

Linear Programming II: Minimization

© 2006 Samuel L. Baker

Assignment 11 is on page 16.

Introduction

A minimization problem minimizes the value of the objective function rather than maximizing it. Minimization problems generally involve finding the least-cost way to meet a set of requirements.

Classic example -- feeding farm animals.

Animals need:

14 units of nutrient A,

12 units of nutrient B, and

18 units of nutrient C.

Two feed grains are available, X and Y.

A bag of X has 2 units of A, 1 unit of B, and 1 unit of C.

A bag of Y has 1 unit of A, 1 unit of B, and 3 units of C.

A bag of X costs \$2. A bag of Y costs \$4.

Minimize the cost of meeting the nutrient requirements.

To solve, express the problem in equation form:

Cost = $2X + 4Y$ objective function to be minimized

Constraints:

$2X + 1Y \geq 14$ nutrient A requirement

$1X + 1Y \geq 12$ nutrient B requirement

$1X + 3Y \geq 18$ nutrient C requirement

↑ ↑

Read vertically to see how much of each nutrient is in each grain.

$X \geq 0, Y \geq 0$ non-negativity

Learning Objective 1: Recognize problems that linear programming can handle.

Linear programming lets you optimize an objective function subject to some constraints. The objective function and constraints are all linear.


Learning objective 2: Know the elements of a linear programming problem -- what you need to calculate a solution.

The elements are

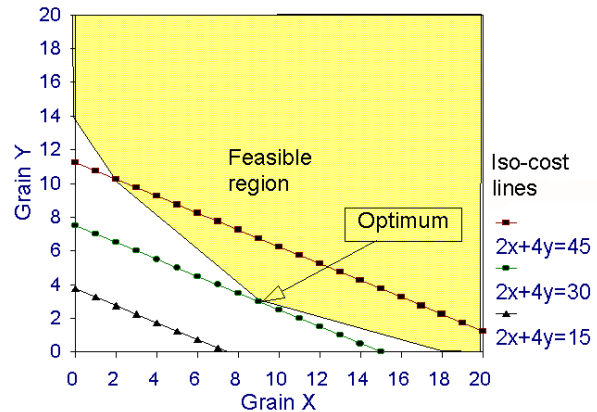
- (1) an objective function that shows the cost or profit depending on what choices you make,
- (2) constraint inequalities that show the limits of what you can do, and
- (3) non-negativity restrictions, because you cannot turn outputs back into inputs.

Graph method of solution

Graph the constraints as equalities, like before. The constraints are now \geq rather than \leq , so the feasible area is everything to the right and above all of the constraint lines. You want to find the lowest cost point of this area. It will be a corner. You superimpose isocost lines that you get from the objective function. When you find an isocost line that just nicks one corner of the area, that corner is the solution.



Feasible region in minimization problem



Isocost lines and optimum

Learning objective 3: Understand the principles that the computer uses to solve a linear programming problem.

The computer uses the simplex method to systematically trace around the feasible area (the simplex) in the figure above. It goes from one corner to the next until it finds the best value for the objective function.

Learning objective 4a: What linear programming problems have no solution?

Those that have no feasible area. This means it's impossible to satisfy all the constraints at once.

Learning objective 4b: What difference does linearity make?

If the constraints are not linear, the feasible area has curved edges. The simplex method doesn't work because you can't be sure that the solution is at a corner. The solution may be in the middle of a curved edge. If the objective function is not linear, the solution may not even be on an edge. It may be in the interior of the feasible area.

The simplex method of solution for minimization problems

The simplex method gets more complicated when you have constraints that are \geq . The method can't start at the origin (with $X = 0, Y = 0$, etc.) because that is not a feasible solution. To find a starting place, we invent some new variables in addition to the slack variables. The new variables are called "artificial" and are conventionally labelled with A's. Artificial variables allow us to automatically come up with a feasible solution to start the simplex method. In this case, the starting feasible solution is $A_1=14, A_2=12, A_3=18$, and all the rest (X, Y , and the S's) =0.

The equations now look like this.

$$\begin{aligned} 2X + 1Y - 1S_1 + 0S_2 + 0S_3 + 1A_1 + 0A_2 + 0A_3 &= 14 \\ 1X + 1Y + 0S_1 - 1S_2 + 0S_3 + 0A_1 + 1A_2 + 0A_3 &= 12 \\ 1X + 3Y + 0S_1 + 0S_2 - 1S_3 + 0A_1 + 0A_2 + 1A_3 &= 18 \end{aligned}$$

Each constraint equation has one slack variable (S) and one artificial variable (A) with non-zero coefficients. The slack variable in each equation has a -1 coefficient because in a minimization problem we have to meet or exceed each constraint. This means we have to subtract the slack variable rather than add it as we did when we were maximizing the objective function.

The objective function, which we want to minimize, is expanded to:

$$2X + 4Y + 0S_1 + 0S_2 + 0S_3 + MA_1 + MA_2 + MA_3 = \text{Cost}$$

M stands for a mammoth number, meaning a very large unit cost for the A's, large enough to be sure that none of the A's will be in the final solution.

In matrix form we have:

X	Y	S ₁	S ₂	S ₃	A ₁	A ₂	A ₃	
2	1	-1	0	0	1	0	0	14
1	1	0	-1	0	0	1	0	12
1	3	0	0	-1	0	0	1	18
-2	-4	0	0	0	-M	-M	-M	0

The first thing we do is to get rid of those -M's that are at the bottom of the three columns for the A's. To do this, multiply each of the constraint rows by M and add, element by element, to the last row.

2	1	-1	0	0	1	0	0	14
1	1	0	-1	0	0	1	0	12
1	3	0	0	-1	0	0	1	18
4M-2	5M-4	-M	-M	-M	0	0	0	44M

The columns for the A's now have one 1 and the rest 0's. The A's, and only them, are in the initial basic feasible solution. The solution is $A_1 = 14, A_2 = 12, A_3 = 18$, all X's and S's = 0. Cost is 44M, a monstrously mammoth number that we want to reduce. We then proceed with the simplex method as before. The pivot column is the second, because it has the biggest number (5M-4) in the bottom row.

2	1	-1	0	0	1	0	0	14
1	1	0	-1	0	0	1	0	12
1	3	0	0	-1	0	0	1	18
4M-2	5M-4	-M	-M	-M	0	0	0	44M

To find the pivot row, we divide each number in the last column by the corresponding number in the second column, to get $14/1=14$, $12/1=12$, and $18/3=6$. 6 is the smallest so the third row is the pivot row. The pivot element is the 3.

The next step is to use that third row to eliminate the other elements in the second column. We multiply the row by $1/3$ and subtract that from the first row. We do the same for the second row. For the bottom row we multiply the third row by $(4/3 - 5M/3)$ and subtract that from the fourth row. For our new third row, we divide the original third row by 3. We get:

X	Y	S ₁	S ₂	S ₃	A ₁	A ₂	A ₃	
1.67	0	-1	0	0.33	1	0	-0.33	8
0.67	0	0	-1	0.33	0	1	-0.33	6
0.33	1	0	0	-0.33	0	0	0.33	6
2.33M-0.67	0	-M	-M-1.33	0.67M	0	0	0	14M+24

This puts Y in the solution, along with A1 and A3. The next iteration adds X to the solution. The iteration after that adds S₁. The final tableau looks like this:

X	Y	S ₁	S ₂	S ₃	A ₁	A ₂	A ₃	
1	0	0	-1.5	0.5	0	1.5	-0.5	9
0	0	1	-2.5	0.5	-1	2.5	-0.5	7
0	1	0	0.5	-0.5	0	-0.5	0.5	6
0	0	0	-1	-1	-M	1-M	1-M	30

This is the final tableau because all the numbers in the bottom row are 0 or negative. Notice that the only M's left are under the A's.

The non-zero variables in the solution are shown by the columns that have one 1 and the rest 0's. The solution amounts are in the right column, and match up with X, Y, and S₁ according to where the 1 in each column is. The solution is: X = 9, Y = 3, S₁ = 7. S₂ and S₃ are not in the final solution, so S₂ = S₃ = 0. The total cost is \$30.

The shadow prices for the requirements are in the bottom row, in columns for S₂ and S₃. Both S₂ and S₃ have a shadow price of -\$1. These mean that it would cost you \$1 if you had to provide one more unit of either nutrient B or nutrient C. Having to provide one more unit of nutrient A costs nothing because you are already providing 7 extra units of A.

Minimizing with Excel

Fortunately, you don't have to do the simplex method by hand anymore. You can get Excel or Quattro Pro to do it for you. The challenge becomes setting the problem up properly.

Learning objective 5: Be able to solve small linear programming problems yourself.

The tricky part is setting up the objective function and the constraints. Write them on paper (without the S's and A's). Then let your spreadsheet crunch them for you.

Here is the minimization example, as set up in Excel:

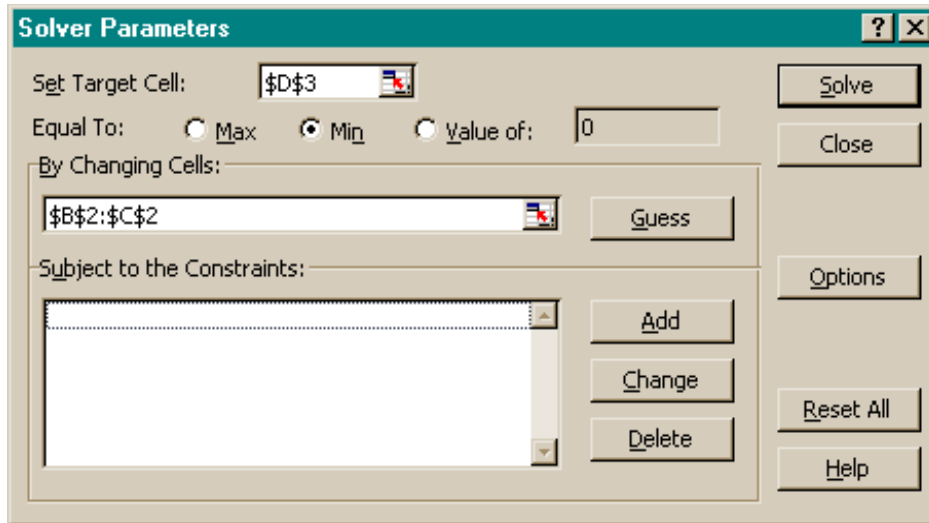
	A	B	C	D	E
1		Grain X	Grain Y	Total	
2	Amount to use	0	0		
3	Cost per bag	2	4	0	
4					
5	Constraints				
6	nutrient A	2	1	0	14
7	nutrient B	1	1	0	12
8	nutrient C	1	3	0	18

Column B is for Grain X. Column C is for Grain Y. Column D has formulas for total cost and total amounts of nutrients provided. If you type the formula in D3 first, and use \$ signs in the right places, you can copy D3 to D6:D8.

	A	B	C	D	E
1		Grain X	Grain Y	Total	
2	Amount to use	0	0		
3	Cost per bag	2	4	=SUMPRODUCT(\$B\$2:\$C\$2,B3:C3)	
4					
5	Constraints				
6	nutrient A	2	1	=SUMPRODUCT(\$B\$2:\$C\$2,B6:C6)	14
7	nutrient B	1	1	=SUMPRODUCT(\$B\$2:\$C\$2,B7:C7)	12
8	nutrient C	1	3	=SUMPRODUCT(\$B\$2:\$C\$2,B8:C8)	18

Here's the same spreadsheet, showing the formulas in the D column.

Select Tools and then Solver... from Excel's menu.



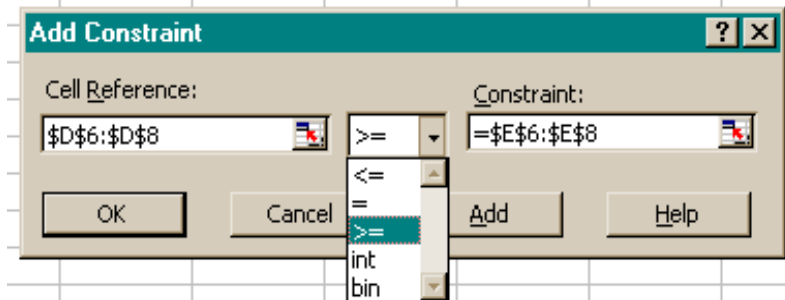
If you are re-using a spreadsheet from an earlier assignment, the safest way to start is to click on the Reset All button. This clears all the old settings and cell choices, preventing some elusive bugs.

Then set the Target Cell to D3 and the Changing Cells to B2:C2.

On the Equal To: line, check Min. If you don't do this, your solution will be unbounded. Solver will fly off looking for the upper-right corner of your boundless feasible region.


Click on the Add button to add constraints.

The nutrient amounts our grain provides, in D6:D8, must be greater than or equal to the required nutrients in E6:E8. Be sure to change the relation in the center to >=.

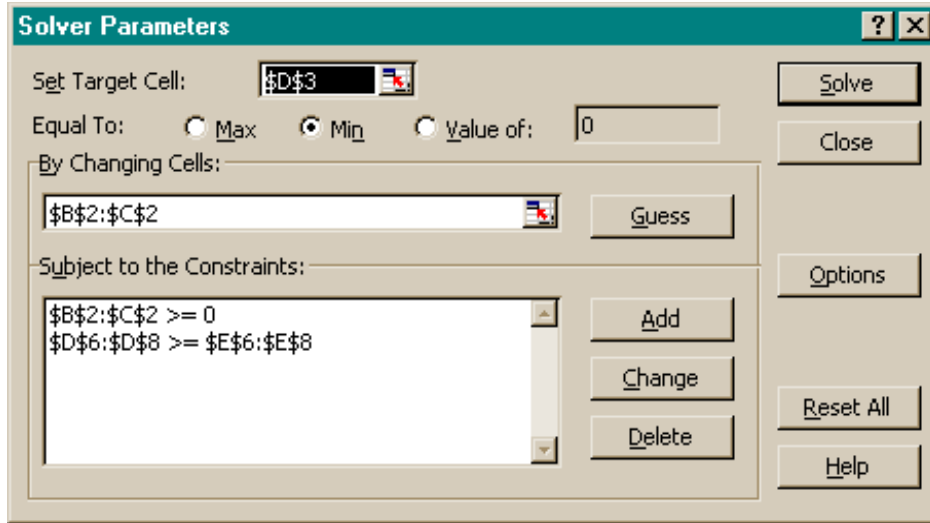


Then click on this box's Add button to save this constraint and get ready to add another one.

This time, it's the non-negativity constraints. B2 and C2 must be >= 0.



Click on OK once this box is filled in as shown, because we're done adding constraints.

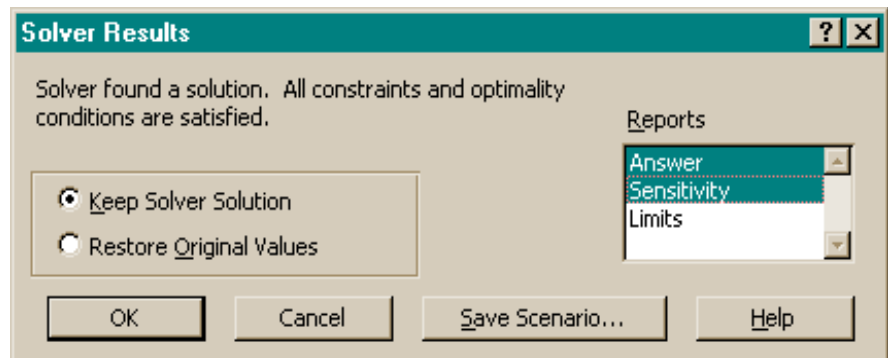


Here's what the Solver Parameters box should look like when you're ready to solve.

Click on Solve.

Done! A Solver Results dialog box pops up in front of your spreadsheet.

In the Reports section of the Solver Results box, select the Answer and Sensitivity Reports. Then click OK.



You will then see that Excel has filled in the best solution, as shown to the right.

This says that we satisfy the requirements for the least possible cost with 9 bags of X and 3 bags of Y. The cost is \$30.

	A	B	C	D	E
1		Grain X	Grain Y	Total	
2	Amount to use	9	3		
3	Cost per bag	2	4	30	
4					
5	Constraints				
6	nutrient A	2	1	21	14
7	nutrient B	1	1	12	12
8	nutrient C	1	3	18	18

Click on the Sensitivity Report tab. This time the shadow prices are called “Lagrange Multipliers.” Why not “Shadow Prices”? My bad! I did not tell you to click the Solver’s Options button and check the Assume Linear box. It does not really matter. The numbers are the same. Only the name is different.

Anyway, the Lagrange Multipliers, which are the Shadow Prices, tell us that if we had to provide another unit of nutrient B, it would cost \$1 more. Providing another unit of C would also cost \$1 more. In other words, if vitamin B or vitamin C pills were available, we would buy them if the price was less than \$1 per unit. We would not buy any vitamin A pills, because the shadow price for A is \$0. The constraint for A is not binding.

Adjustable Cells			
Cell	Name	Final Value	Reduced Gradient
\$B\$2	Amount to use Grain X	9	0
\$C\$2	Amount to use Grain Y	3	0

Constraints			
Cell	Name	Final Value	Lagrange Multiplier
\$D\$6	nutrient A Total	21	0
\$D\$7	nutrient B Total	12	1
\$D\$8	nutrient C Total	18	1

There are two more shadow prices hiding in the Sensitivity Report. They are under the “Reduced Gradient” heading above. In this example, they tell us how much the non-negativity constraints are costing us. Which is nothing, because we would not want to turn nutrients into grain the way these numbers are. In one of the homework problems, a Reduced Gradient number may be important.

Modifications and extensions

A. Multiple Optimal Solutions

If a constraint happens to be parallel to the isocost lines, there will be multiple solutions, if that constraint turns out to be binding. I can convert this example into one that multiple solutions by making the prices of both grains the same. This makes the isocost lines parallel to the nutrient B constraint.


	A	B	C	D	E	F
1		Grain X	Grain Y	Total		
2	Amount to use	0	0			
3	Cost per bag	3	3	0		
4						These two
5	Constraints					are parallel.
6	nutrient A	2	1	0	14	
7	nutrient B	1	1	0	12	
8	nutrient C	1	3	0	18	

Now both grains cost the same, \$3 per bag, and the two grains have the same amount of nutrient B.

When the cost per unit of a nutrient is the same for two grains, the isocost lines are parallel to the constraint for that nutrient.

The shading in this diagram is trying to show that the nutrient B constraint is parallel to the cost objective function.

One edge of the feasible region is now parallel to the iso-cost lines. The diagram to the right shows how the lowest iso-cost line that touches the feasible region touches it along a whole line segment, not just at one corner.



Excel’s solver does not directly tell you that there are multiple solutions. Instead, it seems to pick one solution arbitrarily. When I solved this with Excel, I got an arbitrary value near the middle of this line segment. This was a true solution, but it was only one of many. Any combination of X and Y that sum to 12 from (2,10) to (9,3) is an optimum.

I solved again in Excel, this time first clicking the Options button and checking the Assume Linear Model box. Excel came back with the (9,3) solution, which is at one end of the line segment of optimal solutions.

The sensitivity report, shown here, does have a subtle indication that something is up: The Shadow Price for nutrient C is \$0, even though there is no slack in nutrient C, the amount of C provided being equal to the amount required.

Adjustable Cells

Cell	Name	Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
\$B\$2	Amount to use Grain X	9	0	3	0	2
\$C\$2	Amount to use Grain Y	3	0	3	6	0

Constraints

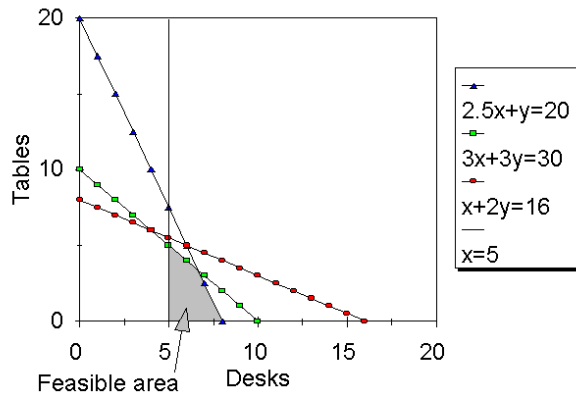
Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
\$D\$6	nutrient A Total	21	0	14	7	1E+30
\$D\$7	nutrient B Total	12	3	12	6	2.8
\$D\$8	nutrient C Total	18	0	18	14	6

The \$0 shadow price for C means that the C constraint is not really active at the (9,3) solution. Ordinarily in a linear programming problem, the number of active constraints (constraints with positive shadow prices) equals the number of variables. We have two variables, X and Y, so ordinarily we would have two active constraints. The fact that we have only one active constraint indicates that we have parallelism, and that the solution displayed is one among many.

B. Mixing \geq and \leq constraints

Actually, we're already doing this with the non-negativity constraints, but here's a further use of the same technique.

Suppose in the desks and tables problem from last week that we have an order for 5 desks that we are obligated to fill. We have to produce at least 5 desks. The effect of adding this constraint is to cut off part of the feasible area, from the left.



Adding constraint that $x \geq 5$

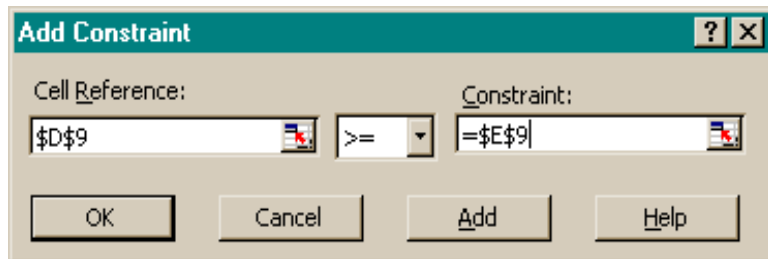
The new optimum turns out to be 5 desks and 5 tables. Profit is \$35. Without the added constraint, 4 desks and 6 tables were produced, and the profit was \$36. The added constraint costs \$1 in lost profit.

Doing this in Excel requires adding another constraint. The best way to do this is to add another constraint row to the spreadsheet. Being obligated to supply 5 desks is expressed as a constraint that $1 \times \text{Desks} + 0 \times \text{Tables} \geq 5$. The 1 and 0 in that equation go in cells B9 and C9. Copy the formula in D8 and paste it to D9.

	A	B	C	D	E
1		Desks	Tables	Total	
2	Amount to make	0	0		
3	Unit profit	3	4	0	
4					
5	Constraints				
6	assembling	2.5	1	0	20
7	buffing	3	3	0	30
8	crating	1	2	0	16
9	obligated	1	0	0	5

The added “obligated” constraint.

Then, in Solver, add the new constraint that $D9 \geq E9$. Be sure to get the greater-than-or-equal-to sign showing in the middle of the Add Constraint dialog box.



Alternatively, you could just add a new constraint that $B2 \geq 5$. The problem

with this is that this new constraint is not visible unless you bring up the Solver dialog box. Six months from now, you might come back to your spreadsheet and not remember that you added this constraint.

After Solving, the Sensitivity Report shows a shadow price for the “obligated” constraint of -1. This means that every extra desk you are obligated to make *reduces* your profit by \$1. This is consistent with

the finding that adding the obligated constraint made the profit \$35, instead of the \$36 you made without the new constraint. (Another advantage of putting this new constraint in the 9th row, instead of just adding a constraint to Solver that $B2 \geq 5$, is that the shadow price shows up properly labeled in the Sensitivity Report.)

Synonyms for Shadow Price

Depending on how you which software you are using, whether you specify a linear model, and how you handle constraints on single amounts (like the $B2 \geq 5$ above), the what-it's-worth-to-you-for-another-unit prices may also be called:

- ▶ Shadow Prices
- ▶ Lagrange Multipliers
- ▶ Opportunity Cost
- ▶ Reduced Cost
- ▶ Reduced Gradient

You may have look up and down the Sensitivity Report to find all the shadow prices in their various guises.

Applications

The computer takes care of the math, so the hard part is figuring out how to set the problem up.

A. Staff Scheduling

Suppose you have these weekly staffing requirements:

Day	Staff Needed
Mon	180
Tue	160
Wed	150
Thu	160
Fri	190
Sat	140
Sun	120

Each employee works 5 days in a row, then he/she is off for two days. Employees may start on any day of the week. Each employee costs \$400 per week. The objective is to minimize the total cost of meeting the requirements.

Here is one possible spreadsheet setup. It uses 0's and 1's to put in the constraints:

	A	B	C	D	E	F	G	H	I	J
1		Mon-Fri	Tue-Sat	Wed-Sun	Thu-Mon	Fri-Tue	Sat-Wed	Sun-Thu	Total	
2	Amounts	0	0	0	0	0	0	0		
3	Cost	400	400	400	400	400	400	400	0	
4										
5	Constraints									
6	Mon	1	0	0	1	1	1	1	0	180
7	Tue	1	1	0	0	1	1	1	0	160
8	Wed	1	1	1	0	0	1	1	0	150
9	Thu	1	1	1	1	0	0	1	0	160
10	Fri	1	1	1	1	1	0	0	0	190
11	Sat	0	1	1	1	1	1	0	0	140
12	Sun	0	0	1	1	1	1	1	0	120

	I
1	Total
2	
3	=SUMPRODUCT(\$B\$2:\$H\$2,B3:H3)
4	
5	
6	=SUMPRODUCT(\$B\$2:\$H\$2,B6:H6)
7	=SUMPRODUCT(\$B\$2:\$H\$2,B7:H7)
8	=SUMPRODUCT(\$B\$2:\$H\$2,B8:H8)
9	=SUMPRODUCT(\$B\$2:\$H\$2,B9:H9)
10	=SUMPRODUCT(\$B\$2:\$H\$2,B10:H10)
11	=SUMPRODUCT(\$B\$2:\$H\$2,B11:H11)
12	=SUMPRODUCT(\$B\$2:\$H\$2,B12:H12)

To get the logic behind this spreadsheet, consider row 6, for example. It's the constraint for Monday. The total number of people on the job on Monday equals the number of people who work Monday-Friday, plus Thursday-Monday, Friday-Tuesday, Saturday-Wednesday, and Sunday-Thursday. Row 6 has 1's in those columns. People who work Tuesday-Saturday and Wednesday-Sunday are not working Monday, so Row 6 has 0's in those columns. Cell I6's SUMPRODUCT formula calculates the total number of people working Monday. It must be greater than or equal to what is in J6.

An alternative way to set up the problem is:

	A	B	C	D	E	F
1	Schedule	Number	Cost	Day	Have	Need
2	Mon-Fri	0	400	Mon	=SUM(B2,B5:B8)	180
3	Tue-Sat	0	400	Tue	=SUM(B2:B3,B6:B8)	160
4	Wed-Sun	0	400	Wed	=SUM(B2:B4,B7:B8)	150
5	Thu-Mon	0	400	Thu	=SUM(B2:B5,B8)	160
6	Fri-Tue	0	400	Fri	=SUM(B2:B6)	190
7	Sat-Wed	0	400	Sat	=SUM(B3:B7)	140
8	Sun-Thu	0	400	Sun	=SUM(B4:B8)	120
9		Total	=SUMPRODUCT(B2:B8,C2:C8)			

The formulas in Column E do what the first layout does with 0's and 1's. The idea behind them is, for example, that the number of staff working Monday equals total number of people who started their 5-day work week on Monday, Thursday, Friday, Saturday, and Sunday. The constraints are that each of the formulas in the E column has to be greater than or equal to the corresponding Need amount in the F

column.

Both methods also need the non-negativity constraints, that the cells containing the amounts to hire are all ≥ 0 .

B. Transportation

You are distributing pharmaceuticals from two locations. Location A has 100 units. Location B has 250 units. You have customers at locations C, D, and E, who want 50, 120, and 180 units respectively. Minimize cost, based on these unit transportation costs:

	From: A	B
To:		
C	\$5	\$3
D	\$10	\$4
E	\$12	\$9

Here is a spreadsheet layout for this transportation problem:

	A	B	C	D	E	F	G	H	I
1	Routes:	AC	AD	AE	BC	BD	BE	Total	Limit
2	Amount	0	0	0	0	0	0		
3	Cost	5	10	12	3	4	9	=SUMPRODUCT(\$B\$2:\$G\$2,B3:G3)	
4	Constraints								
5	100 at A	1	1	1	0	0	0	=SUMPRODUCT(\$B\$2:\$G\$2,B5:G5)	100
6	250 at B	0	0	0	1	1	1	=SUMPRODUCT(\$B\$2:\$G\$2,B6:G6)	250
7	50 to C	1	0	0	1	0	0	=SUMPRODUCT(\$B\$2:\$G\$2,B7:G7)	50
8	120 to D	0	1	0	0	1	0	=SUMPRODUCT(\$B\$2:\$G\$2,B8:G8)	120
9	180 to E	0	0	1	0	0	1	=SUMPRODUCT(\$B\$2:\$G\$2,B9:G9)	180

The key insight is that the variables are the six possible routes, that is, the six possible combinations of the two sources and the three destinations.

The constraints are a mixture of \geq and \leq . Be careful when you enter these to be sure that they are all in the right directions. The constraints are:

- H5:H6 \leq I5:I6
- H7:H9 \geq I7:I9
- B2:G2 \geq 0

If Solver gives you an error message that the solution is unbounded, check the directions of your $>$ and $<$ signs. Also check that you did include the non-negativity constraint, that the amounts are all ≥ 0 .

The “dual” illustrated

Every linear programming problem has a "dual" -- a restatement of the problem that switches the roles of the constraints and the activities. The original problem is called the "primal."

Here is the desks and tables example, showing the primal and the dual. The numbers are the same in the primal and the dual, but the columns in one are the rows in the other, the constraint constants in one are the coefficients of the objective function in the other, and the directions of the inequalities are switched. In the primal version of the desks and tables problem we maximize the objective function. We minimize the objective function in the dual.

Primal:

	A	B	C	D	E
1		Desks	Tables	Total	Limit
2	Amounts	0	0		
3	Profit	3	4	0	
4					
5	Constraints				
6	assemble	2.5	1	0	20
7	buff	3	3	0	30
8	crate	1	2	0	16

The constraints in the primal are:

$$D6:D8 \leq E6:E8$$

$$B2:C2 \geq 0.$$

The objective is to maximize D3.

Dual:

	A	B	C	D	E	F
1		Assemble	Buff	Crate	Total	Limit
2	Values	0	0	0		
3	Available	20	30	16	0	
4	Constraints					
5	desks	2.5	3	1	0	3
6	tables	1	3	2	0	4

The constraints in the dual are:

$$E5:E6 \geq F5:F6$$

$$B2:D2 \geq 0.$$

The objective is to minimize E3.

In the dual you are doing some accounting. You are allocating the profit value of each finished product (desk or table) to the inputs used to make it. These inputs' accounting values tell you how much you should be willing to pay if you can buy more inputs. These accounting values are the shadow prices of the inputs. Because they come from the Dual, they are sometimes called “dual values.”

The $E5 \geq F5$ constraint in the dual says that you must choose values for the three inputs such that the total value of the inputs into a desk is at least as great as \$3, the value of a desk. The $E6 \geq F6$ constraint in the dual says that the values you put on those three inputs must also be such that the total value of the inputs that go into a table is at least as big as \$4, the value of a table. The dual's objective function says that you want to chose values that minimize the total value of your initial stock of available resources. A mind-bender, isn't it?

The solutions are on the next page.

The primal's solution:

	A	B	C	D	E
1		Desks	Tables	Total	Limit
2	Amounts	4	6		
3	Profit	3	4	36	
4					
5	Constraints				
6	assemble	2.5	1	16	20
7	buff	3	3	30	30
8	crate	1	2	16	16

The dual's solution:

	A	B	C	D	E	F
1		Assemble	Buff	Crate	Total	Limit
2	Values	0	0.667	1		
3	Available	20	30	16	36	
4	Constraints					
5	desks	2.5	3	1	3	3
6	tables	1	3	2	4	4

The primal's solution amounts (4 and 6) are the shadow prices in the dual. Likewise, the dual's solution values (0, .667, and 1) are the shadow prices in the primal.

The simplex method solves the primal and the dual simultaneously. This is how Excel gets the numbers it puts in the Sensitivity Report.

Your turn:

The primal's Solver Sensitivity Report

Adjustable Cells

Cell	Name	Final Value	Reduced Cost
\$B\$2	Amounts Desks	4	0
\$C\$2	Amounts Tables	6	0

Constraints

Cell	Name	Final Value	Shadow Price
\$D\$6	assemble Total	16	0
\$D\$7	buff Total	30	0.666667
\$D\$8	crate Total	16	1

The dual's Solver Sensitivity Report

Adjustable Cells

Cell	Name	Final Value	Reduced Cost
\$B\$2	Values Assemble	0	4
\$C\$2	Values Buff	0.667	0
\$D\$2	Values Crate	1	0

Constraints

Cell	Name	Final Value	Shadow Price
\$E\$5	desks Total	3	4
\$E\$6	tables Total	4	6

Assignment 11

Write out your answers in sentences. Show your work. Write out your equations for the constraints and objective function. Include a printout of your spreadsheet results or attach a copy of your spreadsheet file.

1. Use the graph method and then the computer to solve this one:

You are the hospital gardener. The flowers require 28 pounds of nitrogen, 30 pounds of phosphates, and 4 pounds of potash. You have two fertilizers to use. One kind costs \$4 per bag. It contains 4 pounds of nitrogen, 2 pounds of phosphates, and no potash per bag. The other kind costs \$5 per bag. It contains 2 pounds of nitrogen, 3 pounds of phosphates, and 1 pound of potash.

a. How many bags of each kind of fertilizer do you use to minimize your costs?

b. The CEO's brother owns the company that makes the first kind of fertilizer. You must buy at least 6 bags of it. Now how much of each kind of fertilizer do you use? (Be sure to click Solver's Options button and check Assume Linear Model.)

c. For b, what do the shadow prices mean? Explain all of them, including the shadow price for the constraint about having to buy 6 from the CEO's brother. (It may be called Shadow Price, Lagrange Multiplier, Reduced Cost, or Reduced Gradient, depending on how you set up your spreadsheet.)

2. You have to schedule aides to cover the following daily requirements:

Monday	80
Tuesday	85
Wednesday	75
Thursday	80
Friday	90
Saturday	70
Sunday	60

You can start an aide on any day of the week. An aide works 5 consecutive days, then gets 2 days off. Aides cost \$40 per day on weekdays, \$44 on Saturday and Sunday. This means that a work week that includes no weekend days costs \$200. A work week that includes one weekend day costs \$204. A work week that includes both weekend days costs \$208. How do you hire aides so as to minimize your total cost per week? What will your total cost per week be?

3. Let's build on the transportation problem example on page 13. Imagine that more problems have arisen. Depot A and Customers C and D are in Canada. Depot B and Customer E are in the USA. Canada will not allow you to import more than 100 units from the USA. Also, there's a weak bridge along the road from A to C. No more than 40 units can be sent that way. How much do you move from each depot to each customer to meet all the requirements at minimum cost? If you could pay someone to smuggle units in excess of 100 into Canada, how much would it be worth to you per unit?