TABLE 9

Last Month's Sales	Current Month's Sales Good	Bad
Good	.95	.05
Bad	.40	.60

TABLE 10

Last Month's Sales	Current Month's Sales Good	Bad
Good	.80	.20
Bad	.20	.80

expected discounted profit over an infinite horizon (use $\beta = .98$).

a Use the policy iteration method to determine an optimal stationary policy.

b Use linear programming to determine an optimal stationary policy.

c Perform two iterations of value iteration.

d Find a policy that maximizes average profit per month.

3 Suppose you are using the policy iteration method to determine an optimal policy for an MDP. How might you use LINDO to solve the value determination equations?

Group B

4 During any day, I may own either 0 or 1 share of a stock. The price of the stock is governed by the Markov chain shown in Table 11. At the beginning of a day in which I own a share of stock, I may either sell it at today's price or keep it. At the beginning of a day in which I don't own a share of stock, I may either buy a share of stock at today's price

TABLE 11

Today's Price	Tomorrow's Price $0	$1	$2	$3
$0	.5	.3	.1	.1
$1	.1	.5	.2	.2
$2	.2	.1	.5	.2
$3	.1	.1	.3	.5

or not buy a share. My goal is to maximize my expected discounted profit over an infinite horizon (use $\beta = .95$).

a Use the policy iteration method to determine an optimal stationary policy.

b Use linear programming to determine an optimal stationary policy.

c Perform two iterations of value iteration.

d Find a policy that maximizes average daily profit.

5 Ethan Sherwood owns two printing presses, on which he prints two types of jobs. At the beginning of each day, there is a .5 probability that a type 1 job will arrive, a .1 probability that a type 2 job will arrive, and a .4 probability that no job will arrive. Ethan receives $400 for completing a type 1 job and $200 for completing a type 2 job. (Payment for each job is received in advance.) Each type of job takes an average of three days to complete. To model this, we assume that each day a job is in press there is a $\frac{1}{3}$ probability that its printing will be completed at the end of the day. If both presses are busy at the beginning of the day, any arriving job is lost to the system. The crucial decision is when (if ever) Ethan should accept the less profitable type 2 job. Ethan's goal is to maximize expected discounted profit (use $\beta = .90$).

a Use the policy iteration method to determine an optimal stationary policy.

b Use linear programming to determine an optimal stationary policy.

c Perform two iterations of value iteration.

SUMMARY

Key to Formulating Probabilistic Dynamic Programming Problems (PDPs)

Suppose the possible states during period $t + 1$ are $s_1, s_2, \ldots s_n$, and the probability that the period $t + 1$ state will be s_i is p_i. Then the minimum expected cost incurred during periods $t + 1, t + 2, \ldots,$ end of the problem is

$$\sum_{i=1}^{i=n} p_i f_{t+1}(s_i)$$

where $f_{t+1}(s_i)$ is the minimum expected cost incurred from period $t + 1$ to the end of the problem, given that the state during period $t + 1$ is s_i.