Roem-05.qxd 18/11/02 6:12 PM Page 1

5

Applying Stochastic Dynamic Programming to the Valuation of Gas Storage and Generation Assets

Tom Weston

Dynegy

The valuation of energy assets and the contracts structured around them must implicitly involve the flexibility and physical constraints of their operation. Such problems are often best described by recursive treatments, such as the theory of stochastic dynamic programming. To this end, we give an introductory account of stochastic dynamic programming, giving examples relevant to the energy industry, and consider in detail the problems involved in the numerical solution of such examples. We focus on two particular classes of problem: the valuation of natural gas storage, and that of power generation assets.

INTRODUCTION

The theory of pricing financial derivatives, which has proved so successful over the past three decades, has been strongly shaped by certain features of the markets in which the methods have been applied. Some of the most important features of the theory involve the liquidity of underlying markets used for hedging, but as important for the nature of the mathematics involved is the fact that many of the contracts have relatively simple payoff structures. Even exotic options, which provide revenue based upon the average, maximum or standard deviation – or other functions of an underlying quantity – often provide a succinct formula for the value of a derivative at a particular point in time. Although, to then derive the general expression for the value of a contract under different conditions is a taxing problem, this fact almost inevitably leads to a formalism based upon partial differential equations. In contrast, the theory of real options, particularly in the energy markets, provides many examples where sequential decisions, each contingent on the last, must be made correctly to derive the maximum benefit from specific contracts or physical assets. For example, historically, natural gas has often been sold by upstream market participants in the form of "swing" contracts, whereby the flexibility inherent in the process of producing gas is sold on by structuring a contract to allow the buyer to adjust the volumes of gas taken, provided the aggregate volumes over a period of time are above a specified minimum value. This contract evidently involves decisions made sequentially, which are contingent upon previous decisions. For instance, if not enough gas is taken earlier on, the flexibility of the contract is lost as the buyer is then obliged to take sufficient volume to satisfy the minimum requirement. Consideration of the mathematics of this type of problem shows that the structure for the calculation is most naturally described recursively. Such recursive problems are often described in the applied mathematics literature as "dynamic programming", and it is a testament to the deregulation of various industries that the literature originally devoted to problems of minimisation of cost is now being altered to also consider the maximisation of profit. To this end, the purpose of this chapter is to introduce the recursive concepts of dynamic programming with regard to the maximisation of the cashflow to be extracted from an energy asset, where the underlying commodity is tradable with an uncertain price. We first introduce the theory of dynamic programming with the case of known prices. We then discuss the natural extension to stochastic prices, and develop a general formalism that is consistent with popular models for commodity spot prices, addressing the numerical solution of the problem. Then we look at the application of this formalism to two important classes of energy real options – natural gas storage and power generation.

INTRODUCTION TO DYNAMIC PROGRAMMING

Before beginning our exposition of the techniques of dynamic programming, it is necessary to describe the notation that will be used in this chapter. In what follows, the term "asset" will be taken to mean either an item of physical equipment, or a contract that enables either physical production or trading of an energy commodity. The term will also be used synonymously with "system". Also, we encounter variables that will fall under two broad categories.

- *State variables* (eg, x_t) – these variables will determine the precise state of our system as a function of time. For the most part there will be an obvious choice for them in the systems we deal with, eg, storage levels in inventory problems.
- *Control variables* (eg, a_t) – these variables are those that are available to us to change over time to achieve our optimisation, and thus represent the optionality inherent in our asset. Again, for many of the problems considered there will be a natural choice that will coincide with some literal physical control parameter.

The commodity price, which we generically denote S_t, will prove to be a key variable in our analysis. Depending on different approaches, it may be treated either as a known deterministic function of time, or alternatively as a full stochastic process.

We shall now develop the basic dynamic programming formalism. The combination, at each point in time, of system state and action taken (ie, choice of control variable) yields a cashflow, which we write as the function:

$$c_t = K(x_t, a_t, S_t)$$

Our objective is to maximise a weighted sum of the cashflows over a particular period of time, ie, to find the choice of control variables that maximises the quantity:

$$c_t + \beta c_{t+1} + \beta^2 c_{c+2} + \cdots + \beta^{T-t} c_T = \sum_{m=t}^{T} \beta^{m-t} c_m$$

where $\beta \in [0,1]$ is a discount factor[1], subject to the physical constraints of the system. In order to better specify these constraints we introduce an *equation of motion*, which describes the evolution of the system state given a choice of control variables:

$$x_{t+1} = L(x_t, a_t)$$

We now define the *value function* as:

$$V_t(x_t) = \max_{a_t, a_{t+1}, \ldots, a_T} \left(c_t + \beta c_{t+1} + \beta^2 c_{c+2} + \cdots + \beta^{T-t} c_T \right)$$

that is, the discounted sum of cashflow corresponding to the optimal choice of control parameters when starting from a particular state over a particular time interval. This function is a useful object with which to describe an important property of the optimal parameter choice.

The situation, as we have so far described it, is represented in Figure 1. Each edge of the graph represents a possible choice of control parameter, and each vertex an attainable state of the system. Note there is no restriction on whether or not the edges recombine. The path through the graph shown in bold, starting at state A, is the optimal path, in the sense of maximising cashflow with respect to its starting point of A. Now consider the optimum path starting at point B. Since the cashflows are additive along each stage of the path, if any improvement of aggregate cashflow were possible by altering the path to take into account the new starting point of B, then these improvements may also be made to the original path, suggesting it was not originally optimal. This simple argument is the basis of Bellman's celebrated *optimality principle* (Bellman, 1957)[2]. This simple notion has profound consequences for the calculation of

Figure 1 Representation of the changes of state associated with different choices of control parameters. The bold path is optimal

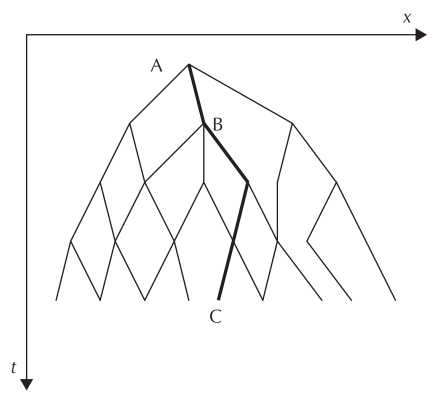

optimal parameters. As a result, our expression for the value faction may be simplified as:

$$\begin{aligned}
&\max_{a_t, a_{t+1}, \ldots, a_T} \left(c_t + \beta c_{t+1} + \beta^2 c_{c+2} + \cdots + \beta^{T-t} c_T \right) \\
&\quad = \max_{a_t} \left(c_t + \max_{a_{t+1}, \ldots, a_T} \left(\beta c_{t+1} + \beta^2 c_{c+2} + \cdots + \beta^{T-t} c_T \right) \right)
\end{aligned}$$

where it is understood that the second maximisation incorporates the state achieved by the first choice of parameter. This simplification arises directly from Bellman's principle: the maximisation from $t + 1$ onwards is independent of choices previous to that time, and only depends upon its initial state. The above expression may be made more precise by using our expressions for cashflow and equation of motion:

$$\begin{aligned}
V_t(x_t) &= \max_{a_t} \left(K(x_t, a_t, S_t) + \beta \max_{a_{t+1}, \ldots, a_T} \left(c_{t+1}(L(x_t, a_t), a_{t+1}) + \beta c_{c+2} + \cdots + \beta^{T-t-1} c_T \right) \right. \\
&= \max_{a_t} \left(K(x_t, a_t, S_t) + \beta V_{t+1}(L(x_t, a_t)) \right) \qquad \textbf{(1)}
\end{aligned}$$

This recursive equation for the value function is known as the *Bellman equation* for the problem and represents a full statement of the optimisation problem.

It might be thought on first glance that the recursive nature of the Bellman equation may make it difficult to solve. However, there are several situations where the recursive structure of the equations will enable a simple solution method. The simplest of these is where there is a known value of the system at some point in the future. This will frequently be the case with fixed term contracts, such as contracts for use of gas storage facilities or tolling agreements for power stations. Beyond the term of these agreements the value is clearly zero, as no cashflows accrue from the contracts. Generalising this case to the situation where the known future value is an arbitrary function of system state, we have:

$$V_T(x_T) = g(x_T)$$

where g is an arbitrary function. In such a case, this expression may be substituted into the right-hand side of the Bellman equation, which may then be solved for $V_{T-1}(x_{T-1})$. This process may be iterated to find the value function at any time. In general, such methods are described as *backwards induction*.

For situations where no such terminal value is known, a simple procedure to solve the Bellman equations is as follows. For a discount parameter β less than unity, the current value of a system state becomes insensitive to the state of the system when considered far enough in the future. Therefore, a simple computational procedure is to adopt an arbitrary (but sensible) final state at some future date far enough away in time, and solve with backwards recursion from this value, checking for the required insensitivity.

An elementary dispatch decision

Having introduced dynamic programming through deterministic prices, let us now examine an elementary example that illustrates the principles. Consider the following extremely simplified scenario. The manager of a power station must operate the plant over the coming week. After exactly one week the plant must be switched off for maintenance. The manager examines available prices and obtains margins (ie, revenue less variable operating costs) for running the plant for each hour over the next week.

The margins are listed in Table 1 and also are plotted out in Figure 2. These are based upon typical profiles of power spot prices and show that the plant is a true "peaker" – it is typically only economically viable to run at times of peak demand (and price). To simplify matters further it is assumed the plant is extremely flexible and may be switched off or back to full capacity rapidly. However, in switching the plant back on a fixed cost is incurred, here taken to be €12/MW.

The task is to optimise the timing of the decision to switch on or off in order to maximise the cashflow. Using the terminology outlined above, the obvious choice of state variable to denote whether the plant is switched on or off is $x \in \{0, 1\}$. With this choice we must choose $a \in \{-1, 0, 1\}$ to denote our choice to switch off, remain in the current state, or switch on, in which case we have $x_{t+1} = L(x_{x,}a_t) = x_x + a_t$. We obviously face the problem of how to enforce the fact that it is impossible to switch off a plant that is already off, etc. Either this logic must be incorporated into the equation of motion, written above, or an alternative computational method is to

Table 1 Example margins, in €/MWh for each hour for operating a hypothetical power plant over the period of one week

	Mon	Tue	Wed	Thu	Fri	Sat	Sun
1	−3.21	−2.01	−4.43	−6.21	−5.84	−5.51	−4.43
2	−2.75	−1.51	−4.09	−5.79	−5.48	−5.19	−4.15
3	−2.68	−1.43	−4.03	−5.72	−5.41	−4.98	−3.96
4	−2.88	−1.65	−4.18	−5.91	−5.58	−5.48	−4.4
5	−3.1	−1.89	−4.34	−6.1	−5.75	−5.67	−4.57
6	−3.44	−2.27	−4.6	−6.42	−6.03	−5.81	−4.69
7	−1.83	−0.49	−3.4	−4.95	−4.73	−5.45	−4.38
8	0.57	1.65	−2.82	2.15	−0.35	−2.47	−0.99
9	4.23	5.5	−0.52	7.22	3.46	−1.61	−0.15
10	2.66	3.85	−1.51	5.04	1.82	0.28	1.7
11	1.66	2.79	−2.14	3.65	0.78	1.39	2.79
12	1.92	3.07	−1.97	4.02	1.05	1.39	2.79
13	3.14	4.35	−1.21	5.7	2.32	1.46	2.85
14	−0.13	0.91	−3.26	1.18	−1.08	−0.38	1.05
15	−0.67	0.35	−3.6	0.43	−1.65	−2.56	−1.08
16	−1.52	−0.55	−4.13	−0.75	−2.53	−3.51	−2.01
17	0.63	1.71	−2.79	2.22	−0.3	−2.88	−1.39
18	3.7	4.94	−0.86	6.48	2.9	0.16	1.58
19	6.13	7.5	0.67	9.84	5.43	5.42	6.73
20	13.28	6.39	9.77	4.7	14.37	1.02	3.96
21	5.21	0.92	3.08	−1.21	5.44	−2.39	0.16
22	2.69	−0.79	1	−3.06	2.65	−3.87	−1.48
23	1.33	−1.71	−0.13	−4.06	1.15	−4.51	−2.19
24	−1.09	−3.36	−2.13	−5.83	−1.53	−5.22	−2.98

Figure 2 Margin (in €/MWh) shown over the course of a week, together with x_t calculated using the Bellman equation procedure

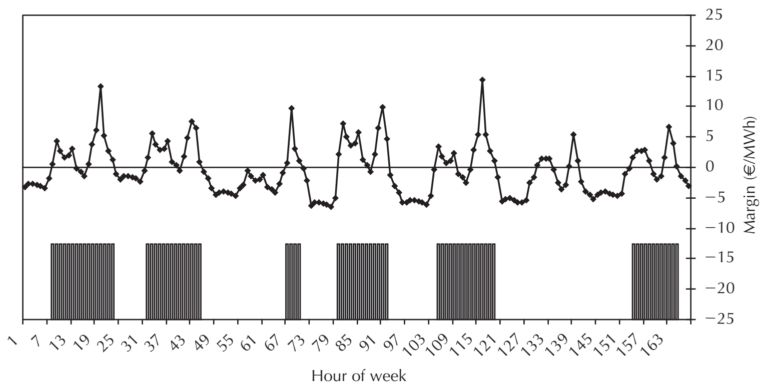

incorporate it into the cashflow function and associate an impossible change of state with an infinite cost. Adopting this latter method:

$$c_t = K(x_t, a_t, m_t) = \begin{Bmatrix} -k_\infty & (x_t = 0,\ a_t = -1) \\ 0 & (x_t = 0,\ a_t = 0) \\ m_t - c_s & (x_t = 0,\ a_t = 1) \\ 0 & (x_t = 1,\ a_t = -1) \\ m_t & (x_t = 1,\ a_t = 0) \\ -k_\infty & (x_t = 1,\ a_t = 1) \end{Bmatrix}$$

where m_t is the margin derived from running in hour t, c_t is the cost of starting the plant, and k_∞ is an infinite cost. Applying these choices to the Bellman equation for the problem, which takes the form (we do not bother with discounting, when considering cashflows over a week):

$$V_t(x_t) = \max_{a_t}(K(x_t, a_t, m_t) + V_{t+1}(L(x_t, a_t))$$

subject to the final condition

$$V_{T+1}(x_{T+1}) = \begin{Bmatrix} 0 & (x_{T+1} = 0) \\ -k_\infty & (x_{T=1} = 1) \end{Bmatrix}$$

provides a simple calculation which can be replicated, for instance, on a spreadsheet.

The results of applying this procedure are shown in Figure 2 in terms of the hours when the plant is switched on. It might be thought that the results of the Bellman equation approach could be encapsulated by a simple rule, for instance, to switch off for any period of consecutive negative margin whose cumulative loss is greater that the start cost. However, the results of the Bellman equation produce a total margin which is over 10% larger than the result of this simple strategy. This difference derives from the asymetrical treatment of the decision to switch on or off in the heuristic treatment: consider the extreme instance where only one hour provides positive margin and this is less than the start cost. The simple argument would suggest switching off twice, before and after this hour, whereas the optimal decision is to never switch on at all.

STOCHASTIC DYNAMIC PROGRAMMING UNDER A ONE FACTOR STOCHASTIC PROCESS

Thus far we have derived the Bellman equation and discussed its solution for systems where all information is known into the future (deterministic dynamic programming). We now turn to the situation of greatest interest in the modelling of energy assets – where the unknown element of the problem derives from the commodity price satisfying a stochastic process. Here we shall take this to mean a stochastic differential equation of the form:

$$\frac{\mathrm{d}S_t}{S_t} = \eta \mathrm{f}(S_t, t)\mathrm{d}t + \sigma(t)\mathrm{d}W_t \tag{2}$$

where $\mathrm{d}W_t$ is a standard Wiener process and $\mathrm{f}(S_t, t)$ is a function that will allow for mean reversion to some (possibly time-dependent) value for the price. A popular choice for this function, which we make considerable use of, is $\ln(\bar{S} / S)$. This leads to the single factor Schwartz model (Schwartz, 1997), for which many analytical results are available. Some of this model's properties are described in the appendix to this chapter.

The main difference between the methods we shall develop to describe dynamic programming under the process above and those in the case of deterministic prices is the way in which we must treat the price itself. Previously we distinguished between state variables and control variables, but the price entered the problem merely as a parameter in the cashflow obtained at each time step. Now, we shall need to promote the price to the set of state variables. The most important reason for this is that consideration of prices in relation to the properties of the distribution of prices implied for them by the above process, leads to opportunities that will affect our view of the value of the system. For instance, consider a storage problem. We have a view of the distribution of prices at the next instant in time. If the price we actually observe is less than our expected mean price, then this suggests there is an enhanced possibility of profit by buying the commodity and storing it for later sale. Thus, in assessing the value at any point in time, we must directly include the currently observed price as definitive of our state. Acknowledging this fact leads us to view our state as being defined by the quantities x_t and S_t.

Another difference between the deterministic and stochastic problems is the statement of cashflow maximisation. Since future prices are now uncertain, we can only hope to maximise *expected* cashflows. However, we must be careful in how we define this expectation. Since we have suggested that our actions will depend upon the comparison of currently observed prices with our expectations, and also that the process for price itself conditions future expectations on the current value, we must carefully use conditional probabilities for our expression of cashflow. The expression for the expectation of cashflow which we shall seek to maximise is:

$$c_t + \beta E\left(c_{t+1} + \beta^2\left(c_{t+2} + \cdots + \beta^{T-t} E\left(c_T \left| S_T x_T\right.\right) \cdots \left| S_{t+2} x_{t+2}\right.\right) S_{t+1} x_{t+1}\right)$$

that is, we consider the expected cashflows in the future as being related to all the information that will become available up to the point considered. We shall consider the expected values to be calculated with respect to absolute probabilities in what follows, but the formalism could equally well apply to risk-neutral probabilities. We may now proceed to a Bellman equation in the equivalent manner; we define a value function as the maximal value of the above expression obtained by varying the control parameters, ie,

$$\begin{aligned} V_t(x_t, S_t) = \max_{a_t, a_{t+1}, \ldots, a_T} & \left(c_t(S_t, x_t, a_t)\right. \\ & \left. + \beta E\left(c_{t+1} + \beta^2 E\left(c_{t+1} + \cdots + \beta^{T-t} E\left(c_T \left| x_t S_T\right.\right) \cdots \left| x_{t+1} S_{t+1}\right.\right) \left| x_t S_t\right.\right)\right) \end{aligned}$$

The previous arguments that led to the optimality principle apply here also, when we consider the expected value with our expanded definition of system state. Hence we may write the Bellman equation directly as:

$$V_t(x_t, S_t) = \max_{a_t}\left(c_t + \beta E\left(V_{t+1}\left(x_{t+1}, S_{t+1} \left| x_t S_t\right.\right)\right)\right)$$

or, more explicitly using the equation of motion and cashflow functions:

$$V_t(x_t, S_t) = \max_{a_t}\left(K(x_t, a_t, S_t) + \beta E\left(V_{t+1}\left(L(x_t, a_t), S_{t+1}\right) \left| S_t\right.\right)\right)$$

This will be the key result for our stochastic dynamic programming analysis of energy assets. Before using it in detail we apply it in a simple illustrative analytic example.

A simplified storage model

Consider a storage problem defined at three time steps, $t = 1$, 2, or 3. At each step we must decide how much to add to our inventory or, alternatively, how much to remove. Choosing our state variable, x, as the amount in store and our control variable as the amount removed or added, then our equation of motion is simply:

$$x_{t+1} = x_t + a_t$$

and the cashflow function becomes

$$c_t = -a_t S_t$$

where we have taken the convention that $a > 0$ corresponds to increasing the amount in store. (The cashflows correspond only to purchasing the commodity to place in store, or selling that removed from store. It would be a simple matter to add extra costs.) We place constraints on both the amount of storage available, and also on the maximum amount that may be removed or added at each time step. Hence:

$$x_{\min} \le x_{t+1} \le x_{\max}$$

and

$$a_{\min} \le a_t \le a_{\max}$$

By identifying the bounding constraint, these two inequalities may be combined with the equation of motion to yield:

$$\max(a_{\min}, x_{\min} - x_t) \le a_t \le \min(a_{\max}, x_{\max} - x_t)$$

To simplify the example further we will assume the prices are not conditionally dependent; we choose them to be independent and identically distributed with a simple density function, $p(S)$, which we will specify later.

We assume further that any of the commodity left in store after $t = 3$ is lost. This boundary condition may be written to coincide with the Bellman equation formalism by extending the set of time steps to include a fourth and noting that no cashflows accrue after the third time step, thus:

$$V_4(x, S) = 0$$

We now proceed to solve the problem through backwards induction, recursively using the Bellman equation, equation (1), three times:

$$\begin{aligned} V_3(x, S) &= \max_{a_3}[-a_3 S + \beta E(V_4(x + a_3, S))] \\ &= \max_{a_3}[-a_3 S] \\ &= -\max(a_{\min}, x_{\min} - x)S \end{aligned}$$

This result is intuitive; we remove as much as possible on the final step and sell it. Applying backwards induction again yields:

$$\begin{aligned} V_2(x, S) &= \max_{a_2}[-a_2 S + \beta E(V_3(x + a_2, S))] \\ &= \max_{a_2}[-a_2 S + \beta E(-\max(a_{\min}, x_{\min} - x - a_2)S)] \\ &= \max_{a_2}[-a_2 S - \beta \bar{S} \max(a_{\min}, x_{\min} - x - a_2)] \end{aligned}$$

where $\bar{S}$ is the mean price. Now, the value of a_2 that maximises value is not so obvious in the event of a large negative value for $a_{\min}$ and, if $S < \bar{S}$ the value is maximised by a_2 taking on its largest permitted value, and vice versa. For intermediate values of $a_{\min}$ a more careful inspection is required.

So as not to make the notation too cumbersome, in the calculation of the final value function we assume that $x_{\min} - x > a_{\min}$ and $x_{\max} - x < a_{\max}$. In this case we find:

$$V_2(x, S) = \begin{Bmatrix} (x - x_{\min})S & (\beta\bar{S} - S < 0) \\ (x_{\max} - x_{\min})\beta\bar{S} + S(x - x_{\max}) & (\beta\bar{S} - S > 0) \end{Bmatrix}$$

In performing the final backwards induction step we shall require the expectation value of the expression for $V_2(x, S)$. For simplicity we take the density function of the distribution to be of lognormal form:

$$p(S) = \frac{1}{\sqrt{2\pi}} \frac{1}{S\sigma} \exp\left[-\frac{(\ln S / \bar{S} + \sigma^2 / 2)^2}{2\sigma^2} \right]$$

Then the expectation value takes the value:

$$E[V_2(x, S)] = (x - x_{\text{min}})\bar{S} + (x_{\text{max}} - x_{\text{min}})\bar{S} \left[\beta N\left(\frac{\ln \beta}{\sigma} + \frac{\sigma}{2} \right) - N\left(\frac{\ln \beta}{\sigma} - \frac{\sigma}{2} \right) \right]$$

Thus the final Bellman equation becomes:

$$\begin{aligned} V_1(x, S) &= \max_{a_1}[-a_1 S + \beta E[V_2(x + a_1, S)]] \\ &= \max_{a_1}[(\beta\bar{S} - S)a_1] + \beta\bar{S}(x - x_{\text{min}}) \\ &\quad + \beta\bar{S}(x_{\text{max}} - x_{\text{min}})\left(\beta N\left(\frac{\ln \beta}{\sigma} + \frac{\sigma}{2} \right) + N\left(\frac{\ln \beta}{\sigma} - \frac{\sigma}{2} \right) \right) \end{aligned}$$

In the first term of this final expression, if $\beta\bar{S} - S > 0$ then evidently the largest permitted value of a_1 is chosen and vice versa. The three terms admit to simple interpretations, the first is due to a straight comparison of the currently observed price against the expected mean, the second is a measure of the intrinsic value of the commodity in store, and the third is a measure of the extrinsic value to be obtained over the three periods. It may be directly observed that this final term is a monotonically increasing function of σ; in general we will find that the value of flexible storage is intimately related to the volatility of the underlying price.

Notwithstanding the previous simplified example, since conditional probabilities are evidently essential to the analysis, we need to describe how to obtain them for our stochastic process. There are several approaches to this problem but, sacrificing elegance, (and possibly computational efficiency) for generality and simplicity, we proceed by outlining a simple grid based approach. The process will prove easier to manipulate if we consider it in terms of the logarithm of the price. Writing $y = \ln S$ and applying Itô's lemma to equation (2) produces the process:

$$\mathrm{d}y = [\eta \mathrm{f}(\exp(y_t), t) - \sigma^2 / 2]\mathrm{d}t + \sigma \mathrm{d}W$$

This process may be written in discrete form as:

$$y_{t+1} = y_t + \eta[\mathrm{f}(\exp(y_t), t) - \sigma^2 / 2]\delta t + \sigma\sqrt{\delta t}\phi_t$$

where the random quantities $\{\phi_t\}$ are independent and drawn from the distribution $N(0, 1)$. In this discrete form, the distribution of y_{t+1} conditional upon y_t can be inferred from the properties of ϕ:

$$\mathrm{prob}\left(y'_{t+1} < y_{t+1}(\phi_t) \middle| y_t\right) = \mathrm{prob}(\phi'_t < \phi_t)$$

Writing this expression in terms of the underlying densities gives us:

$$\int_{-\infty}^{y_{t+1}} p_x\left(y'_{t+1} \middle| y_t\right) \mathrm{d}y'_{t+1} = \int_{-\infty}^{\phi_t} p_\phi(\phi')\mathrm{d}\phi'$$

Differentiating with respect to ϕ_t yields:

$$p_y\left(y_{t+1} \middle| y_t\right) = p_\phi(\phi) / (\mathrm{d}y_{t+1} / \mathrm{d}\phi)$$

Applying this formula to our discrete version of the stochastic process, with the stated Gaussian density function for ϕ, produces:

$$p\left(y_{t+1} \middle| y_t\right) = \frac{1}{\sqrt{2\pi\sigma^2\delta t}} \exp\left[-\frac{(y_{n+1} - y_n - (\eta \mathrm{f}[\exp(y_t)] - \sigma^2 / 2)\delta t)^2}{2\sigma^2\delta t}\right]$$

We shall use this formula for conditional price distributions in the Bellman equation where it is assumed that the time step is small. However, to check the validity of this use, we shall consider the application of the short-term form above to calculate longer duration conditional densities. This will also introduce the grid discretisation technique we shall use.

Since the original stochastic process is implicitly a Markov process, we may construct longer-term conditional density functions from shorter duration ones using the Chapman–Kolmogorov equation (Gardiner, 1985):

$$p\left(y_t \middle| y_0\right) = \int\int\cdots\int p\left(y_t \middle| y_{t-1}\right) p\left(y_{t=1} \middle| y_{t-2}\right) \cdots p\left(y_1 \middle| y_0\right) \mathrm{d}y_{t-1}\mathrm{d}y_{t-2}\cdots\mathrm{d}y_1$$

We now consider each of the individual conditional density functions (which are assumed to be associated with short time intervals) as taking discrete values on a regular grid and approximate each of the integrals by a sum:

$$p\left(y_t^i \middle| y_0^m\right) = \sum_{i,j,k,\ldots,l} p\left(y_t^i \middle| y_{t-1}^j\right) p\left(y_{t-1}^j \middle| y_{t-2}^k\right) \cdots p\left(y_1^l \middle| y_0^m\right) (\delta y)^t$$

where δy is the grid spacing. Effectively this approximation transforms our continuous Markov process into a discrete Markov chain (Cox and Miller, 1965). In order to consider the numerical stability of this method, we note that the above procedure involves successive matrix multiplications, the number of which in the limit $\delta t \to 0$ becomes infinite. In order to prevent divergence, we thus require that the eigenvalues of each matrix have an absolute value less than unity. To describe the consequences of this requirement we shall pursue the Markov chain aspect of our method. In the theory of discrete Markov chains, vectors, which describe probabilities associated with particular states (for us, particular sets of price), are evolved in time by the application of *transition* (or *stochastic*) matrices. From the probabilistic nature of these matrices, they are required to satisfy two properties:

- ❑ each element must be positive, and
- ❑ each row must sum to unity.

An elementary result of the theory of Markov chains (which can be obtained as a simple consequence of the Perron-Frobenius theorem (Grimmett and Stirzaker, 1992)) is that the absolute values of the eigenvalues of a transition matrix are bounded within the range $[0, 1]$. Therefore, if we can demonstrate that our short term conditional densities are at least approximately transition matrices (including a factor of the grid size), we can be assured of stability. From our explicit expression above, evidently each element is positive. The row summation requirement implies:

$$\sum_m p\left(y_{t+1}^m \middle| y_t^n\right) \delta y = \sum_m \frac{\delta y}{\sqrt{2\pi\sigma^2 \delta t}} \exp\left[- \frac{(m\delta y - n\delta y - (\eta f[\exp(y_t^n)] - \sigma^2/2)\delta t)^2}{2\sigma^2 \delta t} \right] \approx 1$$

Noting that the last term in the numerator in the exponential function has a factor of δt (and that the first term does not) it may then

Figure 3 A graph of the function $\frac{\alpha}{\sqrt{2\pi}} \sum_{m=-\infty}^{\infty} \exp(-\alpha^2 m^2/2)$ plotted against α

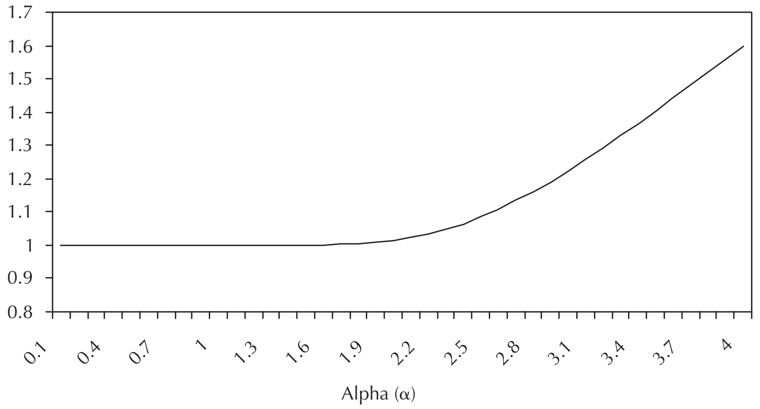

be neglected in the limit of infinitesimal time steps. Also we shall assume that we have chosen a large enough grid, and hence that summation over grid points may be considered infinite (in both directions). These simplifications transform the sum into the form:

$$\sum_{m=-\infty}^{\infty} \frac{\alpha}{\sqrt{2\pi}} \exp(-\alpha^2 m^2 / 2)$$

where we have defined $\alpha = \delta y / \sigma\sqrt{\delta t}$. This quantity may be related to theta functions of the third order (in the notation in Abramowitz and Stegun (1964) it may be written as $(\alpha^2 / \sqrt{2\pi})\vartheta_3(\pi, \exp(-\alpha^2 / 2)))$. We graph it as a function α in Figure 3. Inspection of this plot suggests we will obtain numerical stability if:

$$\alpha = \frac{\delta y}{\sigma\sqrt{\delta t}} < 1$$

This expression relates to the size of time and price grids.

As an example we have applied this method to the Schwartz model, using a set of parameters of similar orders of magnitude to those of the UK gas prices described in the appendix. By an analogy with northern European gas and power price behaviour, we have set the mean price level and volatility to have an elementary

Figure 4 Comparison of approximate conditional probability densities with an analytic expression for the Schwartz single factor model used in simulating price distributions over the course of a year

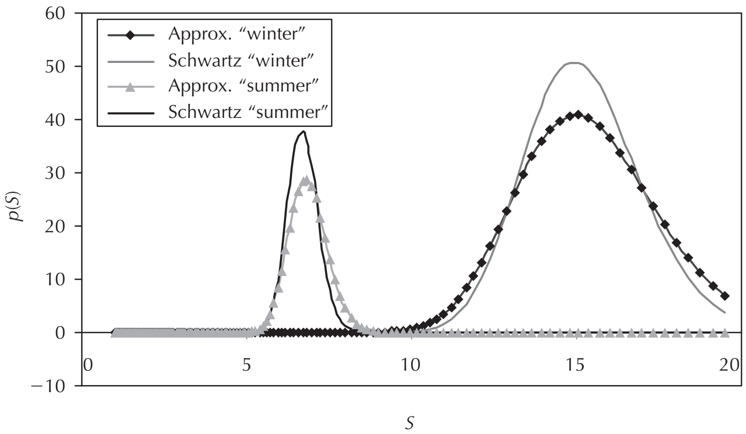

sinusoidal time-dependence over a year, which leads to "summer" (low price and low volatility) and "winter" (high price, high volatility) periods, corresponding to the minimum and maximum of these functions. In Figure 4 we plot probability distributions for both our approximation, iterated over the year and also, using the analytic solution to the Schwartz model (but taking the instantaneous mean and volatility as discussed in the appendix). We can see that, even using a time step of one day, we can obtain approximate solutions that have the correct behaviour. Greater accuracy can be obtained by choosing parameters that reduce the quantity α described above.

The discretisation approach may also be applied to our modelling of state variables in the Bellman equation. Let us consider the state of the system satisfying a set of discrete states $\{x_i\}$, which will also require the control variables to consist of discrete variables. In this case, the value function in the Bellman equation, may be considered to be a matrix of discrete systems states and commodity prices. This form will prove convenient for computational purposes.

Given this discrete approach the Bellman equation is transformed from its continuous form:

$$V_t(x_t, S_t) = \max_{a_t} \left(K(x_t, a_t, S_t) + \beta E \left(V_{t+1}(L(x_t, a_t), S_{t+1}) \middle| S_t \right) \right)$$

to a discrete form

$$V_t\left(x_t^i, S_t^j\right) = \max_{a_t}\left[K\left(x_t^i, a_t, S_t^j\right) + \beta\sum_k p_t\left(S_{t+1}^k \middle| S_t^j\right) V_{t+1}\left(L(x_t, a_t), S_{t+1}^k\right)\delta S\right]$$

(as in the calculation of future probability densities, it will often improve numerical efficiency to use the representation of the logarithm of the price).

ADDITIONAL RESULTS CONCERNING THE METHOD OF DYNAMIC PROGRAMMING

Before moving on to discuss the application of our stochastic dynamic programming formalism, it will prove useful to discuss some additional results suggested by the theory.

Additional constraints

The previous analysis, using stochastic dynamic programming to derive the fair value of an asset by deriving the optimal strategy, gave a useful formalism for considering simple, idealised systems. In the real world of hard energy assets, whose flexibility derives from the underlying engineering of physical processes, there are often many physical constraints, which go beyond a simple single state variable. For instance, the output of power stations is often capped by emissions considerations, or it is often given a floor by "burner tip" take or pay fuel contracts. When switched off, a generation unit may need to remain off for a minimum amount of time before producing power again. Certain tolling (power station lease) agreements may limit the number of times a plant may be started up over a particular time period. In practice there is no simple strategy for capturing all this detail, but in principle each variable specified in this way affects the value observed at each point in time, and therefore must be included in the value function definition. Consequently, each must also be given an equation of motion of its own, and again the state space is increased.

For instance, consider a power station with constrained starts over a year. A natural choice of an extended state space to accommodate this restriction would be to keep the output capacity as one state space, and to include the incremental number of starts of the year as another. Hence the Bellman equation would include this extra variable, y_t:

$$V_t(x,y,S) = \max_{a_t}\left[K(x,a,S) + \beta E\left(V_{t+1}(L(x,a_t), M(y,x,a_t), S')|S\right)\right]$$

where we have included an extra equation of motion to update y_t. This function would be of the form:

$$M(y,x,a) = \begin{cases} y & (x \neq 0 \text{ or } a = 0) \\ y+1 & (x = 0 \text{ and } a > 0) \end{cases}$$

Evidently, the size of the matrices associated with the value function in discrete form increases with the power of the number of state variables. The impact of this fact on the practicality of using dynamic programming for very complex (ie, high dimensional) systems is referred to by some authors as "the curse of dimensionality" (Ljungqvist and Sargent, 2002).

Expected future system state

Having demonstrated the applicability of dynamic programming to derive a value for a particular asset, it is worthwhile considering what other information, in addition to this number, may be extracted. When using dynamic programming it must be noted that the problem of the calculation of the optimum strategy and the corresponding value are simultaneously solved. In this sense, when the Bellman equation, written in the form:

$$V_t(x,S) = \max_{a_t}\left[K(x,a,S) + \beta E\left(V_{t+1}\left(L(a,x), S'|S\right)\right)\right]$$

is recursively solved, we may record the value of $a_t(x, S)$ at each time step. This information shows us directly how to modify the system at any combination of system state, price and time. This is a direct encoding of the optimum strategy and as such contains all the information we require to evaluate the expected value of any relevant quantity. For instance, it is of great use to find the expected value of the system state at some time in the future. To do this, recall that the equation of motion shows the evolution of the system state:

$$x_{t+1} = L(x_t, a_t) = L(x_t, a_t(x_t, S_t))$$

where we have substituted the functional form of a derived from the solution of the Bellman equation. From our discussion of the Chapman–Kolmogorov equations in discrete form, which were applied to the numerical stability of conditional distributions in the previous section, we note that the probability density of the price can be built up independently of the dynamic program as:

$$p_t\left(S_t^i\right) = \sum_{j,k,\ldots,l,m} p_t\left(S_t^i \middle| S_{t-1}^j\right) p_{t-1}\left(S_{t-1}^j \middle| S_{t-2}^k\right) \cdots p_1\left(S_1^l \middle| S_0^m\right) p_0\left(S_0^m\right)$$

where $p_0\left(S_0^m\right)$ is the initial distribution of price (which may be a Kronecker delta function, (Kreyszig, 1999)). We can therefore construct a density representation of price at each time step. Combining this density with the equation above for the propagation of system state, we can now construct a conditional density for the state variable as:

$$p_t\left(x_{t+1}^i \middle| x_t^j\right) = \sum_{x_{t+1}^i - L\left(x_t^j, a_t\left(x_t^j, S_t^k\right)\right)=0} p_t\left(S_t^k\right)$$

where the sum over price probabilities includes only the set which join x_{t+1}^i to x_t^j. Having obtained these conditional probabilities, the density for system state can be directly calculated in a similar way to the price above ie, by successive use of Chapman–Kolmogorov type relations. When we have expressions for both the price density and the system state density at each time, it is a simple matter to evaluate the expected value of $a_t(x_t, S_t)$.

Asset portfolios

The theoretical framework of dynamic programming can sometimes be useful to draw general conclusions about optimisation problems. One situation, which occurs in various forms in energy markets, is where a portfolio of several assets is controlled by a single market participant. As an example, consider the (not uncommon) scenario of a power station that has a fuel supply contract with a swing component, and which also has a certain allowance of emissions credits that it may trade, or use in its own power generation. The fuel swing contract and the emissions allowance may both be considered as contracts with optionality, which will have an optimal strategy in extracting cashflow. A pertinent question is whether the three assets when considered in a portfolio have a different aggregate value than when considered independently. This question can be considered by setting up a dynamic programming formalism for a portfolio containing two systems:

$$V_t\left(x^{(1)}, x^{(2)}, a^{(1)}, a^{(2)}, S^{(1)}, S^{(2)}\right) = \max_{\substack{a_t^{(1)},\ldots,a_T^{(1)} \\ a_t^{(2)},\ldots,a_T^{(2)}}} \left[\sum_{m=t}^{T} \beta^{m-t} K\left(x^{(1)}, x^{(2)}, a^{(1)}, a^{(2)}, S^{(1)}, S^{(2)}\right)\right]$$

where the notation for the two systems is hopefully transparent. From a simple consideration of the problem, it should be evident that it will split into two sub-problems if, and only if we can divide the cashflow function into two independent functions that each have a single control parameter argument:

$$K\left(x^{(1)}, x^{(2)}, a^{(1)}, a^{(2)}, S^{(1)}, S^{(2)}\right) = K^{(1)}\left(x^{(1)}, a^{(1)}, S^{(1)}, S^{(2)}\right) + K^{(2)}\left(x^{(2)}, a^{(2)}, S^{(1)}, S^{(2)}\right)$$

and also that the equations of motion that update each state variable are independent:

$$x_{t+1}^{(1)} = L^{(1)}\left(x_t^{(1)}, a_t^{(1)}\right)$$
$$x_{t+1}^{(2)} = L^{(2)}\left(x_t^{(2)}, a_t^{(2)}\right)$$

Therefore, in the example above, when considered simply, the swing contract, emissions allowance, and the generation part itself each have separate equations of motion and cashflow functions, and thus the valuation of the portfolio is equal to the sum of the values of its parts. However, when there are economies (or diseconomies) of scale and the cashflow functions are not additive, the formalism breaks down. A simple example of this instance would be a fuel contract that is specified as being at the "station gate" ie, a cost is incurred to resend the fuel back to a point where it may be traded.

THE VALUATION OF NATURAL GAS STORAGE ASSETS

Having obtained a formalism for the application of stochastic dynamic programming under a one factor price model by discretising the prices and system states on regular grids, we now apply it to an important example of an energy asset, that of natural gas storage. Gas is an important energy commodity for which the valuation of storage is critical. However, the methods we are going to discuss can readily be applied to pump storage reservoirs, coal stockpiles, and various others examples of storage for energy products.

In common with all energy assets, storage assets come with a large number of real world complexities, but for the purposes of our discussion here we shall use an idealisation with the following characteristics.

- ❏ The storage facility has a finite capacity, defined by the upper bound of x_{space}.
- ❏ Over each time step for which we consider price fluctuations, we may either:
  - ❏ inject an amount x_{inj} of gas, or withdraw an amount x_{with} of gas from store.

In addition to these assumptions, we take the storage to be owned on a lease basis such that, after a time T, any gas left in store is lost. There are frequently occurring complexities, that complicate this simple model, involving the actual rates of withdrawal and injection, which are physically dependent upon the storage levels. Also, there are often options to renew contracts to prevent gas being stranded in store, and markets for gas in store.

However, concentrating on our simplified model, the discrete nature of the problem requires immediate consideration. Our previously derived formalism involves considering the state variable as being defined on a regularly spaced grid; the natural choice is to choose the amount of gas in store as this variable. With this choice, the control variables, which we take to be the amounts withdrawn or injected, must also coincide with the grid. To accomplish this we need to make rational approximations to the parameters described above. In examining this problem we note that often the three parameters in question are ordered as $x_{space} > x_{with} > x_{inj}$. This occurs when the gas is stored at higher pressures than the transmission system, and the compressors fitted to inject gas into store are not powerful enough to reverse the effect of this pressure difference. In this instance, it is natural to measure x_{space} and x_{with} in terms of x_{inj}.

In order to find the best rational approximation to the system we may expand each of these two quotients as continued fractions:

$$\frac{x_{\text{space}}}{x_{\text{inj}}} = m_0 + \cfrac{1}{m_1 + \cfrac{1}{m_2 + \cfrac{1}{m_3 + \cdots}}}$$

with a similar result for $x_{\text{with}}/x_{\text{inj}}$. The quantities in the right hand side of this equation are integers (for details of how to perform the expansion see, eg, Khinchin, 1964), and by truncating the expansion after successive terms we get a sequence of efficient rational approximations. For example, the Rough storage facility off the northeast coast of England sells storage in "bundled units" which consist of $x_{\text{space}} = 66.593407\,\text{kWh}$, $x_{\text{with}} = 1\,\text{kWh/day}$ and $x_{\text{inj}} = 0.351648\,\text{kWh/day}$. In this particular example, the rational approximations taken at successive orders, become:

$$x_{\text{with}} / x_{\text{inj}} = 2, 3, 17/6, \ldots$$

and

$$x_{\text{space}} / x_{\text{inj}} = 189, 379/2, 568/3, \ldots$$

Having chosen the required accuracy, the two rational numbers must be considered in terms of the least common multiple of their denominators, in order to obtain the final set of integers representing the three parameters which we write as (n_{space}, n_{with}, n_{inj}). In this particular example, an efficient choice for this set is (189,3,1).

The choice of state space representation leads directly to our choice of control parameters. Allowing all discrete transitions between allowed injection and withdrawal rates requires the choice of control parameter set as $a \in \{-n_{\text{with}}, -n_{\text{with}}+1, \ldots, 0, \ldots, n_{\text{inj}}\}$. However the results obtained in using such a set are often well approximated by the use of the smaller ternary set decision set $a \in \{-n_{\text{with}}, 0, n_{\text{inj}}\}$.

Having decided upon the discretisation of the state space, the equivalent discretisation of the prices is more elementary; we need only choose end points for the maximum and minimum price to be considered. This choice may be made by consideration of the expression for the forward probability density, such as that derived in the appendix, by choosing our range of prices in relation to the tails of the distribution and then choosing a grid spacing that satisfies our stability criterion. Also required are the boundary condition for the problem and the cashflow expression. The former is derived from our assumption of the expiry of our storage contract at time $t = T$. The direct consequence of this, as discussed, is the condition, $V_{T+1}(x, S) = 0$. In the most elementary case, the cashflow determined at each time step is simply the volume moving in or out of store, multiplied by the current commodity price, ie,

$$K(x, a, S) = -aSx_{inj}$$

noting that our set of control parameters is measured in units of the gas injection volume. Obviously, it is a simple matter to include extra, non-fuel variable costs in this function.

Having obtained these quantities we are now able to directly use the discrete form of the Bellman equation derived in the previous section.

THE VALUATION OF POWER GENERATION ASSETS

We now turn our attention to the application of our stochastic dynamic programming methodology to the valuation of power generation assets. This problem is a considerably more complex topic than that of gas storage, mainly due to the physical complexity of power plants themselves.

Before we examine the issues in greater detail, we note that the most obvious challenge to our existing methodology is that the value of a power plant comes from the spread between its fuel price (or fuel prices, if the plant is a dual fuelled plant) and the power price. Let us address this issue first by simply considering the situation for a gas-fired plant. We must now include both the power price, and the gas price in our spot price model. In order to calculate the appropriate conditional probabilities required by the Bellman equation, we note that power and gas commodity prices are frequently highly correlated. To account for this we model both prices simultaneously using two stochastic processes:

$$\frac{dS_p}{S_p} = \eta_p \ln\left(\bar{S}_p / S_p\right) dt + \sigma_p dW_p$$
$$\frac{dS_g}{S_g} = \eta_g \ln\left(\bar{S}_g / S_g\right) dt + \sigma_g dW_g$$

with $E(dW_p, dW_g) = \rho dt$. The subscripts denote power and gas respectively. We may proceed to the equivalent discrete formulation of the conditional density in a similar manner. Making the substitutions $y_p = \ln S_p$ and $y_g = \ln S_g$ and applying Itô's lemma, and then moving to a discrete formulation of the stochastic processes, yields:

$$y_{t+1}^p = y_t^p + \eta_p \left(\ln \bar{S}_p - y_t^p - \sigma_p^2 / 2\right) \delta t + \sigma_p \sqrt{\delta t} \phi_t^p$$
$$y_{t+1}^g = y_t^g + \eta_g \left(\ln \bar{S}_g - y_t^g - \sigma_g^2 / 2\right) \delta t + \sigma_g \sqrt{\delta t} \phi_t^g$$

where ϕ_p and ϕ_g are joint unit normal variables with correlation ρ. These quantities therefore have a joint density function of bi-normal form:

$$p(\phi_p, \phi_g) = \frac{1}{2\pi\sqrt{1-\rho^2}} \exp\left[-\frac{1}{2}\frac{\phi_p^2 - 2\rho\phi_p\phi_g + \phi_g^2}{1-\rho^2}\right]$$

Using a similar relation between densities as that used previously, we may write:

$$p\left(y_{t+1}^p, y_{t+1}^g \middle| y_t^p, y_t^g\right) = p(\phi_p, \phi_g) \Big/ \left|\frac{\partial\left(y_{t+1}^p, y_{t+1}^g\right)}{\partial(\phi_p, \phi_g)}\right|$$

and thereby proceed to a simple form for the joint conditional density.

To further outline the details of the calculation we must define our state variables. It will prove natural to define $x \in \{0, 1, 2, \ldots\}$ where the integers denote a particular capacity, as measured in MW, at which the plant can produce power in a stable fashion over the time step in question. We will denote the capacities themselves by the quantities $\{c_0, c_1, c_2, \ldots\}$. For instance, a simple example would be a plant that might run at three levels of output, 0 MW (ie, switched off), 100 MW or 300 MW. The lowest non-zero level is often termed the minimum stable generation. This choice of state variables requires the definition of control variables as $a \in \{-2, -1, 0, 1, 2\}$, and our equation of motion for the change in the state variable as $x_{t+1} = L(x_t, a_t) = x_t + a_t$. With these definitions we may proceed to define the cashflow function, $K(x, a, S^p, S^g)$, which is easiest to display in a matrix form:

		a_t				
		-2	-1	0	1	2
x_t	0	$-\infty$	$-\infty$	0	$c_1(S^p - hS^g - v) - s_{c,1}$	$c_2(S^p - hS^g - v) - s_{c,2}$
	1	$-\infty$	0	$c_1(S^p - hS^g - v)$	$c_2(S^p - hS^g - v)$	$-\infty$
	2	0	$c_1(S^p - hS^g - v)$	$c_2(S^p - hS^g - v)$	$-\infty$	$-\infty$

In the above formulae, the quantity h, termed the heat rate, is the number of units of fuel required for each unit of power produced, and v is a non-fuel variable cost. Both of these quantities may change with system state or control variable. The figure "$-\infty$" denotes the infinite cost associated with an impossible transition. The variables $s_{c,1}, s_{c,2}, \ldots$ represent start costs. These variables, which often have a significant impact on valuation, are costs involved in firing up the plant to start operation. They are often dependent on the fuel and/or power price, as the process of starting up the plant generally involves extra fuel use, or using an amount of power from the grid to commence generation. Typically, there are very large start costs if the requirement is to produce maximum output in very rapid times. In fact, extra components may be added to the matrix to specify additional costs particular to the capacity. A common example of a feature that may be incorporated in this way is the *ramp rate*, associated with a change of state. These are published maximum rates at which the capacity of a plant may be increased or decreased. Effectively, this finite rate means the plant is producing more or less than that assumed by a simple instantaneous switch, and this therefore corresponds to a power and/or fuel dependent correction.

As mentioned in the previous section on the theory of stochastic dynamic programming, it is very common for power generation assets to have additional constraints, and we discussed the method for incorporating them into our formalism. The process is to add additional significant state variables (together with their corresponding equations of motion). We briefly list some of the significant examples of variables to consider.

- *Emissions limits*. In many regions the total amount of carbon dioxide (CO_2) and sulphur dioxide (SO_2) produced by combusting fuel for power generation is capped annually. Considering the amount of these pollutants produced per unit of power, effectively this limit caps the amount of production. However, the plant operator still has the option of when to produce this power, and this optionality has a value. We may include this feature by expanding the set of state variables to include one such variable (eg, y_t), which records the cumulative energy output. We then enforce the constraint by including a boundary condition of the form:

$$V_T(x_T, y_T, S) = \left\{ \begin{matrix} -\infty & (y_T > y_{\max}) \\ V_1(x_T, 0, S) & (y_T < y_{\max}) \end{matrix} \right\}$$

  ie, if the limit is exceeded an (infinite) penalty is induced, otherwise the limit is reset in the new year.
- *"Take or pay"/"burner tip" fuel contracts*. Many plants have flexible fuel contracts associated with them, which allow the plant operator to decide how much fuel to take over the year, provided the annual aggregate is between minimum and maximum limits. As discussed previously, these contracts can typically be valued separately. However, they sometimes have a clause that the fuel cannot be resold (ie, it is for the "burner tip"). In contrast to the emissions limits above, this constraint enforces a minimum on the amount of energy production, and as such this effect can be treated in a similar way.
- *"Effective operating hours"*. Some power plants use complex calculations to indicate when expensive maintenance will occur, the details of which often derive from precise contractual arrangements with maintenance providers. A typical calculation

may involve the "effective operating hours" run by the plant; this quantity is calculated through an expression such as $h_{\mathrm{EOH}} = k_h h_{\mathrm{actual}} + k_s n_{\mathrm{starts}}$, where h_{actual} is the actual cumulative number of hours the plant has run and n_{start} is the cumulative number of times the plant has started up. Often the constants have the values in the range $k_h \approx 1$, but k_s may be an order or magnitude larger. When the quantity h_{EOH} passes through key values, significant capital expense is required. The effective operating hour quantity can be readily tracked as an additional state variable and a term added into the cashflow function, which adds the large costs according to the trigger values.

- *Minimum on/off times and hot/warm/cold starts*. A further element of complexity often occurs for power plants concerned with the temperature of the operating equipment. Various parts of the machinery have optimum operating temperatures that are considerably higher than ambient temperature. When the plant is switched off it obviously cools, and thus the costs of starting up power plants, which have been switched off for different times, vary. Often the start costs involved are classed as "hot", "warm", "cold" etc, and correspond to a characteristic time of last production. The converse of this effect is that many power plants have lead times to switching on or off, caused by other technical requirements. Both of these effects may be modelled by introducing the cumulative time the plant has been switched on or off for as a state variable, and then making the start costs a function of these quantities in the cashflow function.

Having derived the cashflow and considered the implications of possible requirements for extra state variables, we now have sufficient information to apply the stochastic dynamic programming formalism, derived in the previous chapter.

CONCLUSION

The nature of many contracts and physical assets in the energy markets implicitly involve the concept of sequential, contingent decision processes. Such processes naturally lead to recursive problems of the type that are frequently described by the theory of dynamic programming. To that end, we described the method of stochastic dynamic programming as a natural extension to deterministic programming, and derived a formalism where the stochastic nature of the problem arises from commodity prices satisfying a stochastic process. The generic form chosen for this process is one which is familiar from the modelling of commodity derivatives. The mathematical nature of the problem allowed some conclusions to be drawn, and further results to be derived. Fundamentally, the methodology provided an equation whose nature forced a numerical solution. We derived a simple, general method to achieve this and discussed the issues of its numerical stability. Finally, we discussed in detail the application of the method to the important problems of gas storage pricing and power generation pricing, with particular attention to the extensions and modifications needed to adapt to physical characteristics of those assets.

APPENDIX – PROPERTIES OF THE SCHWARTZ PROCESS

The one factor model described by the stochastic differential equation:

$$\frac{\mathrm{d}S}{S} = \eta \ln(\bar{S} / S)\mathrm{d}t + \sigma \mathrm{d}W$$

was introduced by Schwartz (Schwartz, 1997) as the simplest of three prototype models for commodity prices. It is particularly useful as many analytical results are available for it, and we summarise some of them here.

Here η represents a mean reversion rate, $\bar{S}$ represents a mean price level and σ a volatility. dW is a standard Wiener process. We assume throughout that $\bar{S}$ and σ are time-dependent, but η is not (although the mathematics is not made significantly more complex if this last constraint is relaxed). Note that in the limit $\sigma = 0$ and constant $\bar{S}$ the solution has the form:

$$\ln S(t) = \ln S(t_0) + \ln \bar{S}(1 - \exp(-\eta(t - t_0)))$$

then evidently the mean reversion parameter may be related to the half-life of the decay process of the logarithm of the price returning to its equilibrium value, that is the time taken for it to fall to half of its initial value: $t_{1/2} = \ln 2/\eta$.

We proceed to solve this process by noting that the substitution $z = \exp(\eta t) \ln S$ reduces it to the form:

$$dz = \exp(\eta t)[(\eta \ln \bar{S} - \sigma^2 / 2)dt + \sigma dW]$$

which may be immediately integrated as:

$$\ln S(t) = \exp(-\eta t) \ln S(0) + \exp(-\eta t) \int_0^t \exp(\eta t') \times [\eta \ln \bar{S}(t') - \sigma(t')^2 / 2]dt' + \exp(-\eta t) \int_0^t \exp(\eta t') \sigma(t') dW(t')$$

To study the properties of this expression we may use two well-known identities of the Itô stochastic calculus; for arbitrary non-anticipating functions f(t) and g(t) then:

$$E\left[\int_0^t f(t')dW(t')\right] = 0$$

$$E\left[\int_0^t f(t')dW(t') \int_0^t g(t')dW(t')\right] = \int_0^t f(t')g(t')dt'$$

(see, for instance Gardiner, 1985). Use of these formulae immediately yields:

$$E[\ln S(t)] = \exp(-\eta t) \ln S(0) + \exp(-\eta t) \int_0^t \exp(\eta t')[\eta \ln \bar{S}(t') - \sigma(t')^2 / 2]dt'$$

$$\text{var}[S(t)] = \exp(-2\eta t) \int_0^t \exp(2\eta t') \sigma^2(t') dt'$$

In order to obtain an explicit form for the forward probability density of $S(t)$, consideration of the mean-square limit for the Itô integral of an arbitrary function g(t) reveals:

$$\int_0^t g(t')dW(t') = \left[W_t \sqrt{\frac{1}{t} \int_0^t g^2(t')dt'} \right]$$

with W_t being the Wiener process. For practical purposes we may consider this quantity as equivalent to a Brownian process

(Neftci, 2000) and hence having the density function:

$$p_W(W) = \frac{1}{\sqrt{2\pi t}} \exp(-W^2 / 2t)$$

This function may be used to directly infer a density function for $S(t)$:

$$p(S(t)) = \frac{1}{\sqrt{2\pi S^2 \exp(-2\eta t)\int_0^t \exp(2\eta t')\sigma^2(t')\mathrm{d}t'}}$$
$$\times \exp\left[-\frac{\left(\ln S(t) - \exp(-\eta t)\ln S(0) - \exp(-\eta t)\int_0^t \exp(\eta t')[\eta \ln \bar{S}(t') - \sigma(t')^2/2]\mathrm{d}t'\right)^2}{2\exp(-2\eta t)\int_0^t \exp(2\eta t')\sigma^2(t')\mathrm{d}t'}\right] \quad \mathbf{(3)}$$

By successive integration by parts a useful asymptotic approximation can be found for terms in the expression for the density in equation (3).

$$\exp(-\eta t)\int_0^t \exp(\eta t')\mathrm{f}(t')\mathrm{d}t' \approx \mathrm{f}(t)/\eta + \mathrm{f}'(t)/\eta^2 + \mathrm{f}''(t)/\eta^3 + \cdots$$

By using this approximation, truncated after the first term, in the density approximation an expression for the density as a lognormal distribution is obtained which has an instantaneous mean of $\bar{S}\exp(-\sigma^2/4\eta)$ and variance of $\bar{S}^2(1-\exp(-\sigma^2/2\eta))$.

A case of particular relevance in the energy markets, particularly that of power, is where the mean reversion level and volatility are periodic functions with frequency ω. For instance suppose $\ln \bar{S}(t) \approx \cos(\omega t)$. In this instance:

$$\exp(-\eta t)\int_0^t \exp(\eta t')\ln \bar{S}(t')\mathrm{d}t' \approx \frac{\eta\cos\omega t + \omega \sin\omega t}{\eta^2 + \omega^2}$$

and hence, for large frequencies the mean and variance are periodic with the same frequency but pick up a phase lag of $\arctan(\omega/\eta)$. This can cause complications when trying to simulate prices which contain high frequency Fourier components.

The most proper way to evaluate parameters for the Schwartz model is through a regression analysis of the price return against the logarithm of the price level, however, the above description of the process in the limit of high mean reversion yields the following approximate formulae for the parameters:

$$\bar{S} \approx \langle S \rangle$$

$$\sigma \approx \sqrt{\left\langle \left(\mathrm{d}S / S - \langle \mathrm{d}S / S \rangle \right)^2 \right\rangle \Big/ \mathrm{d}t}$$

$$\eta \approx \frac{\sigma^2 \langle S \rangle^2}{2\left\langle \left(S - \langle S \rangle \right)^2 \right\rangle}$$

where the angled brackets denote an average over an appropriate period of historical data.

As an example of a highly volatile spot energy price, we plot in Figure 5 the UK system average price (SAP) – an average of within day gas price, in units of pence/therm (Transco, 2002). By fitting a Schwartz model as described above, we find:

❑ $\sigma \approx 230\%$
❑ $\eta \approx 211$

after averaging over the period of observation. Both quantities have been annualised (that is they have units of (*years*)$^{-1/2}$ and

Figure 5 System average price (SAP). The average of UK within day gas prices

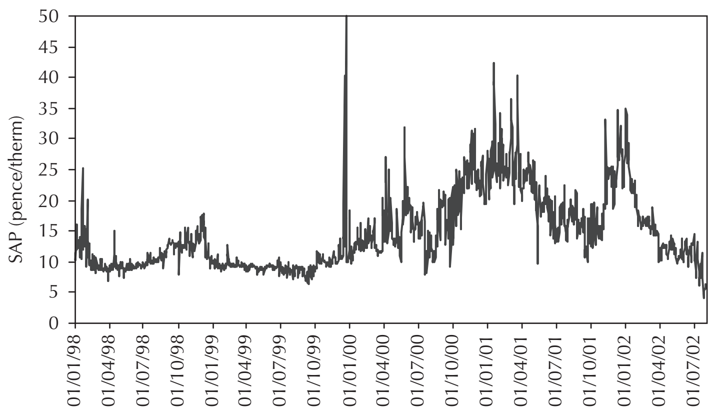

(*years*)$^{-1}$ respectively). The value of the mean reversion parameter suggests a half-life for the relaxation of the logarithm of the price of just over one day.

1 A detailed discussion of the appropriate choice of discount factor (beta) for particular contracts or projects is not within the scope of an introductory account of dynamic programming methods, however a few remarks are appropriate. If a liquidly traded underlying market exists whose tenor coincides with the granularity of the optionality inherent in the asset, then an argument can be formulated based upon portfolio replication and arbitrage, such that a discount factor close to the risk-free interest rate should be used (with a possible extra contribution to compensate for physical operational risk). Where this is not the case, different hedging strategies are possible which lead to different degrees of market risk. Evidently the choice of riskier strategies must be compensated by a higher expected return, and consequently a higher discount factor applied. A method of calculation of this factor is to consider the stock prices of related businesses and apply a model, such as the Capital Asset Pricing Model, to calculate the excess return required by investors at a level of risk appropriate to the hedging strategy. In this instance, since a risk-neutral approach is not being adopted the dynamics of the risk factors need to be considered also.

2 Which may be more succinctly stated as an optimal policy (choice of control parameters) has the property that, whatever the initial action, the remaining choices constitute an optimal policy with respect to the sub-problem that result from the initial actions. (See Bather, 2000 and Dixit and Pindyck, 1993.)

BIBLIOGRAPHY

Abramowitz, M., and I. A. Stegun, 1965, *Handbook of Mathematical Functions* (New York: Dover).

Bather, J., 2000, *Decision Theory* (New York: Wiley).

Bellman, R., 1957, *Dynamic Programming* (NJ: Princeton University Press).

Cox, R. D., and H. D. Miller, 1965, *The Theory of Stochastic Processes* (London: Chapman and Hall).

Dixit, A. K., and R. S. Pindyck, 1993, *Investment under Uncertainty* (NJ: Princeton University Press).

Elton, E. J., and M. J. Gruber, 1995, *Modern Portfolio Theory and Investment Analysis*, Fifth Edition (New York: Wiley).

Gardiner, C. W., 1985, *Handbook of Stochastic Methods*, Second Edition (Berlin: Springer-Verlag).

Grimmett, G. R., and D. R. Stirzaker, 1992, *Probability and Random Processes*, Second Edition (Oxford: OUP).

Kreysig, E., 1999, *Advanced Engineering Mathematics*, Eighth Edition (New York: Wiley).

Khinchin, A. I., 1997, *Continued Fractions* (New York: Dover).

Ljungqvist, L., and T. J. Sargent, 2000, *Recursive Macroeconomic Theory* (Cambridge MA: MIT Press).

Neftci, S., 2000, *Introduction to the Mathematics of Financial Derivatives*, Second Edition (New York: Academic Press).

Schwartz, E. S., 1997, "The stochastic behaviour of Commodity Prices: Implications for Valuation and Hedging" *Journal of Finance*, 52 (3).

Transco, 2002, http://www.transco.uk.com/services/flexhome.htm