

Components of DP formulations

- What are the main elements of a DP formulation?

A dynamic programming model for baseball...

<http://www.footballcommentary.com/bbmodel.htm>



Other applications of DP to sports:

Other Sports: <http://mat.tepper.cmu.edu/blog/?cat=9>

Assignments and Presentations

Homework Assignments:

- Assignment 1 is up on ctools. Due start of class next Thursday
- There will be an assignment every 7-10 days through the semester
- I will grade a sample of the questions but provide detailed solutions for all questions

In-Class Assignments:

- Work in teams of 2-3 in class to complete a problem. Hand in one copy with your names on it.
- Graded on a scale of 0-3

Participation: 10% for in-class assignments and 5% for the mini project presentation

Q. What questions do you have about the course?

- Is this course applied or theoretical? **Yes**
- How do we submit assignments? **Hardcopy at the start of class on the due date. Handwritten is fine. Include a cover page with your name.**
- Can we get the slides before class? **I'll do my best to get them up at least an hour before.**
- Will there be student presentations? **Yes, team presentations (teams of 2-3) on a project of interest to you. You will write a proposal for your project mid-semester but you can start planning anytime.**
- What kind of software is used for this course? **Matlab** Will AMPL be used? **No required**
- When is the midterm and final exam? **See class schedule on canvas**
- Do you “curve the grades”? **No. I may increase the final course grade depending on class performance.**

Q. What questions do you have about the course?

- Will I be ok without the prerequisites for this course? **Depends. You will likely need to do extra work to catchup.**
- What sources do you recommend for a first introduction to LP and Markov Chains? **Operations Research: Applications and Algorithms, by Wayne Winston is an introductory text that covers both. IOE510 and IOE515 are excellent courses**
- Will Matlab be taught? **No**
- Do we need to write our own code or do we use premade Matlab algorithms for dynamic programming? **You need to write your own Matlab code.**
- What are the requirements for groups for in-class assignments? **Groups must be 2 or 3 people. The group does not need to be the same every time.**
- Does this course overlap others? **Yes, it relates to many others including IOE510, IOE515, and others on optimization and stochastic models.**

Q. What do you hope to learn?

- Theory of DP
- How to apply DP to real world problems
- How to formulate DPs
- Computational methods for DPs
- Large scale/big data methods, Approximate dynamic programming
- Increased understanding of stochastic modeling

Applications of Interest

- Supply chain applications
- Finance and revenue management
- Transportation
- Healthcare applications
- Energy systems
- Robotics
- Inventory control
- Games
- Airlines

- The Lot-sizing problem
 - Wagner-Whitin Algorithm

General DP Formulation

A DP has **states**, $s_t \in S$, **actions** $a_t \in A$, **rewards**, $r_t(s_t, a_t)$, and an **optimal value function**, $v_t(s_t)$, defined for stages $t = 1, 2, \dots, T$

Optimality Equations:

$$v_t(s_t) = \max_{a_t \in A} \{r_t(s_t, a_t) + v_{t+1}(s_{t+1})\}, \quad \forall s_t$$

$$v_T(s_T) = R(s_T), \quad \forall s_T$$

$v_t(s_t)$ is the maximum total reward for all stages $t, t + 1, \dots, T$, also called the “optimal value to go”

Transition from s_t to s_{t+1} governed by a **transfer equation**:

$$s_{t+1} = g(s_t, a_t)$$


Often it is reasonable to weight “rewards” received later in time less than rewards received earlier in time

Assume that for some $\lambda < 1$, \$1 received at the beginning of year $t+1$ is equivalent to λ dollars received at the beginning of year t

Example: What is the **net present value** (NPV) of the following stream of payments if $\lambda = 0.95$?

- Year 1: \$100, Year 2: \$200, Year 3: \$300, Year 4: \$400

Add a discount factor to incorporate **discounting**:


$$v_t(s_t) = \max_{a_t \in A} \{r_t(s_t, a_t) + \lambda v_{t+1}(s_{t+1})\}, \quad \forall s_t$$

$$v_T(s_T) = R(s_T), \quad \forall s_T$$

Discount factor $\lambda \in [0,1]$ discounts future rewards

As a result $v_t(s_t)$ is the maximum *net present value* of rewards for all stages $t, t + 1, \dots, T$

Lot-sizing Problem

The dynamic lot-size problem arises in many manufacturing contexts

The decision maker must consider: production (fixed and variable costs), inventory, and demand



Problem Description:

1. Demand d_t during periods (stages) $t = 1, 2, \dots, T$
2. Demand for period t must be met by period t inventory or production
3. Cost of producing a_t units in period t is given by: $c(0) = 0$, and for $a_t > 0$, $c(a_t) = K + ba_t$, where K is a fixed cost, and b is a variable (per-unit) cost
4. At the start of period t , the inventory level s_t is observed
5. An inventory holding cost, h , is incurred for any inventory left at the end of period t

Questions: What are the states, actions, and “rewards” for this problem?

Timing:

- The firm's inventory position is reviewed at the end of each stage and then the production decision is made (referred to as a **periodic review policy**).

Immediate "Reward":

$$r_t(s_t, a_t) = h_t \times (s_t + a_t - d_t) + c(a_t)$$

Transfer Equation:

$$s_{t+1} = s_t + a_t - d_t$$

Optimality Equations:

$$v_t(s_t) = \min_{a_t \in A} \{h_t \times (s_t + a_t - d_t) + c(a_t) + v_{t+1}(s_t + a_t - d_t)\}, \forall s_t, t = 1, \dots, T$$

Boundary Condition: $V_{T+1}(s_{T+1}) = 0, \forall s_t$

The state space is unbounded, making the optimality equations (potentially) difficult to solve for large problems

Limits on production and capacity can be used to simplify the problem by limiting the number of possible states.

Inventory Example

A company knows that demand for a product for the next 4 months is:

Month	Demand
1	1
2	3
3	2
4	4

At the beginning of each month the company decides how many units to produce. During a month in which any units are produced there is a setup cost of \$3. In addition, there is a variable cost of \$1 per unit produced. At the end of each month there is a holding cost of \$0.5 per unit. Manufacturing capacity limits production to 5 units per month, and space limits inventory to 4 units. Determine the optimal production schedule.

- Wagner and Whitin (1958)** recognized some **special properties** of the lot-sizing problem that make it easier to solve via DP when costs are stationary (i.e. they don't vary in time)
- Basic idea: optimal production quantities will always enough to cover an integer number of future periods demand
- As a result, the number of states and actions for the DP is signif

** See “Background Reading” Directory on Canvas: Wagner and Whitin, 1958, “Dynamic Version of the Economic Lotsize Problem”, 1958

The following Theorems are important for the development of the Wagner-Whitin algorithm

Theorem 1: If it is optimal to produce anything during stage t , i.e., if $a_t > 0$ then $s_t = 0$.

Proof:

Suppose an optimal production plan suggests $s_t > 0$ and $a_t > 0$. Then inventory holding cost can be reduced by delaying the production that resulted in s_t until period t . ■

Theorem 2: Suppose it is optimal to produce in stage t . Then for some $j = 0, 1, \dots, T - t$, it is optimal to produce an amount that exactly suffices to meet the demands for stages $t, t + 1, \dots, t + j$.

Proof:

Since all demands must be met, any other action implies there exists a period t' such that $s_{t'} a_{t'} > 0$. This contradicts Theorem 1. ■

Given Theorems 1 and 2, with the possible exception of the first stage, production will occur only during periods in which beginning inventory is zero; during each stage in which beginning inventory is zero (and $d_t \neq 0$), production must occur

Wagner-Whitin Algorithm

Applying Theorems 1 and 2 Wagner and Whitin developed the following recursion:

$$v_t = \min_{j=0,1,2,\dots,T-t} (c_{tj} + v_{t+j+1})$$

Where $v_{T+1} = 0$ and c_{tj} is the total cost to satisfy demands for stages $t, t + 1, \dots, t + j$. In other words:

$$c_{tj} = K + h(jd_{t+j} + (j-1)^+ d_{t+j-1} + \dots + (j - (j+1))^+ d_{t+1})$$

To find an optimal production schedule start by finding v_T and then compute $v_{T-1}, v_{T-2}, v_{T-3}, \dots, v_1$

*Note: Since costs are stationary, the total production cost is constant and can be ignored

In-Class Assignment

Use the Wagner-Whitin algorithm to solve the following problem:

- $K = \$3$
- $h = \$0.50$
- $b = \$1$
- $d_1 = 1$, $d_2 = 3$, $d_3 = 2$, $d_4 = 4$

What is the optimal production policy? What is the total cost associated with this policy?

Plan for next few classes

- Formulate a series of problems that can be solved by deterministic dynamic programming:
 - Equipment replacement
 - Knapsack problem
 - Travelling salesperson problem
 - Pattern recognition
- Return to shortest path problems with special features
- Reading: Chapters 1-3 of Dreyfus and Law