Wagner–Whitin Algorithm and Silver–Meal Heuristic for Dynamic Lot-Size Model

A periodic review inventory model in which each period's demand is known at the beginning of the problem is a **dynamic lot-size model.** A cost-minimizing production or ordering policy may be found via a backward recursion, a forward recursion, the Wagner–Whitin algorithm, or the Silver–Meal heuristic.

The Wagner–Whitin algorithm uses the fact that production occurs during a period if and only if the period's beginning inventory is zero. The decision during such a period is the number of consecutive periods of demand that production should meet.

During a period in which beginning inventory is zero, the Silver–Meal heuristic computes the average cost per period (setup plus holding) incurred in meeting the demand during the next k periods. If k^* minimizes this average cost, then the next k^* periods of demand should be met by the current period's production.

Computational Considerations

Dynamic programming is much more efficient than explicit enumeration of the total cost associated with each possible set of decisions that may be chosen during the T stages. Unfortunately, however, many practical applications of dynamic programming involve very large state spaces, and in these situations, considerable computational effort is required to determine optimal decisions.

REVIEW PROBLEMS

Group A

1 In the network in Figure 14, find the shortest path from node 1 to node 10 and the shortest path from node 2 to node 10.

FIGURE 14

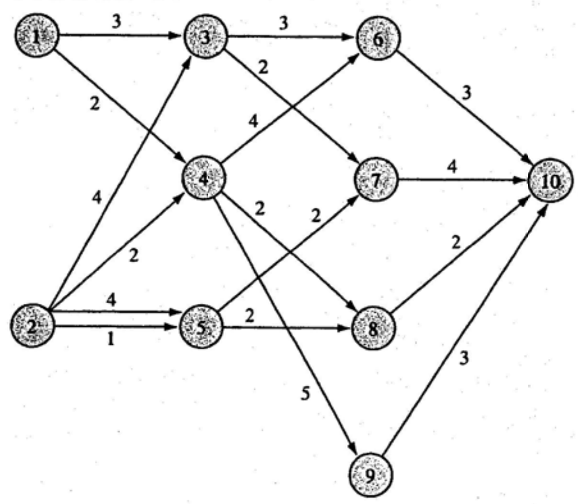

2 A company must meet the following demands on time: month 1, 1 unit; month 2, 1 unit; month 3, 2 units; month 4, 2 units. It costs $4 to place an order, and a $2 per-unit holding cost is assessed against each month's ending inventory. At the beginning of month 1, 1 unit is available. Orders are delivered instantaneously.

a Use a backward recursion to determine an optimal ordering policy.

b Use the Wagner–Whitin method to determine an optimal ordering policy.

c Use the Silver–Meal heuristic to determine an ordering policy.

3 Reconsider Problem 2, but now suppose that demands need not be met on time. Assume that all lost demand is backlogged and that a $1 per-unit shortage cost is assessed against the number of shortages incurred during each month. All demand must be met by the end of month 4. Use dynamic programming to determine an ordering policy that minimizes total cost.

4 Indianapolis Airlines has been told that it may schedule six flights per day departing from Indianapolis. The destination of each flight may be New York, Los Angeles, or Miami. Table 20 shows the contribution to the company's profit from any given number of daily flights from Indianapolis to each possible destination. Find the optimal number of flights that should depart Indianapolis for each destination. How would the answer change if the airline were restricted to only four daily flights?

TABLE 20

	Profit per Flight ($)					
	Number of Planes					
Destination	1	2	3	4	5	6
New York	80	150	210	250	270	280
Los Angeles	100	195	275	325	300	250
Miami	90	180	265	310	350	320

5 I am working as a cashier at the local convenience store. A customer's bill is $1.09, and he gives me $2.00. I want to give him change using the smallest possible number of coins. Use dynamic programming to determine how to give the customer his change. Does the answer suggest a general result about giving change? Resolve the problem if a 20¢ piece (in addition to other United States coins) were available.

6 A company needs to have a working machine during each of the next six years. Currently, it has a new machine. At the beginning of each year, the company may keep the machine or sell it and buy a new one. A machine cannot be kept for more than three years. A new machine costs $5,000. The revenues earned by a machine, the cost of maintaining it, and the salvage value that can be obtained by selling it at the end of a year depend on the age of the machine (see Table 21). Use dynamic programming to maximize the net profit earned during the next six years.

7 A company needs the following number of workers during each of the next five years: year 1, 15; year 2, 30; year 3, 10; year 4, 30; year 5, 20. At present, the company has 20 workers. Each worker is paid $30,000 per year. At the beginning of each year, workers may be hired or fired. It costs $10,000 to hire a worker and $20,000 to fire a worker. A newly hired worker can be used to meet the current year's worker requirement. During each year, 10% of all workers quit (workers who quit do not incur any firing cost).

a With dynamic programming, formulate a recursion that can be used to minimize the total cost incurred in meeting the worker requirements of the next five years.

b How would the recursion be modified if hired workers cannot be used to meet worker requirements until the year following the year in which they are hired?

8 At the beginning of each year, Barnes Carr Oil sets the world oil price. If a price p is set, then $D(p)$ barrels of oil will be demanded by world customers. We assume that during any year, each oil company sells the same number of barrels of oil. It costs Barnes Carr Oil c dollars to extract and refine each barrel of oil. Barnes Carr cannot set too high a price, however, because if a price p is set and there are currently N oil companies, then $g(p, N)$ oil companies will enter the oil business [$g(p, N)$ could be negative]. Setting too high a price will dilute future profits because of the entrance of new companies. Barnes Carr wants to maximize the discounted profit the company will earn over the next 20 years. Formulate a recursion that will aid Barnes Carr in meeting its goal. Initially, there are 10 oil companies.

9 For a computer to work properly, three subsystems of the computer must all function properly. To increase the reliability of the computer, spare units may be added to each system. It costs $100 to add a spare unit to system 1, $300 to system 2, and $200 to system 3. As a function of the number of added spares (a maximum of two spares may be added to each system), the probability that each system will work is given in Table 22. Use dynamic programming to maximize the probability that the computer will work properly, given that $600 is available for spare units.

Group B

10 During any year, I can consume any amount that does not exceed my current wealth. If I consume c dollars during a year, I earn c^a units of happiness. By the beginning of the next year, the previous year's ending wealth grows by a factor k.

a Formulate a recursion that can be used to maximize total utility earned during the next T years. Assume I originally have w_0 dollars.

b Let $f_t(w)$ be the maximum utility earned during years $t, t + 1, \ldots, T$, given that I have w dollars at the beginning of year t; and $c_t(w)$ be the amount that should be consumed during year t to attain $f_t(w)$. By working backward, show that for appropriately chosen constants a_t and b_t,

$$f_t(w) = b_t w^a \quad \text{and} \quad c_t(w) = a_t w$$

Interpret these results.

11 At the beginning of month t, farmer Smith has x_t bushels of wheat in his warehouse. He has the opportunity to sell wheat at a price s_t dollars per bushel and can buy wheat at p_t dollars per bushel. Farmer Smith's warehouse can hold at most C units at the end of each month.

a Formulate a recursion that can be used to maximize the total profit earned during the next T months.

b Let $f_t(x_t)$ be the maximum profit that can be earned during months $t, t + 1, \ldots, T$, given that x_t bushels of

TABLE 21

	Age of Machine at Beginning of Year		
	0 Year	1 Year	2 Years
Revenues ($)	4,500	3,000	1,500
Operating Costs ($)	500	700	1,100
Salvage Value at End of Year ($)	3,000	1,800	500

TABLE 22

Number of Spares	Probability That a System Works		
	System 1	System 2	System 3
0	.85	.60	.70
1	.90	.85	.90
2	.95	.95	.98

or

$$p_H w_H \geq p_S w_S(1 + p_H - p_S) \quad (7)$$

For example, if $p_H = .60$, $p_S = .90$, $w_H = .55$, and $w_S = .50$, then (5) and (7) are both satisfied, and Martina should serve hard on her first serve and soft on her second serve. On the other hand, if $p_H = .25$, $p_S = .80$, $w_H = .60$, and $w_S = .45$, then both serves should be soft. The reason for this is that in this case, the hard serve's advantage from the fact that w_H exceeds w_S is outweighed by the fact that a hard serve on the first serve greatly increases the chances of a double fault.

To complete our analysis, we must consider the situation where (5) does not hold. We now show that if

$$p_H w_H \geq p_S w_S \quad (8)$$

Martina should serve hard on both serves. Note that if (8) holds, then $f_2 = \max\{p_H w_H, p_S w_S\} = p_H w_H$, and Martina should serve hard on the second serve. Now (6) implies that Martina should serve hard on the first serve if

$$p_H w_H + (1 - p_H) p_H w_H \geq p_S w_S + (1 - p_S) p_H w_H$$

Upon rearrangement, the last inequality becomes

$$p_H w_H(1 + p_S - p_H) \geq p_S w_S$$

Dividing both sides of the last inequality by $p_S w_S$ shows that Martina should serve hard on the first serve if

$$\frac{p_H w_H}{p_S w_S}(1 + p_S - p_H) \geq 1$$

After noting that $p_H w_H \geq p_S w_S$ and $(1 + p_S - p_H) > 1$ (because $p_S > p_H$), we see that the last inequality holds. Thus, we have shown that if $p_H w_H \geq p_S w_S$, Martina should serve hard on both serves. This is reasonable, because if it is optimal to serve hard on the second (and this requires $p_H w_H \geq p_S w_S$), then it should be optimal to serve hard on the first serve, because the danger of double-faulting (which is the drawback to the hard serve) is less immediate on the first serve. Of course, Example 4 could have been solved using a decision tree; see Problem 10 of Section 13.4.

In our solution to Example 4, we have shown how Martina's optimal strategy depends on the values of the parameters defining the problem. This is a kind of sensitivity analysis like the one applied to linear programming problems in Chapters 5 and 6.

PROBLEMS

Group A

1 Vladimir Ulanowsky is playing Keith Smithson in a two-game chess match. Winning a game scores 1 match point, and drawing a game scores $\frac{1}{2}$ match point. After the two games are played, the player with more match points is declared the champion. If the two players are tied after two games, they continue playing until someone wins a game (the winner of that game will be the champion). During each game, Ulanowsky can play one of two ways: boldly or conservatively. If he plays boldly, he has a 45% chance of winning the game and a 55% chance of losing the game. If he plays conservatively, he has a 90% chance of drawing the game and a 10% chance of losing the game. Ulanowsky's goal is to maximize his probability of winning the match. Use dynamic programming to help him accomplish this goal. If this problem is solved correctly, even though Ulanowsky is the inferior player, his chance of winning the match is over $\frac{1}{2}$. Explain this anomalous result.

2 Dickie Hustler has \$2 and is going to toss an unfair coin (probability .4 of heads) three times. Before each toss, he can bet any amount of money (up to what he now has). If heads comes up, Dickie wins the number of dollars he bets; if tails comes up, he loses the number of dollars he bets. Use dynamic programming to determine a strategy that maximizes Dickie's probability of having at least \$5 after the third coin toss.

Group B

3 Supppose that Army trails by 14 points in the Army–Navy football game. Army's guardian angel has assured the Army coach that his team will have the ball two more times during the game and will score a touchdown (worth 6 points) each time it has the ball. The Army coach has also been assured that Navy will not score any more points. Suppose a win is assigned a value of 1, a tie is .3, and a loss is 0. Army's problem is to determine whether to go for 1 or 2 points after each touchdown. A 1-point conversion is always successful, and a 2-point conversion is successful only 40% of the time. The Army coach wants to maximize the expected reward earned from the outcome of the game. Use dynamic programming to determine an optimal strategy. Then prove the following result: *No matter what value is assigned to a tie, it is never optimal to use the following strategy: Go for a 1-point conversion after the first touchdown and go for a 2-point conversion after the second touchdown.* Note that this (suboptimal) strategy is the one most coaches follow!

19.4 Further Examples of Probabilistic Dynamic Programming Formulations

Many probabilistic dynamic programming problems can be solved using recursions of the following form (for max problems):

$$f_t(i) = \max_a \left\{(\text{expected reward during stage } t|i, a) + \sum_j p(j|i, a, t)f_{t+1}(j)\right\} \tag{9}$$

In (9), $f_t(i)$ is the maximum expected reward that can be earned during stages $t, t+1, \ldots$ end of the problem, given that the state at the beginning of stage t is i. The max in (9) is taken over all actions a that are feasible when the state at the beginning of stage t is i. In (9), $p(j|i, a, t)$ is the probability that the next period's state will be j, given that the current (stage t) state is i and action a is chosen. Hence, the summation in (9) represents the expected reward from stage $t + 1$ to the end of the problem. By choosing a to maximize the right-hand side of (9), we are choosing a to maximize the expected reward earned from stage t to the end of the problem, and this is what we want to do. The following are six examples of probabilistic dynamic programming formulations.

EXAMPLE 5 Sunco Oil Drilling

Sunco Oil has D dollars to allocate for drilling at sites $1, 2, \ldots, T$. If x dollars are allocated to site t, the probability is $q_t(x)$ that oil will be found on site t. Sunco estimates that if site t has any oil, it is worth r_t dollars. Formulate a recursion that could be used to enable Sunco to maximize the expected value of all oil found on sites $1, 2, \ldots, T$.

Solution This is a typical resource allocation problem (see Example 1). Therefore, the stage should represent the number of sites, the decision for site t is how many dollars to allocate to site t, and the state is the number of dollars available to allocate to sites $t, t+1, \ldots, T$. We therefore define $f_t(d)$ to be the maximum expected value of the oil that can be found on sites $t, t+1, \ldots, T$ if d dollars are available to allocate to sites $t, t+1, \ldots, T$.

We make the reasonable assumption that $q_T(x)$ is a nondecreasing function of x. If this is the case, then at stage T, all the money should be allocated to site T. This yields

$$f_T(d) = r_T q_T(d) + (1 - q_T(d))0 = r_T q_T(d)$$

For $t < T$,

$$f_t(d) = \max_x \{r_t q_t(x) + f_{t+1}(d - x)\}$$

where x must satisfy $0 \leq x \leq d$. The last recursion follows, because $r_t q_t(x)$ is the expected value of the reward for stage t, and since Sunco will have $d - x$ dollars available for sites

5 At the beginning of each week, a machine is either running or broken down. If the machine runs throughout the week, it earns revenues of \$100. If the machine breaks down during a week, it earns no revenue for that week. If the machine is running at the beginning of the week, we may perform maintenance on it to lessen the chance of a breakdown. If the maintenance is performed, a running machine has a .4 chance of breaking down during the week; if maintenance is not performed, a running machine has a .7 chance of breaking down during the week. Maintenance costs \$20 per week. If the machine is broken down at the beginning of the week, it must be replaced or repaired. Both repair and replacement occur instantaneously. Repairing a machine costs \$40, and there is a .4 chance that the repaired machine will break down during the week. Replacing a broken machine costs \$90, but the new machine is guaranteed to run throughout the next week of operation. Use dynamic programming to determine a repair, replacement, and maintenance policy that maximizes the expected net profit earned over a four-week period. Assume that the machine is running at the beginning of the first week.

6 I own a single share of Wivco stock. I must sell my share at the beginning of one of the next 30 days. Each day, the price of the stock changes. With probability $q(x)$, the price tomorrow will increase by x% over today's stock price (x can be negative). For example, with probability $q(5)$, tomorrow's stock price will be 5% higher than today's. Show how dynamic programming can be used to determine a strategy that maximizes the expected revenue earned from selling the share of Wivco stock. Assume that at the beginning of the first day, the stock sells for \$10 per share.

Group B

7 The National Cat Foundling Home encourages people to adopt its cats, but (because of limited funds) it allows each prospective owner to inspect only four cats before choosing one of them to take home. Ten-year-old Sara is eager to adopt a cat and agrees to abide by the following rules. A randomly selected cat is brought for Sara to see, and then Sara must either choose the cat or reject it. If the first cat is rejected, Sara sees another randomly selected cat and must accept or reject it. This procedure continues until Sara has selected her cat. Once Sara rejects a cat, she cannot go back later and choose it as her pet. Determine a strategy for Sara that will maximize her probability of ending up with the cat she actually prefers.

8 Consider the following probabilistic inventory model:

a At the beginning of each period, a firm observes its inventory position.

b Then the firm decides how many units to produce during the current period. It costs $c(x)$ dollars to produce x units during a period.

c With probability $q(d)$, d units are demanded during the period. From units on hand (including the current period's production), the firm satisfies as much of the demand as possible. The firm receives r dollars for each unit sold. For each unit of demand that is unsatisfied, a penalty cost p is incurred. All unsatisfied demand is assumed to be lost. For example, if the firm has 20 units available and current demand is 30, a revenue of $20r$ would be received, and a penalty of $10p$ would be incurred.

d If ending inventory is positive, a holding cost of \$1 per unit is incurred.

e The next period now begins.

The firm's inital inventory is zero, and its goal is to minimize the expected cost over a 100-period horizon. Formulate a dynamic programming recursion that will help the firm accomplish its goal.

9 Martha and Ken Allen want to sell their house. At the beginning of each day, they receive an offer. We assume that from day to day, the sizes of the offers are independent random variables and that the probability that a given day's offer is for j dollars is p_j. An offer may be accepted during the day it is made or at any later date. For each day the house remains unsold, a maintenance cost of c dollars is incurred. The house must be sold within 30 days. Formulate a dynamic programming recursion that Martha and Ken can use to maximize their expected net profit (selling price − maintenance cost). Assume that the maintenance cost for a day is incurred before the current day's offer is received and that each offer is for an integer number of dollars.

10 An advertising firm has D dollars to spend on reaching customers in T separate markets. Market t consists of k_t people. If x dollars are spent on advertising in market t, the probability that a given person in market t will be reached is $p_t(x)$. Each person in market t who is reached will buy c_t units of the product. A person who is not reached will not buy any of the product. Formulate a dynamic programming recursion that could be used to maximize the expected number of units sold in T markets.

11 Georgia Stein is the new owner of the New York Yankees. Each season, Georgia must decide how much money to spend on the free agent draft. During each season, Georgia can spend any amount of money on free agents up to the team's capital position at the beginning of the season. If the Yankees finish in ith place during the season, their capital position increases by $R(i)$ dollars less the amount of money spent in the free agent draft. If the Yankees finished in ith place last season and spend d dollars on free agents during the off-season, the probability that the Yankees will finish in place j during the next season is $p_{ij}(d)(j = 1, 2, \dots, 7)$. Last season, the Yankees finished in first place, and at the end of the season, they had a capital position of D dollars. Formulate a dynamic programming recursion that will enable the Yankees to maximize their expected cash position at the end of T seasons.

12 Bailey Bliss is the campaign manager for Walter Glenn's presidential campaign. He has D dollars to allocate to T winner-take-all primaries. If x_t dollars are allocated to primary t, then with probability $p_t(x_t)$, Glenn will win primary t and obtain v_t delegates. With probability $1 - p_t(x_t)$, Glenn loses primary t and obtains no delegates. Glenn needs K delegates to be nominated. Use dynamic programming to help Bliss maximize Glenn's probability of being nominated. What aspect of a real campaign does the present formulation ignore?

13 At 7 A.M., eight people leave their cars for repair at Harry's Auto Repair Shop. If person i's car is ready by time

Linear Programming

In a maximization problem, $V(i)$ for each state may be determined by solving the following LP:

$$\min z = V_1 + V_2 + \cdots + V_N$$

$$\text{s.t.} \quad V_i - \beta \sum_{j=1}^{j=N} p(j|i, d)V_j \geq r_{id} \qquad \text{(For each state } i \text{ and each } d \in D(i))$$

All variables urs

If the constraint for state i and decision d has no slack, then decision d is optimal in state i.

Value Iteration, or Successive Approximations

Let $V_t(i)$ be the maximum expected discounted reward that can be earned during t periods if the state at the beginning of the current period is i. Then

$$V_t(i) = \max_{d \in D(i)} \left\{ r_{id} + \beta \sum_{j=1}^{j=N} p(j|i, d)V_{t-1}(j) \right\} \qquad (t \geq 1)$$

$$V_0(i) = 0$$

As t grows large, $V_t(i)$ will approach $V(i)$. For t sufficiently large, the decision that is optimal in state i for a t-period problem is also optimal in state i for an infinite-horizon problem.

REVIEW PROBLEMS

Group A

1 A company has five sales representatives available for assignment to three sales districts. The sales in each district during the current year depend on the number of sales representatives assigned to the district and on whether the national economy has a bad or good year (see Table 12). In the Sales column for each district, the first number represents sales if the national economy had a bad year, and the second number represents sales if the economy had a good year. There is a .3 chance that the national economy will have a good year and a .7 chance that the national economy will have a bad year. Use dynamic programming to determine an assignment of sales representatives to districts that maximizes the company's expected sales.

TABLE 12

No. of Sales Reps Assigned to District	Sales (millions) District 1	District 2	District 3
0	\$1, \$4	\$2, \$5	\$3, \$4
1	\$2, \$6	\$4, \$6	\$5, \$5
2	\$3, \$7	\$5, \$6	\$6, \$7
3	\$4, \$8	\$6, \$6	\$7, \$7

2 At the beginning of each period, a company must determine how many units to produce. A setup cost of \$5 is incurred during each period in which production takes place. The production of each unit also incurs a \$2 variable cost. All demand must be met on time, and there is a \$1 per-unit holding cost on each period's ending inventory. During each period, it is equally likely that demand will equal 0 or 1 unit. Assume that each period's ending inventory cannot exceed 2 units.

a Use dynamic programming to minimize the expected costs incurred during three periods. Assume that the initial inventory is 0 units.

b Now suppose that each unit demanded can be sold for \$4. If the demand is not met on time, the sale is lost. Use dynamic programming to maximize the expected profit earned during three periods. Assume that the initial inventory is 0 units.

c In parts (a) and (b), is an (s, S) policy optimal?

3 At Hot Dog Queen Restaurant, the following sequence of events occurs during each minute:

a With probability p, a customer arrives and waits in line.

b Hot Dog Queen determines the rate s at which customers are served. If any customers are in the restaurant, then with probability s, one of the customers completes