Dynamic Programming determines optimal strategies among a range of possibilities typically putting together 'smaller' solutions. In some cases it is little more than a careful enumeration of the possibilities but can be organized to save effort by only computing the answer to a small problem once rather than many times.

There are ways to adapt Dynamic Programming to a stochastic event. In our cases we envision that the demand for a product varies stochastically. We seek optimal strategies that optimize the expected return.

Let $E(v_i(s_i))$ denote the expected value of being in state s_i with i periods to go. That is we choose a decision and then get an expected return.

$$E(v_i(s_i)) = \max\{E(f_i(d_i, s_i)) + E(v_{i-1}(t_i(d_i, s_i))) \,:\, d_i \in D_i(s_i)\}$$

We consider the following example from the book *Applied Mathematical Programming* by Bradley, Hax and Magnanti. We have a company that must operate for a period of months, producing and storing a commodity for which the monthly demand varies stochastically.

monthly demand	probability
0	.25
1	.40
2	.20
3	.15

There are various rules. At most 3 units can be stored from one monthto the next with an inventory cost of \$100 charged to the month sending the inventory forward. Each unit costs \$1000 to produce and will sell for \$2000 if there is demand. Each unsold unit can be salvaged for \$500 at the end of the time period.

We use I_i in place of s_i to denote the state. Here I_i denotes the inventory with i months to go. What follows is the optimal calculation for one month to go.

The spreadsheet dynamicsales.xls does these calculations for us in an awkward spreadsheet language. Fortunately you need only type in aminimal number of these computations since Excel is made for such purposes. We can obtain the table for 3 months to go by copying the table for 2 months to go a suitable number of spaces directly below the table for 2 months to go so that everything lines up. You can type in the optimal decisions (or do it by excel if you prefer!) The basic data for the stochastic demand is given as A5,A6,A7,A8 with corresponding probabilities B5,B6,B7,B8. The initial table for $v_0(I_0)$

I_0	$v_0(I_0)$
0	0
1	500
2	1000
3	1500

is stored at F5,F6,F7,F8. From then on each month is stored

I_1	d_1	sell	prob	I_0	prod	sales cost	Inv	$v_0(I_0)$	expected value	expected profit
0	0	0	1.00	0	0	0	0	0	0	0
	1	0	.25	1	−1000	0	−100	500	−150	600*
		1	.75	0	−1000	2000	0	0	750	
	2	0	.25	2	−2000	0	−200	1000	−300	560
		1	.40	1	−2000	2000	−100	500	160	
		2	.35	0	−2000	4000	0	0	700	
	3	0	.25	3	−3000	0	−300	1500	−450	200
		1	.40	2	−3000	2000	−200	1000	−80	
		2	.20	1	−3000	4000	−100	500	280	
		3	.15	0	−3000	6000	0	0	450	
1	0	0	.25	1	0	0	−100	500	100	1600*
		1	.75	0	0	2000	0	0	1500	
	1	0	.25	2	−1000	0	−200	1000	−50	1560
		1	.40	1	−1000	2000	−100	500	560	
		2	.35	0	−1000	4000	0	0	1050	
	2	0	.25	3	−2000	0	−300	1500	−200	1200
		1	.40	2	−2000	2000	−200	1000	320	
		2	.20	1	−2000	4000	−100	500	480	
		3	.15	0	−2000	6000	0	0	600	
2	0	0	.25	2	0	0	−200	1000	200	2560*
		1	.40	1	0	2000	−100	500	960	
		2	.35	0	0	4000	0	0	1400	
	1	0	.25	3	−1000	0	−300	1500	50	2200
		1	.40	2	−1000	2000	−200	1000	720	
		2	.20	1	−1000	4000	−100	500	680	
		3	.15	0	−1000	6000	0	0	750	
3	0	0	.25	3	0	0	−300	1500	300	3200*
		1	.40	2	0	2000	−200	1000	1120	
		2	.20	1	0	4000	−100	500	880	
		3	.15	0	0	6000	0	0	900	

You might expect we settle down to a fixed optimal strategy the further backwards in time we go. In this example this happened very quickly at 3 months to go with the optimal decisions (depending on the initial state) as follows

I_i	$d_i^*(I_i)$
0	3
1	2
2	1
3	0

Interestingly you can summarize this strategy as saying ordering so that you have 3 units on hand so that (I imagine) you don't lose any profitable sales. If $v_i = j$ then $d_i^*(j) = 3 - j$. If you use excel to compute the values $v_i(I_i)$ for many months to go you discover that $v_i(I_i) - v_{i-1}(I_{i-1}) = 1075$ suggesting that each month you can make a profit of \$1075.

Given the fixed optimal decisions, it is easy to compute the vector of entries $v_i(I_i)$ from the values $v_{i-1}(I_{i-1})$. We compute the expected revenues from sales that month−inventory costs as

$$.25 \times (-300) + .4 \times 1800 + .2 \times 3900 + .15 \times 6000 = 2325$$

We are ignoring purchasing costs of goods as well as $\mathbf{v}_{i-1}$; the value of being $i-1$ months to go with specifed inventory. The contribution to expected revenue from the value of the left over inventory of j with 1 fewer month to go would be $v_{i-1}(j)$. So the expected revenue from this source is

$$.25 \times v_{i-1}(3) + .40 \times v_{i-1}(2) + .20 \times v_{i-1}(1) + .15 \times v_{i-1}(0).$$

After including the purchasing costs we obtain the matrix equation

$$\begin{bmatrix} v_i(0) \\ v_i(1) \\ v_i(2) \\ v_i(3) \end{bmatrix} = \begin{bmatrix} -3000 + 2325 \\ -2000 + 2325 \\ -1000 + 2325 \\ 0 + 2325 \end{bmatrix} + \begin{bmatrix} .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \end{bmatrix} \begin{bmatrix} v_{i-1}(0) \\ v_{i-1}(1) \\ v_{i-1}(2) \\ v_{i-1}(3) \end{bmatrix}$$

or

$$\mathbf{v}_i = \begin{bmatrix} -675 \\ 325 \\ 1325 \\ 2325 \end{bmatrix} + \begin{bmatrix} .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \end{bmatrix} \mathbf{v}_{i-1},$$

where we define

$$\mathbf{v}_i = \begin{bmatrix} v_i(0) \\ v_i(1) \\ v_i(2) \\ v_i(3) \end{bmatrix}.$$

We now do some matrix algebra. The matrix

$$A = \begin{bmatrix} .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \end{bmatrix}$$

is stochastic (has row sums =1) and so has the vector of 1's as an eigenvector of eigenvalue 1. Other eigenvalues are smaller so the eigenvector of eigenvalue 1 dominates.

Let

$$\mathbf{b} = \begin{bmatrix} -675 \\ 325 \\ 1325 \\ 2325 \end{bmatrix}$$

so that $\mathbf{v}_i = \mathbf{b} + A\mathbf{v}_{i-1}$ We are interested in $\mathbf{v}_i - \mathbf{v}_{i-1}$ for largish i. In this example we are OK for $i \geq 3$ so largish doesn't mean much here. Using the idea of the dominant eigenvector of eigenvalue 1 we have $A^k(\mathbf{v}_i) \approx A^{k-1}(\mathbf{v}_i)$. So

$$\mathbf{v}_i = \mathbf{b} + A\mathbf{b} + A^2\mathbf{b} + \cdots A^{t-1}\mathbf{b} + A^t\mathbf{v}_{i-t}, \qquad \mathbf{v}_{i-1} = \mathbf{b} + A\mathbf{b} + A^2\mathbf{b} + \cdots A^{t-2}\mathbf{b} + A^{t-1}\mathbf{v}_{i-t}$$

from which we deduce, using $A^t(\mathbf{v}_{i-t}) \approx A^{t-1}(\mathbf{v}_{i-t})$, that $\mathbf{v}_i - \mathbf{v}_{i-1} \approx A^{t-1}\mathbf{b}$. Now again using dominant eigenvector arguments we can estimate $A^{t-1}\mathbf{b}$ by writing $\mathbf{b}$ as a linear combination of the eigenvectors where we are only interested in the eigenvector of eigenvalue 1. For this particular A, everything is quite easy. The eigenvalues of A are 1 with multiplicity 1 and eigenvalue 0 with multiplicity 3 and so we may write

$$\mathbf{b} = \begin{bmatrix} -675 \\ 325 \\ 1325 \\ 2325 \end{bmatrix} = 1075 \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} + \begin{bmatrix} -1750 \\ -750 \\ 250 \\ 1250 \end{bmatrix}$$

where the second vector is an eigenvector of A of eigenvalue 0:

$$\begin{bmatrix} .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \\ .15 & .20 & .40 & .25 \end{bmatrix} \begin{bmatrix} -1750 \\ -750 \\ 250 \\ 1250 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

This is a simple case; in a more typical case we would have several eigenvectors of different eigenvalues but they all would have eigenvalues less than 1 in magnitude. We conclude

$$\mathbf{v}_i - \mathbf{v}_{i-1} \approx A^{t-1}\mathbf{b} = 1075 \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}.$$

The special features of A that the other eigenvalues are all 0 in fact means this is an equality.

Why do we interpret this as making \$1075 per month? We have $\mathbf{v}_i(j) - \mathbf{v}_{i-1}(j) \approx 1075$ (actually equality for $i > 3$) .and so the difference of the value of starting with j units on hand with i months to go compared with $i-1$ months to go is \$1075; the extra month has given \$1075 more value.