

Constructing an Initial Basic Feasible Solution

We will use the previous numerical example to illustrate the methods. In algebraic form, our problem is:

$$\begin{array}{rllllllll}
 \text{Minimize} & 3x_{11} & +2x_{12} & & +x_{21} & +5x_{22} & & +5x_{31} & +4x_{32} & & \\
 \text{Subject to:} & & & & & & & & & & \\
 & x_{11} & +x_{12} & +x_{13} & & & & & & & = 45 \\
 & & & & x_{21} & +x_{22} & +x_{23} & & & & = 60 \\
 & & & & & & & x_{31} & +x_{32} & +x_{33} & = 35 \\
 & x_{11} & & & +x_{21} & & & +x_{31} & & & = 50 \\
 & & x_{12} & & & +x_{22} & & & +x_{32} & & = 60 \\
 & & & x_{13} & & & +x_{23} & & & +x_{33} & = 30 \\
 & x_{ij} \geq 0 & \text{for } i = 1, 2, 3 \text{ and } j = 1, 2, 3; & & & & & & & &
 \end{array}$$

and in tableau form, this problem is specified as:

		Sinks			
		1	2	3	
	1	3	2	0	45
	2	1	5	0	60
Sources	3	5	4	0	35
		50	60	30	

An inspection of the algebraic form of the problem shows that we have a set of 6 equality constraints in 9 variables. We will first argue that it is possible to reduce the number of constraint equations by 1. This is a consequence of the following special feature of all *balanced* transportation problems. Notice that the three supply constraints sum up to

$$x_{11} + x_{12} + x_{13} + x_{21} + x_{22} + x_{23} + x_{31} + x_{32} + x_{33} = 140 ;$$

and that the three demand constraints also sum up to this same expression. It follows that it is sufficient to work with any subset of 5 out of these 6 equations. To understand what this means more explicitly, suppose we are given a proposed solution to the original equations, and we are asked to check whether or not the given solution is feasible. To answer this, we would have to substitute the proposed solution into the equations one by one. Now, observe that if the proposed solution is verified to satisfy the first 5 equations, then the equality of the sum of the supply equations and the sum of the demand equations implies that the proposed solution must also satisfy the last equation. Moreover, it should be clear that this

observation holds regardless of the order in which the equations are checked. Thus, indeed, we can choose to remove any one of the given 6 equations, and doing so will not alter the feasible set in any way.

The next question of course is: Which equation should we remove? Interestingly, the answer is that we don't have to commit ourselves at this point. Intuitively, the reason behind this answer is that doing so would give us more flexibility; this will be made more apparent a little bit later.

With one equation removed (conceptually, that is), a basic feasible solution will have 5 basic variables and 4 nonbasic variables. If we are to solve this problem by the standard Simplex method, then our next task is to introduce 5 artificial variables (Actually, 4 is sufficient. Why?) and begin Phase I of the solution procedure. As noted earlier, it seems desirable to avoid introducing these additional variables (since we have to drive them out of the basis eventually). We are, therefore, motivated to explore other, hopefully more-direct, approaches. For this purpose, we will now switch to the simpler tableau representation of the problem.

Our first observation is that it is quite easy to construct a feasible solution to the problem. The general idea is to arbitrarily choose an empty cell in the above tableau and assign a value to the x_{ij} in that cell; after having done that, we then repeat the same process until every cell is assigned an explicit x_{ij} value. Now, as we begin this assignment process, notice that we do need to enforce the supply and demand constraints by making assignments that are such that: (i) the sum of the x_{ij} 's in every row equals the specified supply at the right margin of that row; and (ii) the sum of the x_{ij} 's in every column equals the specified demand at the bottom margin of that column. In other words, we are only interested in assignments with "correct" row sums and column sums. As an example, our earlier solution, namely

		Sinks			
		1	2	3	
	1	3	2	0	
1		20	20	5	45
	2	1	5	0	
Sources 2		20	20	20	60
	3	5	4	0	
3		10	20	5	35
		50	60	30	

satisfies these requirements; and in fact, it can be viewed as the outcome of such an assignment procedure. It should be clear now that this procedure is indeed very easy to implement; and you should attempt to construct a different feasible solution yourself.

Our goal, however, is to construct a basic feasible solution. This means that we actually have requirements that are stronger than just having correct row sums and column

sums. Therefore, the question now is: What additional stipulations should be added to this procedure to ensure that the outcome is a basic feasible solution?

This brings us to an important observation. Earlier, we made the point that out of a total of 9 variables, there are only 5 basic variables in every basic feasible solution. Since all nonbasic variables are assigned the value 0, a basic feasible solution must have at least 4 of its values equal to 0. (It is possible to have more than 4, since the basic solution may be degenerate.) It follows that we should avoid introducing “too many” positive entries into the tableau. For example, the feasible solution above does not have any 0 entry; therefore, it is not a basic solution.

An intuitive idea that seems to be helpful in this regard is that whenever we are about to assign a value into a cell, we should attempt to make that assignment as large as possible. Suppose the cell into which we are about to assign a value is cell (i, j) ; and let us refer to this cell as the *entering cell*. Notice that as a result of having possibly assigned previous entries into row i and into column j , part of the original supply from Source i has been depleted, and part of the original demand at Sink j has been met. Therefore, assigning the largest possible value into cell (i, j) means that this value should equal to either the *remaining supply* from Source i or the *remaining demand* at Sink j , whichever is smaller. It turns out that this additional stipulation is almost enough. Why “almost”? It is because we may go too far and end up with an insufficient number of basic variables.

Our next task, therefore, is to develop one more (and last) stipulation that is designed to eliminate the possibility of not having enough basic variables. For this purpose, we will now go through a careful examination of our assignment procedure.

Suppose we begin the assignment process with cell $(1, 1)$ as the entering cell. Since this is our first attempt at an assignment, the remaining supply from Source 1 is 45, and the remaining demand at Sink 1 is 50. According to the above discussion, we should, therefore, assign 45 into cell $(1, 1)$ as the value of x_{11} . This immediately implies that x_{11} will serve as a basic variable. Moreover, since the remaining supply from Source 1 is exhausted as a result of this assignment, the key concept at this point is to also think of this action as designating x_{11} as the basic variable associated with equation (1). Since every equation is to have exactly one basic variable, we shall say that equation (1) has now received its “quota” of one basic variable. Having just filled its quota, we should therefore prevent equation (1) from “receiving” another basic variable. Mechanically, this is achieved by crossing out, or removing, the first row. Doing this will result in a new tableau with one less row, which will then serve as the starting point for a continuation of the same assignment procedure. Notice that to facilitate the process of making further assignments, we should now reduce the remaining demand at Sink 1 to 5, from 50. This completes what we shall refer to as an *assignment cycle*.

In summary, our first assignment cycle results in the revised tableau below.

		Sinks			
		1	2	3	
Sources	1*	3	2	0	45 0
		45			
	2	1	5	0	
3	5	4	0	35	
		50 5	60	30	

Thus, we have: (i) entered the value 45 into cell (1, 1) as x_{11} ; (ii) revised the original supply from Source 1 and the original demand at Sink 1 to a remaining supply of 0 and a remaining demand of 5, respectively; (iii) marked the first source (or row) with an “*” (asterisk) to indicate its removal; and (iv) crossed out the remaining supply from the first source to indicate that it has been exhausted. (Here, for typographical ease, we choose to put an asterisk next to the name of the source to indicate its removal. For hand calculations, it is better to draw an explicit line across row 1.)

The fact that exactly one row or one column is removed at the completion of an assignment cycle is critical, in that it is the additional stipulation that ensures that a correct number of basic variables is generated at the end of the entire assignment process. To comprehend this claim fully, let us now iterate our procedure to completion. Suppose cell (2, 1) is chosen as the next entering cell. Then, after assigning 5 as x_{21} , revising the remaining supply from Source 2 and the remaining demand at Sink 1 to 55 and 0 (respectively), marking the removal of Sink 1, and crossing out the remaining demand from Sink 1, we obtain the new tableau below.

		Sinks			
		1*	2	3	
Sources	1*	3	2	0	45 0
		45			
	2	1	5	0	
3	5	4	0	35	
		50 5 0	60	30	

At this point, the remaining cells are (2, 2), (2, 3), (3, 2), and (3, 3). Suppose cell (2, 2) is chosen as the next entering cell. Then, after assigning 55 as x_{22} and going through another

round of routine updates, we obtain

		Sinks				
		1*	2	3		
Sources	1*	3		2	0	45 0
	45					
	1		5		0	
2*		5	55			
5		4		0	35	
3						
		50 5 0	60 5	30		

and, at this point, only cells (3, 2) and (3, 3) remain. Suppose we choose to enter a 5 in cell (3, 2); then, the new tableau is

		Sinks				
		1*	2*	3		
Sources	1*	3		2	0	45 0
	45					
	1		5		0	
2*		5	55			
5		4		0	35 30	
3			5			
		50 5 0	60 5 0	30		

and we now have a single remaining cell, cell (3, 3). Notice that the remaining supply and the remaining demand are both at 30. Clearly, we should assign 30 as x_{33} ; and this assignment exhausts both the remaining supply and the remaining demand simultaneously. In a tied situation like this, we can choose to remove either Source 3 or Sink 3. Suppose Source 3 is chosen; then, a final round of updates leads to the tableau below.

		Sinks				
		1*	2*	3		
Sources	1*	3		2	0	45 0
	45					
	1		5		0	
2*		5	55			
5		4		0	35 30 0	
3*			5	30		
		50 5 0	60 5 0	30 0		

This completes the entire assignment procedure, with the explicit assignments $x_{11} = 45$, $x_{21} = 5$, $x_{22} = 55$, $x_{32} = 5$, and $x_{33} = 30$. Cells without an explicit assignment are considered nonbasic; and therefore, their x_{ij} values are all equal to 0. The objective-function

value of this solution is easily computed as:

$$3 \times 45 + 1 \times 5 + 5 \times 55 + 4 \times 5 + 0 \times 30 = 435.$$

That the final remaining supply and the final remaining demand are exhausted simultaneously is a direct consequence of the fact that we have a balanced transportation problem. It also confirms our earlier claim that out of the original 6 constraint equations, only 5 are necessary to define the feasible set.

More importantly, observe that after the removal of Source 3, which concludes the assignment process, Sink 3 remains as the only source or sink that has not been removed. (This is despite the fact that there is no more remaining cell.) Recall that the explicit assignment of an x_{ij} is tantamount to the designation of that x_{ij} as the basic variable associated with the constraint equation whose right-hand-side supply or demand is exhausted by this assignment. It follows that this observation is what allows us to assert that a correct number of basic variables is assigned at the end of this assignment procedure. In this example, this variable count equals 5, which comes from the total number of equations, $3 + 3$, minus the number of redundant equations, 1. In general, for a balanced transportation problem with m sources and n sinks, the correct basic-variable count is equal to $m + n - 1$.

Since equation (6) is the only equation that did not receive an assignment of an associated basic variable, we will consider it as the redundant equation. It should be clear that this identification is a result of the choice of our particular sequence of entering cells. For other choices, one could end up with the identification of a different equation as the redundant equation. A little bit of reflection now reveals that the fact that we did not declare at the outset of the solution procedure a specific equation as the redundant equation is quite helpful, in the sense that doing so indeed offered more flexibility in our effort to construct an initial basic feasible solution.

Another related interesting observation is that if we had chosen cell (2, 2) as the first entering cell in the assignment process, then we would immediately encounter a tied situation. According to our tie-breaking strategy, we should assign 60 as x_{22} and remove either Source 2 or Sink 2 (but not both). Regardless of which of these is selected for removal, the remaining supply from the other source or the remaining demand at the other sink is also reduced to 0 at the same time. For the purpose of discussion, let us choose to remove Source 2; then, the fact that Sink 2, which has no remaining demand, continues to be available for further assignments implies that somewhere down the road in the remainder of the assignment process, we will be forced into assigning a 0 as the x_{ij} value in the second column. Such an *explicit* 0-assignment is important, in that it corresponds to the explicit declaration of the variable whose value is being assigned as a degenerate basic variable. Therefore, we should be sure not to miss, or skip, any such assignments. Otherwise, we would end up with an insufficient number of basic variables.

The assignment procedure described above is called the *northwest-corner method*. This name originates from the fact that at the start of each assignment cycle, the entering cell is always chosen to be the one that is located at the northwest corner of all remaining cells. Since this strategy makes no reference whatsoever to the c_{ij} values, it cannot, in general, be expected to produce good initial basic feasible solutions. (For example, the objective-function value of the above basic feasible solution equals 435, which is worse than that of the earlier feasible solution, at 350.) We will next describe two better procedures that take into account the fact that we are interested in minimizing $\sum_{i=1}^m \sum_{j=1}^n c_{ij}x_{ij}$.

The Least-Cost Method

The only difference between the least-cost method and the northwest-corner method is in the choice of entering variables. Here, the strategy is to always select the cell with the smallest c_{ij} value among all remaining cells as the entering cell. Ties are, as usual, broken arbitrarily.

Below, we will briefly go over an implementation of this new strategy.

We will start with the initial tableau. Since all three c_{ij} 's in column 3 are equal to 0, we can choose any one of the cells in column 3 as the first entering cell. Let cell (1, 3) be our choice. Then, after assigning 30 to x_{13} and going through a round of updates, we obtain the tableau below.

		Sinks			
		1	2	3*	
	1	3	2	0	
Sources	2	1	5	0	30
	3	5	4	0	
		50	60	30 0	

The next entering cell is cell (2, 1). After assigning 50 to x_{21} and going through necessary updates, we obtain the tableau below.

		Sinks			
		1*	2	3*	
	1	3	2	0	
Sources	2	1	5	0	30
	3	5	4	0	
		50 0	60	30 0	

With only column 2 remaining, it is easily seen that we should then (sequentially) assign 15 as x_{12} , 35 as x_{32} , and finally 10 as x_{22} . This yields the final tableau below.

		Sinks				
		1*	2	3*		
Sources	1*	3	2	0	45	1/5 0
			15	30		
2*	1	50	5	10	60	1/0 0
3*	5		4	0	35	0
		50 0	60 45	30 0		
		1/0 0				

Note that after assigning 10 to x_{22} , at the last step, we chose (arbitrarily) to remove Source 2, as opposed to removing Sink 2.

In conclusion, the assignments produced by the least-cost method are: $x_{12} = 15$, $x_{13} = 30$, $x_{21} = 50$, $x_{22} = 10$, and $x_{32} = 35$. This basic feasible solution has an objective-function value of 270, which is significantly better than the previous one.

The Vogel's Approximation Method

We have seen that the least-cost method (typically) offers a significant improvement over the northwest-corner method. The method, however, may be construed as being myopic, in that the choice of the entering cell, at every iteration, is based solely on the location of the smallest available c_{ij} . The Vogel's approximation method will be slightly less myopic.

The strategy for the selection of entering cells in the Vogel's method is as follows. At the start of every assignment cycle, the first step is to compute (or update) a set of *penalties*, one for each row and one for each column, and then, in the second step, to select the entering cell on the basis of the relative magnitudes of these row and column penalties. Specifically, the penalty associated with a row or a column is defined to be the difference between the second-lowest cost and the lowest cost in that row or column; and the entering cell is chosen to be the cell with the smallest c_{ij} in the row or column with the greatest penalty. All ties are broken arbitrarily.

The intuitive basis behind this strategy is as follows. Suppose we wish to assign a value to an x_{ij} in a given row (column). If there are no existing preconditions, then the best possible choice, clearly, is to make an assignment into the least-cost cell in that row (column). However, in the course of an assignment process, we may have made earlier assignments into various parts of the given tableau; and as a result of these earlier assignments, we may not be able to assign anything into that least-cost cell. If this indeed happens to

be the case, we would then be forced into an assignment in a higher-cost cell in that row (column). Now, observe that the incremental cost (rate) of such a “forced” assignment is no less than the difference between the second-lowest cost and the least cost in the given row (column). It follows that the penalty associated with a row (column) can be interpreted as the incremental cost that would incur if we are unable to make an assignment into the least-cost cell in that row (column). With this interpretation, it should now be clear that the underlying intent of the strategy in the Vogel’s method is to reduce, in every assignment cycle, the potential incremental costs associated with forced assignments.

Again, we will use the previous example to illustrate the method. It is easily seen that row 1 has a penalty of $2 - 0 = 2$, row 2 has a penalty of $1 - 0 = 1$, row 3 has a penalty of $4 - 0 = 4$, column 1 has a penalty of $3 - 1 = 2$, column 2 has a penalty of $4 - 2 = 2$, and finally column 3 has a penalty of $0 - 0 = 0$. The results of these calculations are displayed on the right and bottom margins of the tableau, as shown below.

		Sinks				
		1	2	3	Penalty	
		3	2	0		
1					45	2
Sources		1	5	0		
2					60	1
		5	4	0		
3					35	4
		50	60	30		
Penalty		2	2	0		

Since row 3 has the greatest penalty, the first entering cell will be in that row. Since cell (3, 3) is the least-cost cell in row 3, it will be the entering cell. After assigning 30 to x_{33} and going through a round of standard updates, we obtain the new tableau below.

		Sinks				
		1	2	3*	Penalty	
		3	2	0		
1					45	2
Sources		1	5	0		
2					60	1
		5	4	0		
3				30	35 5	4
		50	60	30 0		
Penalty		2	2	0		

(Notice that together with the removal of column 3, we have also crossed out the penalty associated with that column.) The fact that column 3 has been removed implies that the

row penalties (but not the column penalties) should now be revised. Specifically, the penalty associated with row 1 should be revised to $3 - 2 = 1$; and similarly, those for row 2 and row 3 should be revised to $5 - 1 = 4$ and $5 - 4 = 1$, respectively. These revisions will be indicated on the right margin of the tableau, as shown below.

		Sinks				
		1	2	3*		
Sources	1	3	2	0	45	2 1
	2	1	5	0	60	1 4
	3	5	4	0	3 5	4 1
		50	60	30 0		
Penalty		2	2	0		

With these revisions, row 2 now has the greatest penalty; therefore, the next entering cell is cell (2, 1). After assigning 50 as x_{21} and updating in the usual manner, we obtain the tableau below.

		Sinks				
		1*	2	3*		
Sources	1	3	2	0	45	2 1
	2	1	5	0	60 10	1 4
	3	5	4	0	3 5	4 1
		50 0	60	30 0		
Penalty		2	2	0		

With only column 2 remaining, there is no need to update the penalties further. After assigning (sequentially) 45 as x_{12} , 5 as x_{32} , and 10 as x_{22} , we arrive at the final tableau

below.

		Sinks				
		1*	2	3*		
Sources	1*	3	2	0	45 0	2 1
	2*	1	5	0	60 10 0	1 4
	3*	5	4	0	35 5 0	4 1
		50 0	60 15	30 0		
		10 0				
Penalty		2	2	0		

Note that after assigning 10 to x_{22} , at the last step, we chose (arbitrarily) to remove Source 2, as opposed to removing Sink 2.

In conclusion, the assignments produced by the Vogel's method are: $x_{12} = 45$, $x_{21} = 50$, $x_{22} = 10$, $x_{32} = 5$, and $x_{33} = 30$. This basic feasible solution has an objective-function value of 210, which happens to be an improvement over the one produced by the least-cost method. While it is true that solutions produced by the Vogel's method are typically better than those produced by the least-cost method, such an outcome cannot be guaranteed in general.

Discussion

Several other methods for constructing initial basic feasible solutions can be found in our text. These methods offer some differences in terms of total computational effort and in terms of the quality of the produced initial basic feasible solutions. In general, it is difficult to achieve a perfect balance between effort and quality. In fact, it may not even be desirable to do so, since constructing an initial basic feasible solution is only the first phase in the solution of a problem. In other words, it is the total solution effort for a problem that matters in the end. We will, therefore, not attempt to dwell upon a detailed discussion of these other methods.

In some applications, the objective may be to maximize the overall profit derived from the shipments. That is, in place of the c_{ij} 's, we could be given a set of p_{ij} 's, where p_{ij} is the profit per unit of shipment from Source i to Sink j . Clearly, the methods discussed above can be easily adapted to handle such problems. For example, the obvious counterpart of the least-cost method would select, in every assignment cycle, the cell with the greatest available p_{ij} as the entering cell. The Vogel's method can also be modified in a similar manner.