

LP Geometry and Solution

The intuition provided by the geometric solution of two-variable LP's serves well for problems in higher dimensions. This note offers a quick review of the terminology of LP graphical solution, verbally translating the elements of the solution techniques from two to three to " n " dimensions. To get the most out of the note, we suggest that as you read it, you look at illustrative slides from the lectures or "draw along."

The number of decision variables in an LP determines the problem's *dimensionality*. Two-variable problems can be represented in a *space* of two dimensions by drawing a set of two axes, one perpendicular to the other. Three-variable problems can be represented in a 3-D space with a set of three mutually perpendicular axes. n -variable problems can be represented by n -dimension spaces based on a set of n mutually perpendicular axes. While most of us can picture two- and three-dimensional problems in our minds, the visualization of more dimensions is beyond our capabilities. Nevertheless, the geometric intuition carries through and it often suffices to think of higher-dimensional problems as "more complicated" three-dimensional problems.

Each inequality constraint (\leq or \geq) divides the problem space into two parts, called *half-spaces*. All of the points in one of the half-spaces satisfy the constraint, and none of the points in the other satisfy the constraint. The points on the boundary of the two half-spaces define the equality constraint ($=$) which corresponds to the inequality. Thus, in two dimensions an inequality constraint divides the 2-D plane into two parts, while an equality constraint is satisfied only on a line. In three dimensions, an inequality divides volumetric space into two, and an equality is satisfied only on a plane. (In higher dimensions, equality constraints define n -dimensional *hyperplanes*!)

The feasible solution of an LP must simultaneously satisfy all of the problem's constraints. If the problem has only inequality constraints, the set of feasible solutions is, therefore, the intersection of the half-spaces define by the problem's constraints.

The problem's constraints may even have *no* common intersection so that the set of feasible points is empty. In this case the problem is *infeasible* (which may mean that the problem is formulated incorrectly).

If we increase the right-hand side of a ' \leq ' constraint or decrease the right hand side of a ' \geq ' inequality, we *relax* the constraint and enlarge the feasible region to include additional points that simultaneously satisfy all of the LP constraints. This action can only make the optimal Objective Function Value improve or stay the same. That is, the inclusion of new points in the feasible region does not remove any of the original feasible points—including the original optimal solution. In addition, one of the new feasible points may even obtain an

OFV that improves upon the old optimal solution's. Therefore, the new optimal OFV must be at least as "good" as the original. (The opposite actions *tighten* constraints, reduce the set of points included in the feasible region, and can only make the optimal OFV deteriorate or stay the same.)

If the underlying problem space has n dimensions then every corner point, or *vertex*, of the feasible region can be defined by exactly n constraints.¹ Thus, the vertices of 2-D problems can be defined by two constraints, and the vertices of 3-D problems can be defined by three constraints (For example, think of the corners of a cube.) It may also be that more than n constraints pass through a vertex in n -space, in which case the vertex is said to be *degenerate*. For a such a degenerate vertex any subset of n distinct constraints from the original set is sufficient to define the vertex. By "distinct" we mean to exclude different versions of the same constraint. For example, $4x_1 + 6x_2 \geq 5$ and $8x_1 + 12x_2 \geq 10$ look different from each other but are satisfied by the same feasible regions.

For every feasible problem there exists an optimal solution which is located at a *vertex* of the feasible region.² A constraint is *binding* if it passes through this optimal vertex, and *nonbinding* if it does not. If the constraint is binding these actions will most often change the optimal solution and OFV. In general, as the RHS of the binding constraint is changed, the optimal vertex "slides along" the intersection of changing constraint, and as the optimal vertex moves, the optimal OFV changes. An exception exists, however, when the optimal vertex is degenerate. In this case the vertex may or may not move; whether or not it does depends on the specific position of the vertices ($> n$) binding constraints and is problem-specific.

If a constraint is *not* binding, then tightening it (a bit) or relaxing it (as much as you please) will not change the optimal solution or the optimal OFV. In particular, the *slack* of a nonbinding ' \leq ' constraint is defined to be the difference between its right hand side and the value of its left hand side at the optimal vertex. Formally it's defined as

$$\text{slack} = \text{RHS} - \text{LHS}.$$

¹ It may also be that more than n constraints pass through a vertex in n -space, in which case the vertex is said to be *degenerate*. For a such a degenerate vertex any subset of n distinct constraints from the original set is sufficient to define the vertex. By "distinct" we mean to exclude different versions of the same constraint. For example, $4x_1 + 6x_2 \geq 5$ and $8x_1 + 12x_2 \geq 10$ look different from each other but are satisfied by the same feasible regions.

² This is a fact which we hope the geometric solutions presented in class have convinced you is true. For those who are interested, a formal analytical explanation can be found in an advanced text on LP.

Similarly, the *surplus* associated with a nonbinding ' \geq ' constraint is the extra (*i.e.* surplus) value which may be reduced from the constraint's left-hand-side function before the constraint becomes binding and the LHS equals the RHS. Its formal definition is

$$\text{surplus} = \text{LHS} - \text{RHS}.$$

Note that by definition slack and surplus are always greater than zero.

Within an *allowable range*, change of an objective function coefficient will *not* change the vertex at which the optimal solution is found. In two dimensions, for example, changing one of the objective function coefficients causes the objective function line to rotate (around the solution vertex). If the change is large enough, the line will become parallel to one of the binding constraints, and the LP will have multiple optimal solutions. If the change in the coefficient is larger still, the optimal solution "jumps" to another vertex.

Changes to objective function coefficients do, however, change the optimal OFV.

Sensitivity Analysis

One benefit of using an LP model is that a table of sensitivity analysis, *an analysis of the sensitivity of the model's solution to changes in the problem's assumptions*, is provided to you "for free" every time you run an LP. This part of the note reviews the definitions of the terms included in most LP sensitivity reports, as well as the geometric concepts that behind the definitions. We begin with a recapitulation of some definitions.

Preliminaries for Sensitivity Analysis

It is important to distinguish between the *optimal solution*, *i.e.*, the values of the decision variables at optimality (often denoted as x_1^* , x_2^* , x_3^* , *etc.*, where the asterisk (*) indicates the optimal x), and the optimal *objective function value* (OFV), which is simply the value of the objective function when evaluated at the optimal solution.

We say that a constraint is *relaxed* or *loosened* when, for a \leq constraint, the right-hand side (RHS) is *increased*, or when, for a \geq constraint, the RHS is *decreased*. A change in the opposite direction is called a *tightening* or *restriction* of the constraint. We say that an objective function value is *improved* when *increased in a maximization* problem (e.g., increasing profit is an improvement) or *reduced in a minimization* problem (e.g., reducing cost is an improvement).

A reliable intuition is that the *relaxation* of a constraint can only *improve* the OFV or leave it unchanged. Conversely, the *tightening* of a constraint can

only *worsen* the OFV or leave it unchanged. This intuition is valuable, and contributes to straightforward interpretations of complex problems.

Shadow Prices and Allowable Ranges for the RHS

A natural economic interpretation of the degree to which the change in the right hand side of a constraint affects the optimal OFV is that of *marginal cost* or *marginal benefit*. We call this degree of change the *shadow price* of the constraint, and more formally define shadow price to be the *amount of improvement in the optimal OFV that is obtained by relaxing the right hand side by one unit*.³ Equivalently, the shadow price is the rate of *deterioration* in the OFV obtained by *restricting* that constraint. Note that a nonbinding constraint always has a shadow price of zero, since a change in its RHS does not affect the optimal solution or OFV at all.

The shadow price of a constraint is defined for a “one unit” change in the constraint. This “one unit” idea not only tells us that the shadow price is the rate of change of the OFV with respect to changes in the constraint, but also indicates that this shadow price is only locally accurate; if we make dramatic changes in the constraint, naively multiplying the shadow price by the magnitude of the change may mislead us. In particular, the shadow price reported by the spreadsheet holds only within an *allowable range* of changes to the constraint’s right-hand side; outside of this allowable range the shadow price may change. This allowable range is composed of two components. The *allowable increase* is the amount by which the RHS may be *increased* before the shadow price can change; similarly, the *allowable decrease* is the corresponding *reduction* that may applied to the RHS before a change in the shadow price can take place.⁴ (Whether this increase or decrease corresponds to a tightening or a relaxation of the constraint depends on the direction of the constraint’s inequality.)

For a binding constraint, the geometric intuition behind these definitions is as follows. By changing the RHS of a constraint, we change the optimal solution as it slides along the other binding constraints. Within the allowable range of changes to the RHS, the optimal vertex slides in a straight line, and the optimal OFV changes at a constant rate. Once the RHS hits the limit of its allowable increase or decrease, however, the optimal vertex’s slide changes. The vertex may turn a corner, continuing its straight-line slide in a new direction, in which

³ The Excel spreadsheet package defines the shadow price to be the *increase* in the optimal OFV that is obtained from a unit *increase* in the RHS of the constraint (irrespective of the direction of the constraint or of whether the objective function maximizes or minimizes). It may therefore attach a minus sign to the shadow price. In such a case, we must use the heuristic of “relax and improve” to determine the meaning of the minus sign. If a lower OFV (e.g., cost) is better, a negative shadow price might be attached to a \geq constraint on production quantity. *Relaxing* the production requirement *improves* the OFV.

⁴ Note that the OFV may change even though we stay within the allowable increase or decrease; it’s the *shadow price* which is guaranteed to stay constant.

case the optimal OFV changes at a new, constant rate. Or the constraint whose RHS is altered may become non-binding so that its shadow price drops to zero. Or, by tightening the constraint, the problem may become infeasible, in which case the shadow price is not even well defined!

For a nonbinding constraint (which, we remember, will always have a shadow price of zero), we make these observations: Further relaxation of the constraint will never make the constraint binding. One of the allowable limits will thus be infinite—the shadow price will remain zero no matter how much we *relax* the constraint. There always exists, however, an allowable limit on the *tightening* of the constraint beyond which the constraint becomes binding and its shadow price becomes non-zero.

Reduced Costs and Allowable Ranges for Objective Function Coefficients

The reduced cost of a decision variable is defined as the amount by which the *objective function coefficient* of that variable must be *improved* for that decision variable to take a positive value in the optimal solution. Equivalently, reduced cost represents the amount by which the optimal OFV will *deteriorate* if a unit of that variable (currently at zero) were to be forced into the solution.

An intuitive way to think about reduced costs is as follows. If the optimal solution to an LP indicates that the optimal level of a particular decision variable is zero⁵, it must be because the objective function coefficient of this variable (e.g., its unit contribution to profits or unit cost) is not “attractive” enough to justify its “inclusion” in the decision. The reduced cost of that decision variable tells us the amount by which the objective function coefficients must *improve* for the decision variable to become “attractive enough to include”—and take on a non-zero value in the optimal solution. Hence the reduced costs of all decision variables that take non-zero values⁶ in the optimal solution are, by definition, zero—no further enhancement to their attractiveness is needed to get us to use them, as they’re already “included.”

When the reduced cost of a decision variable is non-zero (implying that the value of that decision variable is zero in the optimal solution), the reduced cost is also reflected in the *allowable range* of its objective coefficient. In this case, one of the allowable limits is always infinite (because making the objective coefficient “less attractive” will never cause the optimal solution to include the decision variable in the optimal solution); and the other limit, by definition, is the reduced cost (for it is the amount by which the objective coefficient must “improve” before the optimal solution changes).

⁵ More generally, at its lower or upper limit as specified in the constraints.

⁶ More generally, values strictly *between* their upper and lower limits.

If the objective function coefficients are changed—one at a time—within their respective allowable ranges, the *optimal solution* does not change. (In 2-dimensional LPs, the objective function line rotates at the same vertex.) However, *any* change in the objective coefficient of a decision variable that has a nonzero optimal value changes the OFV, since we're directly changing the weights on the solution variables without changing the values of these variables.⁷ Note that changing the weights *is* enough to change the OFV). Changing an objective coefficient beyond the allowable range causes the optimal solution to jump to another vertex.

⁷ If we change the coefficient on a variable has a zero optimal value (i.e., one which has a nonzero *reduced cost*) by an amount inside its *allowable range*, our change will not have made the variable attractive enough to include—an application of the interaction between *reduced cost* and the *allowable range*.