

Lesson 15. Introduction to Stochastic Dynamic Programming

1 Motivation

- In the dynamic programs we have studied so far, the transitions from one state to the next are **deterministic**
- For example, the knapsack problem:
 - Suppose we are in stage t and state n (deciding whether to take metal t with n kg of space remaining)
 - If we decide to take metal t in stage t , we know exactly what state we will be in stage $t+1$:
- What if the transitions between states are subject to some randomness or **stochasticity**?

2 An example

The Dijkstra Brewing Company is planning production of its new limited run beer, Primal Pilsner, over the next 2 months. Based on some market analysis studies, the company has determined that the demand for the new beer in each month will be:


Demand (batches)	Probability
0	1/4
2	3/4

Each batch of beer costs \$3,000 to produce. Batches can be held in inventory at a cost of \$1,000 per batch per month. Each month, the company can produce either 0 or 1 batches, due to capacity limitations. In addition, the size of the company's warehouse restricts the ending inventory for each month to at most 2 batches. The company has 1 batch ready to go in inventory.

Due to contractual obligations, there is a penalty of \$5,000 for each batch of demand not met. Any batches produced that cannot be stored in the company's warehouse gets thrown away, and cannot be used to meet future demand.

The company wants to find a production plan that will minimize its total production and holding costs over the next 3 months.

- Let's think about the decision-making process starting at month 1
- Let:
 - Node t_n represent month t with n batches in inventory
 - x_t represent the number of batches to produce in month 1
 - d_t represent the number of batches in demand in month 1
- We can draw the following diagram (that looks like a graph) that models the decision-making process



- We can diagram the entire 2-month process in a similar fashion:

$x_1 = 1$

$x_2 = 1$

1_2

2_2

3_2

$x_1 = 0$

$x_2 = 0$

$x_1 = 1$

$x_2 = 1$

1_1

2_1

3_1

$x_1 = 0$

$x_2 = 0$

$x_1 = 1$

$x_2 = 1$

1_0

2_0

3_0

$x_1 = 0$

$x_2 = 0$

- Consider the following production policy:
 - In month 1, produce 1 batch
 - In month 2:
 - ◊ If there are 2 batches in inventory, produce 0 batches
 - ◊ If there are 0 batches in inventory, produce 1 batch
- What is the expected cost of this policy?
- Working backwards:
 - Expected cost in month 2 with 2 batches in inventory (node 2_2):

- Expected cost in month 2 with 0 batches in inventory (node 2_1):

- Expected cost in month 1 (node 1_1):

3 Things to think about

- The policy above gives **contingency plans**
- The diagram we drew on page 3 sort of looks like a shortest path problem, but it's not!
- We cannot solve this example as a shortest path problem, since the edges “are random”
- We can, however, still write a recursion to represent this example problem, and others like it
- We'll explore these ideas in the next few lessons