (following fluctuations in wind in the time horizon of several hours). Whilst this level of modeling is mathematically fundamental, and provides many insights in its own right, it can and should be developed further to include an optimal control model, in order to maximise the value of stochastic storage and generators used jointly.

The structure of the paper is as follows: in section II we give the economic background and develop the model as a system of stochastic differential equations (SDEs); in section III we discuss how the same problem may alternatively be solved by simulation or a PDE specified outside the time domain, together with the necessary solution procedures for each, reserving technical detail for the Appendix; in section IV we present a comparison of results from the SDE simulation and PDE approaches, showing the dynamics of the uncontrolled case, varying both the theoretical assumptions and the various system parameters. Section V presents some conclusions.

II. THE STOCHASTIC STORAGE SYSTEM

We value a stochastic storage system for electricity, intended to absorb BM disturbances to the forecast of wind power output. As a base case we assume that the storage unit is left indefinitely to absorb the disturbances, filling when wind power exceeds forecast and emptying when it falls short, bounded only by the unit's maximum energy capacity and by its power ratings for charge and discharge (our model includes different power ratings for charge and discharge, but for simplicity our worked examples assume these to be identical). This assumes that an unforeseen surplus of wind will be wasted if it cannot be stored, and that when the storage unit meets an unexpected deficit of wind power, it avoids the use of fossil fuel generation to supply the deficit. The benefit of storage is equivalent to the cost of fuel saved, to which might be added the value of emissions saved. An equivalent way of valuing the storage system that we adopted for the purpose of this paper, is to calculate the earnings that storage makes while it is discharging electricity, and the system's capital value is the net present value (NPV) of all the electricity it will ever discharge.

A. Analogy with real options theory

There are two principal approaches (time domain and PDE) to finding the capital value of the storage system, as defined above, both of which we will use, and we will show that the agreement between the two is excellent. In the time domain approach we start from an identified time origin, for example $t = 0$ and sum all expected earnings to $t = \infty$. This can be undertaken in principle by direct integration (not analytically feasible in this case), or by using simulation to sample the required expectation. This is performed by simulating repeated long paths through the time domain, summing the NPV of income over each path, and averaging this NPV over a large sample of such paths. The result is an estimate of the system's value at time $t = 0$ which we denote as V_0.

The second valuation method which we use (and which turns out to be far superior) is to express the problem as a PDE which can be solved numerically. The PDE approach gives us the freedom to eliminate time as an explicit variable (although we could introduce terms in the time domain to the PDE if needed, for example to represent daily cycles in electricity demand). In financial terms, the elimination of time treats the storage system as a perpetual option, which our PDE method values directly. More precisely we treat the storage system as an infinite bundle of options to charge or discharge, one of which matures at every instant of time out to infinity. Each such option is exercised automatically at the instant it matures, but only the discharges ever produce income. The value of this bundle of options is conditional only on the observed values, at the moment of valuation, of two related Markovian variables X and Q, which we next define.

The error in forecasting the rate of wind generation (actual − forecast) is X, a physical quantity measured in gigawatts. When $X < 0$ the storage system discharges, which drives an instantaneously variable flow of income. This income is analogous to a flow rate of wealth, or to an instantaneous rate of income, but not to a stock of wealth, or to a capital value, which is generally made up of the value of any income received at the present instant, plus a right to receive some set of subsequent, possibly stochastic, payments. This is a different situation from the standard Black-Scholes model, where the (geometric) BM variable, S, is already a traded capital value, being the price of a unit of stock, (and hence in general it is already the integral of the stock's discounted future income and capital flows). In our model we must compute the option's capital value V directly from its own income flow, which means simultaneously valuing an infinity of options, whereas in standard Black-Scholes applications we can define V in terms of its right to exchange for some other capital value S whose current value is exogenously given. We will see below that this actually makes the derivation of our PDE slightly simpler than that of the standard Black-Scholes PDE [10].

B. Model development

At the lead times that will eventually be used for optimal management of the storage unit (several hours) the behavior of X over the time increment dt can be well approximated as a BM. The reason for this is the well known persistence effect, meaning that the short term forecast of future wind power output is the present output. This is generally valid for a horizon of up to a few hours [1], which is the typical lead time considered in this paper. For this to be true at all lead times, up to some practical horizon, wind strength must be a close approximation to a BM up to that horizon. However in practice

the output of the generating system will be adjusted at lead times of minutes to hours, in response both to the actual wind power output, and also (if storage is used) to the actual quantity of electricity in store. At such lesser lead times persistence applies, and a BM model can be used as an approximation, namely

$$\mathrm{d}X = \sigma \mathrm{d}W, \tag{1}$$

where σ denotes the volatility of X, and $\mathrm{d}W$ denotes a random variable, drawn from a normal distribution with zero mean and variance $\mathrm{d}t$, the (small) timestep.

From the viewpoint of option valuation, the system behaves differently when charging and when discharging. Hence for PDE purposes (only) we subdivide the function for the system's total value at some time $t = 0$ (call this $V_0(X_0, Q_0)$ as at time zero) into two alternative value functions one or other of which applies at every time t, namely $G(X, Q)$ if the system is charging ($X > 0$) and $F(X, Q)$ if the system is discharging ($X < 0$). We suppress time dependence for G and F since we will define them as the values of perpetual options, but we retain time dependence for V_0 since we define this as an integral in the time domain of expected discounted income, from $t = 0$ to infinity. We discuss later why these alternative specifications give the same value.

C. Charging region: $X > 0$

When $X > 0$ the wind power output is larger than forecast. We assume that the excess electricity is available to charge the storage system, if not already full, so we call this region of X the charging region. The quantity of electricity in store is given by $Q \in [0, Q_{\max}]$, and is measured in gigawatt hours; here $Q_{\max}$ is the maximum capacity of the storage system, and its minimum capacity is zero.

The charge rate (at which electricity enters the system) is at every instant a function of X and Q, denoted by $\mathcal{L}_c(X, Q)$ which will depend on the properties of the storage system. All storage devices, however, must satisfy the physical constraint that when the storage device is full, the charge rate is zero, $\mathcal{L}_c(X, Q_{\max}) = 0$. We first assume a simple piecewise linear form for the charge rate.

Initially we define $\mathcal{L}_c(X, Q)$ away from the boundary $Q = Q_{\max}$ as

$$\mathcal{L}_c(X, Q) = \min\{X, X_c\}. \tag{2}$$

Here X_c is the maximum rate at which the storage system can be charged, known as the (charging) power rating. It is easy to implement more complex specifications in order to reflect more complex system responses, for example a dead zone around $X = 0$, or a more general non-linear $\mathcal{L}_d(X, Q)$.

We did in practice implement a further complexity, by specifying $\mathcal{L}_d(X, Q)$ as a function of both X and Q. The main motive for this was to improve numerical stability (but this could also model the damped responses by a physical storage system), as follows:

$$\mathcal{L}_c(X, Q) = \min\{X, X_c, \lambda_c(Q_{\max} - Q)\}. \tag{3}$$

Here λ_c is a positive constant with dimensions of hours^{-1} used to adjust the point at which the charge rate begins to

decrease as the storage device approaches full capacity. It can be justified on a number of grounds, including numerical expediency (no additional artificial boundary conditions are required to close the problem). However, in reality, some intensive numerical testing revealed the solution to be highly insensitive to the prescribed values. Even by changing the value by a factor of 20 made less than a 1% change to the solution.

In the charging region, the dynamics of the quantity in the storage system, Q, are given by:

$$\mathrm{d}Q = \mathcal{L}_c \mathrm{d}t. \tag{4}$$

In the time domain we can use the same notation for value changes in both the charge region and the discharge region. For a time origin of $t = 0$ we have at all times in the charging region ($X > 0$)

$$\mathrm{d}V_0(X_0, Q_0) = 0. \tag{5}$$

Note that the units of V_0 are taken to be (net present) gigawatt hours, because we are not yet assigning a price to electricity.

D. Discharging region: $X < 0$

When $X < 0$ the output of wind generation is smaller than forecast and there is a deficit of wind-generated electricity. We assume here (before considering the optimal joint operation of electricity storage and conventional generation) that all deficits of wind power will be met from the storage system if $Q > 0$. The storage system will thus be partially discharged, and so this region of X is termed the discharging region.

The discharge rate of the storage system is denoted by $\mathcal{L}_d(X, Q)$ and is, in general, different from $\mathcal{L}_c(X, Q)$. The required physical constraint is that the discharge rate must be zero when the storage system is empty, $\mathcal{L}_d(X, 0) = 0$. We again assume a piecewise linear form for the discharge rate in the X dimension, and the discharge rate is by convention negative, to reflect the positive sign of X in the discharging region:

$$\mathcal{L}_d = -\min\{|X|, X_d, \lambda_d Q\}. \tag{6}$$

Here the positive constant X_d is the system's power rating (maximum rate) for charging, and λ_d is a positive constant with dimensions of hours^{-1} which is used to adjust the point at which the charge rate begins to decrease when the storage device is nearly empty (in a manner analogous to the $\mathcal{L}_c \to 0$ limit as $Q \to Q_{\max}$ in the charging region). The use of a suitable value of λ_d in this region avoids an instantaneous change of the (dis)charging rate from X or X_d to zero as Q approaches its boundary ($Q = 0$), thereby avoiding both physical unrealism and problems of numerical convergence. A similar specification is also useful for Monte Carlo simulation in the time domain (discussed below) where it helps to ensure that randomly simulated realizations of $\mathcal{L}_d$ and $\mathcal{L}_c$ do not violate the bounds on Q.

Figure 1 shows a surface plot of $\mathcal{L}_c$ and $\mathcal{L}_d$ (versus X and Q), illustrating the various regimes (plateaus and linear behaviors) of the (dis)charge rates. The topology of this function comprises plateau regions ($\mathcal{L}_c = 1$, $\mathcal{L}_d = -1$), connected in a linear fashion to prescribed boundary values of zero (as
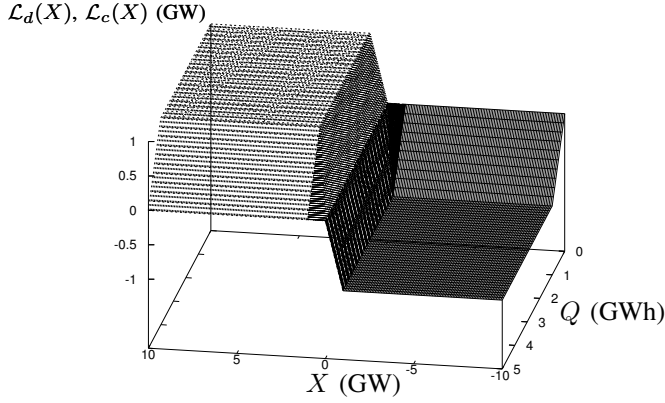
Fig. 1. Variation of $\mathcal{L}_d(X)$, $\mathcal{L}_c(X)$, for $Q_{\max} = 5$GWh, $X_c = X_d = 1$GW, $\lambda_c = \lambda_d = 1\text{h}^{-1}$

a result of the introduction of λ_c and λ_d) and to each other. It should however be emphasised that the methodology in this paper permits entirely arbitrary topology of these functions.

The dynamics of the quantity in the storage system, Q, in the discharging region are then given by:

$$\mathrm{d}Q = \mathcal{L}_d \mathrm{d}t. \tag{7}$$

We model the electrical efficiency of storage by the parameter k where $0 < k < 1$. Whilst the system is discharging, any previously input charge is being depleted at the rate $\mathcal{L}_d$, but the actual rate at which discharged electricity is available to users is only k per unit (gigawatt-hour) discharged. Hence, the system's income during a small time interval $\mathrm{d}t$ of discharge is the positive quantity $-k\mathcal{L}_d\mathrm{d}t$. The constant k reflects the imperfect electrical efficiency of a single cycle into and out of the storage unit, and is independent of time. This is an arbitrarily chosen model of efficiency, in which losses occur only during discharge, but we may note that a slightly higher valuation results if the total efficiency of k per cycle is factored into a product of efficiency on charging and efficiency on discharging.

When valuing in the time domain, the NPV (as at time $t = 0$) of the income generated from discharging during the time interval $\mathrm{d}t$ (subject to $X < 0$) is therefore, if the time increment ends at time t,

$$\mathrm{d}V_0 = -k\mathcal{L}_d\mathrm{e}^{-rt}\mathrm{d}t, \tag{8}$$

where r denotes the interest rate (assumed constant) and the subscript zero denotes that the present value is as at time $t = 0$. The units of $\mathrm{d}V_0$ (and of V_0) are taken to be gigawatt hours, for reasons already given. V_0 is simply the integral of the expectation of the last expression, i.e.

$$\mathrm{E}\left(-k\int_0^\infty \mathcal{L}_d\mathrm{e}^{-rt}\mathrm{d}t\right), \tag{9}$$

where $\mathrm{E}()$ is the expectation operator. Because the above integral cannot be evaluated explicitly, we next consider Monte Carlo simulations of this integral in the time domain.

E. Monte Carlo simulations

It is conceptually straightforward to perform a Monte Carlo simulation using the stochastic system given by equations (1) – (8). The simulation calculates an estimate of the NPV of the system, $V_0(X_0, Q_0)$, by repeatedly generating a realization of a path through (X, Q) space in the time domain starting at (X_0, Q_0) at time zero, and proceeding in prescribed (finite but small) time increments $\mathrm{d}t$ out to a long time horizon. Equation (1) is used to calculate the increment in X ($\mathrm{d}W$ is determined from a standard library package) and the value of X is then updated. If $X > 0$, $\mathrm{d}Q$ is evaluated using (4), and $\mathrm{d}V_0 = \mathrm{d}G_0$ follows from (5). If $X < 0$, $\mathrm{d}Q$ is evaluated from (7) and $\mathrm{d}V_0$ is calculated via (8). If $\lambda_c > 1/\mathrm{d}t$ or $\lambda_d > 1/\mathrm{d}t$, then Q could leave the domain $0 \le Q \le Q_{\max}$. If realized Q is predicted to drop below zero during the time increment, then Q is set to be zero at its end, whereas if Q is predicted to rise above $Q_{\max}$, its value is set to $Q_{\max}$ (this proved generally not to be an issue).

For the parameter range of interest, it was found that time horizons of the order of 200 years were necessary in order to obtain a 'steady state' or terminal NPV for each sampled path, and large numbers of sampled paths were required to ensure stability of their mean. We defer discussion of detailed numerical results thus obtained until section IV, and will first describe the alternative (PDE) approach.

III. THE DIFFERENTIAL EQUATIONS FOR THE VALUE OF THE STORAGE SYSTEM

The behavior of the value $V(X, Q)$ of the storage system differs between the charging and the discharging regions, so we will define

$$V(X, Q) = \begin{cases} G(X, Q), & \text{charging region } (X > 0), \\ F(X, Q), & \text{discharging region } (X < 0). \end{cases}$$

Over the time increment $\mathrm{d}t$ there is stochastic fluctuation in X, which in a deterministic way changes Q. These changes in X and Q change the capital value of the system in two ways. Firstly they change the capital value of the system's options to discharge and to recharge in future. Secondly (in the discharging region only) they contribute income by discharging stored electricity. Current income is a tiny fraction of total capital value, which is dominated by the discounted expectation of future income. Using Itô's Lemma to determine the approximate change in value of the stochastic terms, we obtain

$$\mathrm{d}G = \frac{1}{2}\sigma^2\frac{\partial^2 G}{\partial X^2}\mathrm{d}t + \frac{\partial G}{\partial X}\mathrm{d}X + \frac{\partial G}{\partial Q}\mathrm{d}Q, \tag{10}$$

$$\frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial X^2}\mathrm{d}t + \frac{\partial F}{\partial X}\mathrm{d}X + \frac{\partial F}{\partial Q}\mathrm{d}Q - k\mathcal{L}_d\mathrm{d}t, \tag{11}$$

where the term $-k\mathcal{L}_d\mathrm{d}t$ is the income term in the discharging region.

Note that the above forms of the expressions for $\mathrm{d}F$ and $\mathrm{d}G$ are standard in the finance literature. An intuition for these is that they resemble expansions by Taylor's theorem, but with the addition of terms of the order of $\mathrm{d}t$. A simple explanation

of how Itô's Lemma produces these deterministic terms when X is a BM variable can be found for example in [10].

A PDE for the value of the system can be derived from the above in several ways, again using methods standard in the finance literature (although the equation that results is novel). It is essential to eliminate the stochastic term $\mathrm{d}X$, which is unknown. It is not sufficient merely to replace $\mathrm{d}X$ with its expectation of zero, since it can be shown that any resulting valuation is exposed to unknown and highly variable risks, and would therefore require an unknown and highly variable discount rate. The standard approach is to eliminate the effect of $\mathrm{d}X$ over $\mathrm{d}t$ entirely, by constructing some kind of hedge transaction. In the standard case the stochastic variable X is a capital money value (such as a stock price), and the hedge involves buying or selling ahead some fixed quantity of X.

Our system is subject to a stochastic X which is physical (and $\mathrm{d}X$ is a rate of flow of value, rather than a stock of value). To solve this non-standard problem, we construct a hedge against a change in the physical flow rate itself, rather than hedging against the changes in the value of a stock variable (since the X alternates between negative and positive, it has no useful natural equivalent as a stock). Note in this model, physical capital value changes in response to changes in a physical rate of flow of stock, holding price fixed, whereas in standard financial option and real option models capital value changes due to a fluctuating price for a fixed physical quantity (one unit) of stock. The mathematical formalism is identical in either case, provided we can define a meaningful hedge transaction, with (in theory) a negligible implementation cost.

Clearly it is reasonable to assume that any form of energy can be traded. To achieve the required hedge of the physical flow rate X we can take an ownership position, long or short as needed, over the time increment $\mathrm{d}t$, in the physical resource which the storage system is itself intended to hedge, namely the instantaneous wind power output. For example we can sell, for the time $\mathrm{d}t$ ahead (some fraction Δ of) the forecast quantity of wind-generated energy, and we buy ahead (the same fraction of) the actual physical quantity of wind-generated energy that will be generated. This portfolio gives a surplus or deficit of electricity supply rate at the end of $\mathrm{d}t$ proportional to $\mathrm{d}X$, whose sign and magnitude can be selected to offset the change in capital value which the storage unit will have undergone due to $\mathrm{d}X$.

We therefore consider a portfolio of the following form (illustrating the discharge region, where $X < 0$)

$$\Pi = G + \Delta X.$$

In order for the stochastic terms to balance, $\Delta = \partial G/\partial X$, when $X > 0$ and $\Delta = \partial F/\partial X$, when $X < 0$, so

$$\frac{1}{2}\sigma^2\frac{\partial^2 G}{\partial X^2}\mathrm{d}t + \frac{\partial G}{\partial Q}\mathrm{d}Q = rG\mathrm{d}t, \quad X > 0,$$

$$\frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial X^2}\mathrm{d}t + \frac{\partial F}{\partial Q}\mathrm{d}Q - k\mathcal{L}_d\mathrm{d}t = rF\mathrm{d}t, \quad X < 0.$$

We use the expressions (4) and (7) for $\mathrm{d}Q$ in the appropriate

regions and divide the equations by $\mathrm{d}t$ to obtain:

$$\frac{1}{2}\sigma^2\frac{\partial^2 G}{\partial X^2} + \mathcal{L}_c\frac{\partial G}{\partial Q} - rG \quad = 0, \quad X > 0, \quad (12)$$

$$\frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial X^2} + \mathcal{L}_d\frac{\partial F}{\partial Q} - rF - k\mathcal{L}_d \quad = 0, \quad X < 0. \quad (13)$$

At the boundary between the charge and discharge regions, where $X = 0$, we require continuity of the values F and G and their derivatives in the X-direction (value matching and smooth pasting):

$$G(0, Q) = F(0, Q), \quad (14)$$

$$\frac{\partial F}{\partial X}(0, Q) = \frac{\partial G}{\partial X}(0, Q). \quad (15)$$

Equations (12) and (13) are somewhat novel. The form of the equation for each region resembles a diffusion, but the first order partial differential term is w.r.t. Q rather than w.r.t. time (which does not appear). Also the boundary values for value along $Q = Q_{\max}$ and $Q = 0$ are not specified by the problem, as they are in standard Black-Scholes derivations. The only externally verifiable input to the system's economic value, in either region, is the income term $-k\mathcal{L}_d\mathrm{d}t$ which exists only in the discharging region as a deterministic payment, and then only over the current time increment $\mathrm{d}t$. All the remaining dynamics describe the diffusion of value towards or away from the various (X, Q) states where deterministic income can be earned. The entire value in the charging region, G in (12), arises from diffusion terms, describing the diffusion of G itself, from $G = 0$ at $X = \infty$ to the smooth pasting condition at $X = 0$ where $G = F$. More simply, the system's value G in the charging region is purely the value of its option to move into the discharging region, F, starting from the present values of Q and X. The system's value F in the discharging region in (13) is driven, in addition, by a deterministic income term as well as by the diffusion of F itself. The latter diffusion represents (most importantly) the 'option' which a discharging system has, to move into the charging region when X next becomes positive. In both equation (12) and (13) the first term is the contribution to total option value from the future stochastic diffusion of X, and the second terms arise from the locally deterministic changes to capital value that $\mathcal{L}_c$ and $\mathcal{L}_d$ make to option value by respectively increasing and decreasing Q. The third term in each equation ensures that in every state of X and Q, and over every time increment $\mathrm{d}t$, the combined effects of the system's capital gain (or loss) on its future options, together with any current income, will generate the risk free rate of return r.

A. Equivalence between the PDE solution and solution by simulation in the time domain

An informal explanation of why the PDE solution is equivalent to a sufficiently detailed simulation in the time domain is as follows: the PDE is a model of events over a general time increment $\mathrm{d}t$ (recalling that in its derivation we extract a common factor $\mathrm{d}t$). Consider (13) and restrict attention to the last two terms on the left-hand-side. This is simply the ODE whose solution is the present value at the beginning of $\mathrm{d}t$ of

a deterministic payment $-k\mathcal{L}_d \mathrm{d}t$ received at the end of the time increment $\mathrm{d}t$. Clearly if this deterministic payment were known to arrive after a larger number of time increments $\mathrm{d}t$, this ODE will correctly discount its present value over all of them. However in the present model, all income payments arising after the current time increment $\mathrm{d}t$ are stochastic. The presence of the first two terms (diffusion terms) in the equations (12) and (13), allows for stochastic diffusion, in addition to the accrual of a rate of return over time, for any payments that may arise after $\mathrm{d}t$. In effect therefore the joint solution to the PDEs for F and G replaces all stochastic payments after the end of $\mathrm{d}t$ with the discounted value of the future distribution over their expected value.

However because the option is a perpetual one, the PDE solution at any given pair of X, Q values holds for every $\mathrm{d}t$. Since there is an infinite sequence of possible time increments $\mathrm{d}t$, the solution must hold, as a discounted expected value, for all of them, and therefore holds over any and every path of infinite length starting from a given pair of X, Q values, at any time t. A special case of this is when the values X and Q are observed at some specific time $t = 0$. The resulting value is what is estimated by Monte Carlo simulation in the time domain: there is no explicit calculation of the future distribution of expected values, just a random drawing of many paths of income over a long time horizon, and at a high frequency of sampling. Convergence with the PDE solution occurs if the sample size, time horizon and sampling rate are sufficiently high. The PDE solution simultaneously calculates the solution at every point in the chosen (X, Q) domain, whereas the simulation result holds only at a single chosen (X, Q) starting point. The PDE approach therefore has significant practical advantages.

B. Practical stages in solving the PDE

Since the units of F and G are both gigawatt hours, for consistency in applications the volatility σ and interest rates must be expressed hourly. An adjustment for the selling price of electricity could easily be incorporated by multiplying the income term $k\mathcal{L}_d$ by the appropriate price, *i.e.* there is a linear relationship between the capital value of the system and the electricity price (and indeed with the efficiency k).

Consider a numerical solution procedure for the PDE system (12) - (15). From the mathematical point of view, the intriguing feature is that the overall system is forwards/upwards (increasing Q) parabolic if $X < 0$, but backwards/downwards (decreasing Q) parabolic if $X > 0$. The difference is a consequence of the different signs of $\mathcal{L}_d$ and $\mathcal{L}_c$, which crucially multiply the $\frac{\partial F}{\partial Q}$ and $\frac{\partial G}{\partial Q}$ terms, respectively, and it must be reflected in the solution procedure.

Initially, a preliminary guess was made, typically $F(X, Q) = kQ$ and $G(X, Q) = 0$, or, if available, the converged results from a previous neighboring calculation. A standard Crank-Nicolson marching procedure was adopted, with a rectangular set of mesh points in X, Q space, coupled with line relaxation. Typically 3201 X points and 101 Q points were employed. For a numerical solution, the X domain must be truncated to a finite range, $[-X_{\max}, X_{\max}]$ and

appropriate conditions should be applied on the boundaries $X = \pm X_{\max}$. The behavior of the solution as $|X| \rightarrow \infty$ is discussed in detail in the Appendix, and is used to formulate the (Robin-type) boundary conditions applied at these boundaries, namely equations (20), (21), (23). These were found to work extremely effectively, and enabled the computational task to be significantly reduced through the use of very modest values of $X_{\max}$.

The physical constraints on $\mathcal{L}_d$ when $Q = 0$ and $\mathcal{L}_c$ when $Q = Q_{\max}$, imply that the Q-derivative terms vanish on these boundaries and equations (12) and (13) may be solved explicitly (to within an arbitrary constant). The procedure begins with an arbitrary explicit solution of (13) in the region $-X_{\max} \leq X < 0$ (where $X_{\max}$ was taken to be suitably large, typically 10). The effect of grid spacing was routinely checked to confirm the integrity of the results. The solution was extended upwards in Q towards $Q = Q_{\max}$, with (14) imposed at $X = 0$ and (20) and (21) at $X = -X_{\max}$. Along each line of grid points $Q = \text{constant} > 0$, the resulting algebraic system (for $-X_{\max} \leq X \leq 0$), may be written in tridiagonal form, and hence routinely solved (using Gaussian elimination, for example). Upon reaching $Q = Q_{\max}$, the procedure switches to the $X > 0$ zone. Along this line, the explicit solution of (12) provides starting data for a Crank-Nicolson procedure that marches downwards towards $Q = 0$, in an analogous manner to than adopted for the upwards march in $X < 0$. For the $X = 0$ condition, (15) was used, together with (21) at the truncated domain location $X = +X_{\max}$). The procedure was then repeated, marching upwards in Q for $X < 0$ and downwards in Q for $X > 0$ until the maximum change in values between successive iterates was less than some prescribed tolerance (typically 1×10^{-6}). Overall, the adopted procedure is analogous to that of [11].

IV. DISCUSSION OF TYPICAL SOLUTION DYNAMICS

A. Comparison of Monte Carlo and PDE results

In order to illustrate the developed methodology we analyse a case with some 10GW of aggregate wind power capacity. It is assumed that the reserved requirements will be driven by the assumption time horizons longer than say 3-4 hours will be managed by additional units (which is well within the dynamic capabilities of modern gas-fired technologies). Over that time horizon, the unpredicted changes in wind output could be about 25% to 30% of the installed wind capacity [2], [9], and a value of volatility of $\sigma = 113$ per annum$^{\frac{1}{2}}$ (this was converted for the computations into hours$^{\frac{1}{2}}$) was generally adopted for the purpose of this paper; a sensitivity analysis on σ is presented later in this section. We first consider simulations made in the time domain. Figure 2 shows the time development of the mean NPV for successively lengthening time horizons up to 200 years, for 11 choices of initial X_0 (as indicated), all with $Q_0 = Q_{\max} = 5$. $r = 0.04$ (per annum), $X_d = X_c = 1\text{GW}$, $\lambda_c = \lambda_d = 1\text{h}^{-1}$, $k = 0.7$. All these computations involved 2000 independent path simulations, each with a timestep of approximately 315 seconds (0.0875 hours, i.e. 10^{-5} years) - simulation trials suggested that an acceptable approximation to the fully continuous PDE solution
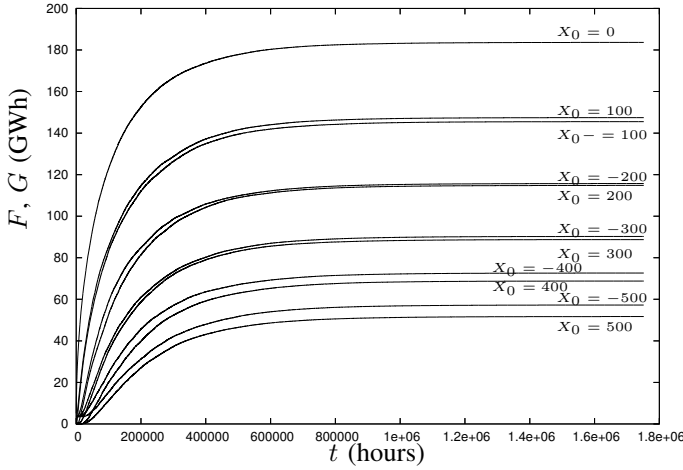
Fig. 2. Simulated income with time t, along $Q_0 = Q_{max} = 5$GWh, $X_c = X_d = 1$GW, $\lambda_c = \lambda_d = 1$h^{-1}, $r = 0.04$ per annum, $\sigma = 113$ per (annum)$^{\frac{1}{2}}$, $k = 0.7$; units of X_0: GW

Fig. 3. Comparison of Monte Carlo results (points) and PDE results, parameters as for figure 2

for these parameter values would require a sample size of several tens of thousands, using a time horizon and sampling interval similar to those above. Computations were rather lengthy - each (X_0) case took approximately 3.5 hours on a 2412MHz AMD Athlon processor to progress the desired 2000 path simulations over 200 years. The figure shows an appreciable gain in NPV over horizons up to about 75 years, and thereafter a slight but detectable increase in NPV. Such qualitative behavior is expected of any income-earning asset, but it appears that the valuations which result when X_0 is farther from $X = 0$ not only have lower limiting NPVs, but take proportionately longer to approach these limits. The reason seems to be that a stochastic storage unit most easily earns income when X is in the neighborhood of $X = 0$, which gives is a high probability of rapid alternation between charging and discharging. Paths for which X_0 is far from $X = 0$ are expected to take a considerable time before first reaching the neighborhood of $X = 0$, but then to linger there for some time. The expected non-discounted value of income per unit of time along a path starting from $X = 0$ declines exponentially, but the expected non-discounted value of income along a path starting far from $X = 0$ (which is dominated by income earned when the path eventually nears $X = 0$) is a delayed, low, but long-lasting pulse - a similar effect is well known from heat diffusion.

Figure 3 shows a comparison of the Monte Carlo simulation results after 200 years (indicated by circles and error bars), as presented in figure 2, and the PDE results, in both cases starting from a full store $Q = Q_{max}$. The error bars for the simulation results indicate the 95% confidence intervals, and strongly encompass the PDE values. We also applied Hotelling's multivariate T Squared test, which confirmed that the 11 means of simulated points did not differ significantly from the PDE solution when tested jointly (detail omitted). Typical PDE calculations take approximately 6 hours (depending on the initial guess) on a 2412MHz AMD Athlon processor (but note this procedure gives the solution across
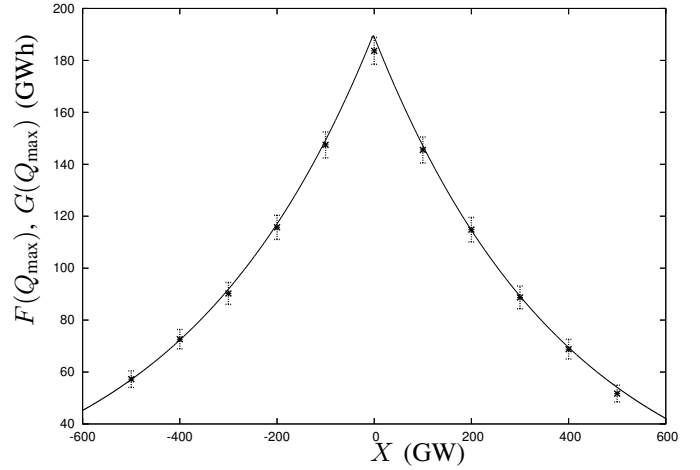
the entire (X, Q) domain, which is some 30,000 points). The agreement between the results obtained from the two quite disparate approaches is excellent (and indicates the high degree of robustness of both solution methods). Figure 3 clearly reveals for both methods the limiting exponential behavior of the 'tails' of the distribution as $|X| \to \infty$ (see the Appendix), along with a sharp peak in the value close to $X = 0$. However it is quite clear that in order to obtain detailed and accurate results the PDE approach is the superior method (by far), since it captures the entire solution domain, at a level of precision and detail which would take in excess of 10^6 hours by simulation.

It should be noted that totally unrealistic values of X have been deliberately taken for the comparison of the two methods, in order to thoroughly validate two, very distinct methodologies. Now that this has been convincingly confirmed, all further results will focus exclusively on practically interesting and viable choices of parameter. Further, the Monte Carlo simulation is impractically slow, so that sensitivity analysis is only feasible by using the PDE approach; results follow.

B. Sensitivity to variations in volatility, σ

Figure 4 shows a magnified view of the figure 3 results along $Q = Q_{max}$ (together with the results along $Q = 0$), close to $X = 0$. This reveals a generic feature of the results, namely, a maximum value at (slightly) negative values of X along $Q = Q_{max}$, whilst along $Q = 0$ the maximum value occurs at slightly positive values of X. Clearly if the system starts full, the most profitable initial state is discharging. If, on the other hand the store starts empty, the most profitable initial state is to be charging - the usefulness of the PDE solution is that it tells us the optimal quantity in either case. An interesting feature is that that whether the store is full or empty, value is maximized when the initial X saturates the power rating by a considerable margin e.g. in the discharge region the value-maximising X is near -3GW, which greatly violates the discharge power rating $X_d = -1$GW. An explanation of this property, which seems to be generic, is that a random walk which starts at the larger
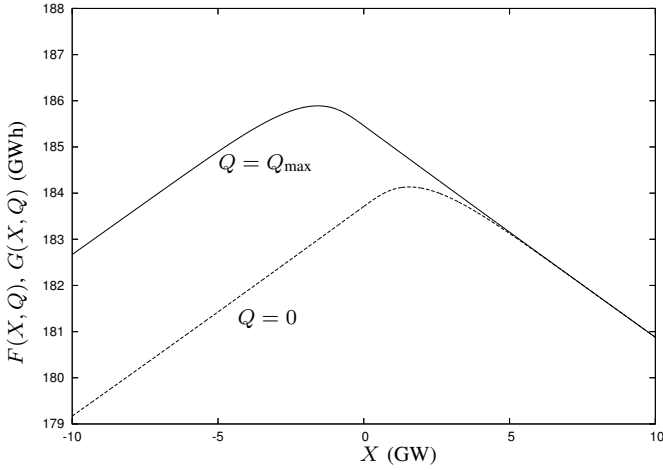
Fig. 4. Results along $Q = 0$ and $Q = Q_{\max}$ close to $X = 0$; parameter details as for figure 3
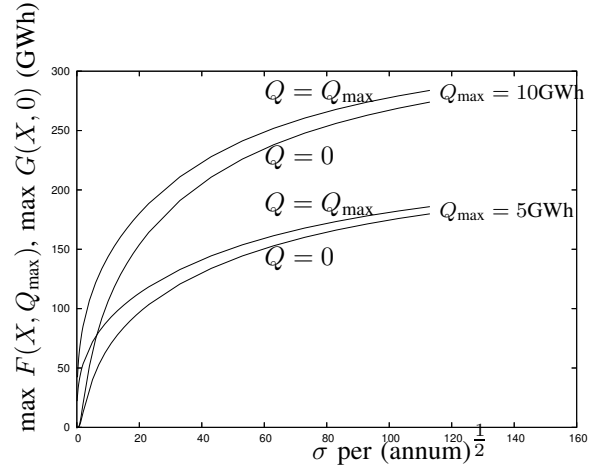

Fig. 5. Results for the maximum values of F, G along $Q = Q_{\max}$ and $Q = 0$ respectively, for varying volatility, $Q_{\max} = 5, 10$; other parameters as for figure 3

deficit of $X = -3\text{GW}$ has the same initial discharge rate of -1GW, but optimally increases the delay before firstly, the discharge rate falls below -1GW and secondly the sign of X changes to begin recharging.

For the parameters used here, and over a wide range of X from $-600GW$ to $600GW$, the total value of the system is dominated by movement in the X dimension; intuitively, total supply must be close enough to total demand, even when storage is used. However in the region of maximum value, where we want the system to operate (near $X = 0$), the value maximizing sign of X is reversed, depending on whether the storage is full or empty ($Q = Q_{\max}$ or $Q = 0$). Hence optimum decision rules are functions of both X (rate) and Q (level). Visual inspection of figure 4 suggests that between $X = -5$ and $X = 5$, control in X and control in Q are roughly equally important in maximizing system value.

A further readily observable, and indeed generic feature of the PDE solution, is that as positive X increases, the results become virtually independent of Q. This is consistent with the observation found by the authors that the asymptotic ($|X| \to \infty$) results, as described in the Appendix, are achieved at remarkably modest values of X (taking just the $n = 0$ term in each series, corresponding to the slowest decaying component of the solution as $|X| \to \infty$, with the corresponding coefficient (A_0, D_0) 'measured' from the numerical results of the PDE).

Inspection of figure 4 reveals that the region of dynamic and economic interest is approximately $-5\text{GW} < X < 5\text{GW}$. Only inside this range is there an appreciable chance of gaining sustained high income in the near future, through alternating between charge and discharge, due to sign reversals of X. Nonetheless the income over this range of interest is remarkably insensitive to X in this range, and equally, perhaps surprisingly (but also confirmed in the Appendix), in the discharging zone, the difference between the fully discharged ($Q = 0$) and fully charged ($Q = Q_{\max}$) states is generally $kQ_{\max}$.

The behavior shown in figure 4 was found to be canonical in shape, only the magnitudes of F and G were found to vary with changes in the various system parameters. Figure 5
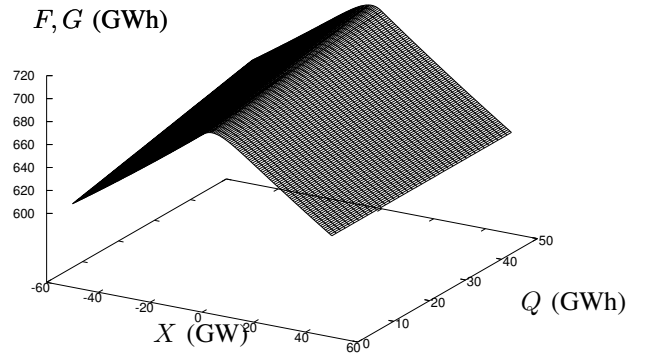

Fig. 6. 3D plot of PDE results, details as for figure 3, except $Q_{\max} = 50$ GWh

shows how volatility variations affect the peak values along both $Q = 0$ and $Q = Q_{\max}$, for both $Q_{\max} = 5\text{GWh}$ and $Q_{\max} = 10\text{GWh}$. Each such pair of points, for a given volatility and $Q_{\max}$, proxies the complete solution domain for the given values of error volatility and storage capacity. This suggests diminishing returns to both volatility and capacity. For example, when compared with the base case value ($\sigma = 113$), the cases $\sigma = 50$ and 10 only reduce the maximum system values to respectively about 81% and 50% of the base case value. These results imply that for any volatility, much of the system's value is generated by small and brief excursions of X around $X = 0$. Hence a storage system could have significant capital value even if the levels of wind forecasting error are fairly low - a result which conflicts with industry intuition.

C. Variations in capacity, Q_{max}

In figure 6 we present a 3D plot of the (PDE) solution across the entire (X, Q) domain for the very large capacity $Q_{\max} = 50\text{GWh}$ (other details as for figure 3). This type of picture is quite universal - note in particular the positive slope
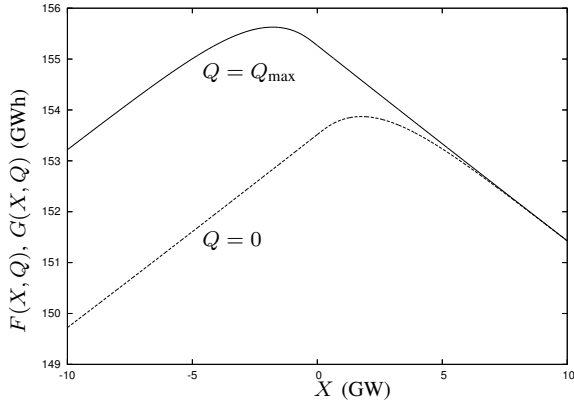
Fig. 7. Results for $X_c = X_d = 0.5$, other parameters as for figure 3

in the Q direction for $X < 0$ and the zero slope for $Q > 0$. They are relatively easily observed in this case, due to the largeness of Q_{max}. However very large excursions of X indeed are required to approach the boundary conditions, where the G solution for very large X tends to the flat plane $G(X, Q) = 0$, and the solution for large negative X tends to the sloping plane $F(X, Q) = kQ$. These (and other) results indicate that again there are diminishing returns to increases in capacity, implying that most of the system's value is generated by cumulatively small and therefore perhaps also short-lived, excursions about $X = 0$.

D. Sensitivity to variations in charging/discharging power rating

The maximum rates at which electricity can be charged/discharged are set by parameters X_c and X_d; the effect of varying these parameters in a correlated manner is shown in figure 7, depicting results for $X_c = X_d = 0.5$ GW. Comparing these results with the base results (figure 4), we see that reducing the power ratings (X_c, X_d) has remarkably little effect on the value of the storage: 50% of the base case power rating gives 84% of the system's value. This suggests again that small excursions around $X = 0$ (slow in- and out-flows, which presumably alternate rapidly) contribute a large fraction of the system's total value. This has implications for storage system design.

V. CONCLUSIONS

In this paper we developed a novel technique to assess the value of a generic, fully flexible, energy storage technology when used to absorb random fluctuation of wind power in real time. Based on standard techniques from financial and real option modeling, we have suggested a new class of PDE for modeling stochastic storage systems, together with a numerical method for solving it.

The innovations are to treat the stochastic storage system as a perpetual option (thus eliminating time from the model when realistic to do so), to treat it as an Asian option (thus eliminating the future level of inventory as an explicit variable in the model), and to regard the system as an instrument whose only source of value is a stochastic flow of income,

which arises from the successive expiry of an infinite sequence of mutually interacting options to charge and discharge. We have shown that there is close agreement between the PDE solution and simulation results, but the PDE numerical solution is several orders of magnitude faster than simulation. The paper contributes a systematic framework for applying PDEs to many problems of stochastic storage which are analytically intractable, and which could not be solved within a practically acceptable time by simulation methods.

Our results indicate that there are diminishing returns to increases in capacity of such a storage system, implying that most of the system's value is generated by cumulatively small and therefore short-lived excursions. This is important for the design of storage systems for managing wind-power fluctuations. We have illustrated the method for purely physical storage systems, subject to linear Brownian motion of only one physical variable. The method can clearly be generalized to systems with multiple physical and/or price variables, and with considerably more complex stochastic and physical dynamics than we have assumed. If we use the present model to value a stochastic storage system at several fossil fuel prices, the result would be the payoff function for a real call option to build such a unit. Conventional real option methods could then be used to calculate how high the fossil fuel price has to rise to justify building the unit.

Although the presented model is mathematically fundamental, an important limitation of the results which we presented here is that we have modeled and solved only the value of the uncontrolled state of a stochastic storage system. Future progress towards optimal control will fall into two main areas, stochastic and economic. In stochastic methods, it will be necessary to generalize some existing regimes of optimal stochastic control. It will be necessary to specify regimes which are simultaneously optimal for a stochastic rate variable $\mathcal{L}_c$ (a modified function of the Brownian motion X variable) and a stochastic level variable Q (a heavily modified form of the integral of $\mathcal{L}_c$). In economics, it will be necessary to adapt more detailed economics of a particular generation mix and more realistic modeling of storage. For example, in order to capture the cost performance of fossil fuel generation it will be necessary to model the costs and time delays involved in warming up the units, the fuel penalties of operating part loaded, the existence of minimum stable generation etc. Furthermore, more realistic modeling of electricity storage (and many other possible applications of stochastic storage) may require generalizations of the stochastic and physical dynamics of our present model.

APPENDIX

APPENDIX : DETAILS OF $|X| \rightarrow \infty$ BEHAVIOR OF THE SOLUTION

In the limit $X \rightarrow -\infty$, $\mathcal{L}_d = -\min\{X_d, \lambda_d Q\}$ and governing equation (13) takes one of two forms, depending

on the value of Q:

$$\frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial X^2} - \lambda_d Q\frac{\partial F}{\partial Q} - rF + k\lambda_d Q = 0, \qquad Q < X_d/\lambda_d, \tag{16}$$

$$\frac{1}{2}\sigma^2\frac{\partial^2 F}{\partial X^2} - X_d\frac{\partial F}{\partial Q} - rF + kX_d = 0, \qquad Q > X_d/\lambda_d. \tag{17}$$

Arbitrage arguments lead us to insist that F is continuous along $Q = X_d/\lambda_d$.

For $Q < X_d/\lambda_d$, the governing equation is (16) and we are then able to write the solution as a Taylor series expansion about $Q = 0$,

$$F^- = \sum_{n=0}^{\infty} Q^n \mathcal{F}_n^-(X) = \frac{k\lambda_d Q}{\lambda_d + r} + \sum_{n=0}^{\infty} A_n Q^n e^{\sqrt{\frac{2(n\lambda_d+r)}{\sigma^2}}X}, \tag{18}$$

where the A_n are constants.

In order to properly connect with (18) at $Q = X_d/\lambda_d$, the solution to (17) for $Q > X_d/\lambda_d$ must have the form

$$F = e^{\frac{r}{\lambda_d} - \frac{rQ}{X_d}}\left(\frac{kX_d}{\lambda_d + r} - \frac{kX_d}{r}\right) + \frac{kX_d}{r}$$
$$+ \sum_{n=0}^{\infty} A_n \left(\frac{X_d}{\lambda_d}\right)^n e^{-n + \frac{n\lambda_d Q}{X_d}} e^{\sqrt{\frac{2(n\lambda_d+r)}{\sigma^2}}X}. \tag{19}$$

Note that (18) and (19) indicate that not only is F continuous at $Q = X_d/\lambda_d$, but so too is $\frac{\partial F}{\partial Q}$.

Although the A_n terms in (18) and (19) all decay exponentially as $X \to -\infty$, the decay rates can be quite slow in practice. Neglecting these terms can lead to significant and spurious domain truncation effects. The dominant (i.e. slowest) exponential decay arises from the $n = 0$ terms. Therefore, in order to impose the $X \to -\infty$ conditions at a finite (and reasonable) value of X (say $X = -X_{\max}$), we construct a Robin-type condition that eliminates A_0 and follows directly from (18) and (19), namely

$$\frac{\partial F}{\partial X} - \sqrt{\frac{2r}{\sigma^2}}F = -\frac{kX_d}{r}\sqrt{\frac{2r}{\sigma^2}}\left(e^{\frac{r}{\lambda_d} - \frac{rQ}{X_d}}\left(\frac{r}{\lambda_d+r} - 1\right) + 1\right), \tag{20}$$

$$\frac{\partial F}{\partial X} - \sqrt{\frac{2r}{\sigma^2}}F = -\sqrt{\frac{2r}{\sigma^2}}\frac{k\lambda_d Q}{\lambda_d + r}, \tag{21}$$

for $Q > X_d/\lambda_d$ and $Q < X_d/\lambda_d$ respectively. This is equivalent to taking (just) the $n = 0$ terms in (18) and (19); in practice, computationally this approximation was found to work exceedingly well.

The solution for $X \to \infty$ is a little simpler — the lack of any forcing terms implies that only exponentially decaying (in X) terms are present,

$$G = \sum_{n=0}^{\infty} D_n(Q)e^{-\sqrt{\frac{2(n\lambda_c+r)}{\sigma^2}}X}. \tag{22}$$

Again, the condition as $X \to \infty$ is imposed at a finite distance $X = X_{\max}$ using a (Robin-type) condition that takes into account the dominant exponential decay:

$$\frac{\partial G}{\partial X} + \sqrt{\frac{2r}{\sigma^2}}G = 0. \tag{23}$$

Again, this is tantamount to taking just the $n = 0$ term in (22), a procedure that nonetheless proved to be numerically very accurate. Finally, it is worth noting that the solution becomes predominantly independent of Q as $X \to \infty$.

REFERENCES

[1] R. Doherty and M. O'Malley, "A new approach to quantify reserve demand in systems with significant installed wind capacity." *IEEE Transactions on Power Systems*, vol. 20, pp. 587 – 595, 2005.
[2] G. Strbac, "System cost of additional renewables." Tech. Rep., 2002, report for Department of Trade and Industry (UK), http://www.dti.gov.uk/energy/developep/080scar_report_v2_0.pdf.
[3] D. E. A. GmbH, "Planning of the grid integration of wind energy in germany onshore and offshore up to the year 2020 (dena grid study," Tech. Rep., 2005, http://www.dena.de.
[4] M. Ilic and P. Skantze, "Electric power systems-operation by decision and control: The case revisited." *IEEE Control Sys. Mag.*, vol. 20, pp. 25–39, 2000.
[5] P. Skantze and M. Ilic, *Valuation, Hedging and Speculation in Competitive Electricity Markets: A Fundamental Approach*. Norwell, MA: Kluwer, 2001.
[6] E. Allen and M. Ilic, *Price-Based Commitment Decisions in the Electricity Market*. New York: Springer-Verlag, 1999.
[7] L. Clewlow and C. Strickland, "Energy derivatives: Pricing and risk management." Tech. Rep., 2000, laCima Publications.
[8] M. Denton, A. Palmer, R. Masiello, and P. Skantze, "Managing market risk in energy," *IEEE Transactions on Power Systems*, vol. 18, pp. 494 – 502, 2003.
[9] G. Strbac and M. Black, "The future value of storage in the uk with generator intermittency," Tech. Rep., 2005, report for Department of Trade and Industry (UK), http://www.dti.gov.uk/renewables/publications/pdfs/dgdt100400000.pdf.
[10] P. Wilmott, S. Howison, and J. Dewynne, *The Mathematics of Financial Derivatives*. Cambridge University Press., 1995.
[11] P. W. Duck, T. W. Marshall, and E. Watson, "First-passage times for the uhlenbeck-ornstein process." *J. Phys. A. (Maths and Gen.)*, vol. 19, pp. 3545–3558, 1986.