Sensitivity analysis

Big picture

- *The setting:* You have solved a linear program with the simplex method, and you have an optimal tableau.
- *The question:* If one of the numbers in the LP were changed slightly, how would that change affect the optimal solution?
- *The answer:* Sensitivity analysis provides two kinds of information.
  - A maximum increase and a maximum decrease that can be made to the number in the LP before the change becomes a "major" change. (A change within these limits is a "minor" change.)
  - If the change is a "minor" change, sensitivity analysis gives the optimal solution and the optimal objective value for the changed LP.

Correspondences

- Every **constraint** has a corresponding slack, surplus, and/or artificial variable:
  - A $\leq$ constraint has a corresponding **slack variable**.
  - An $=$ constraint has a corresponding **artificial variable**.
  - A $\geq$ constraint has a corresponding **surplus variable** and a corresponding **artificial variable**.
- Every **constraint** also has a corresponding **dual variable**.
- Every **variable** has a corresponding **column** in the simplex tableau.
- Every **basic variable** has a corresponding **row** in the simplex tableau.
- Note: Rows in the simplex tableau do *not* correspond to constraints!

Meanings of the numbers in the objective row

The numbers in the objective row of a simplex tableau can be interpreted in two different ways:

- In *any* tableau, a number in the objective row in a nonbasic column indicates how much the value of the objective function will change if the value of the corresponding variable is increased by 1, but the sign is flipped: a *negative* entry indicates an *increase* in the objective value, and a *positive* entry indicates a *decrease*.
  - This is why, in the simplex algorithm, we pivot on columns with negative entries in the objective row. A negative number means that giving the corresponding variable a positive value (i.e., making that column basic) will *increase* the value of the objective function.
  - Note: This interpretation of the entries in the objective row gives information about the effects of a *change to the basis* of the current solution, without changing the LP itself; it gives information about what would happen if you move from the current corner of the feasible region to a neighboring corner of the feasible region by changing which variables are basic (and thereby changing the values of the variables).
- In an *optimal* tableau, the numbers in the objective row in columns corresponding to *slack* and *artificial* variables (*not* surplus variables) are the optimal values of the corresponding **dual variables.**
  - The optimal value of a dual variable is the amount that the optimal objective value would change if the right-hand side of the corresponding constraint were *increased* by 1. (There is no sign flip here: a *positive* dual variable indicates an *increase* in the objective value, and a *negative* dual variable indicates a *decrease*.)
  - Note: This interpretation gives information about the effects of a *change to the LP itself.*

Right-hand side sensitivity analysis

When the proposed change to the LP is to the *right-hand side of a constraint*.

- If the constraint is **tight** in the optimal solution:
  1. Look at the column for the corresponding *slack* or *artificial* variable (*not* the surplus variable). Call this column the "working column."
  2. Form quotients by dividing the entries in the *right-hand column* of the optimal tableau by the entries in the working column (not including the entries in the objective row). [Mnemonic: Use the *right-hand column* for *right-hand side sensitivity analysis*.]
  3. To find the **maximum increase**: Consider the quotients from step 2 that use *negative* entries from the working column. Flip the signs of those quotients to make them positive, and then choose the smallest.
     - If the working column has no negative entries, then there is no maximum increase, i.e., there is no limit to the amount the right-hand side of the constraint can be increased.
  4. To find the **maximum decrease**: Consider the quotients from step 2 that use *positive* entries from the working column. Choose the smallest such quotient.
     - If the working column has no positive entries, then there is no maximum decrease, i.e., there is no limit to the amount the right-hand side of the constraint can be decreased.
  5. Within the limits identified in steps 3 and 4, every unit *increase* in the right-hand side of the constraint will change the right-hand column of the optimal tableau by the corresponding value in the working column. (This includes the value of the objective function, for which the per-unit change is given by the value of the optimal dual variable, found at the bottom of the working column.) Every unit *decrease* in the right-hand side of the constraint will have the opposite effect.
     - So to get the entries in the right-hand column of the new optimal tableau, multiply the *change* in the right-hand side of the constraint by each entry in the working column, and add those products to the corresponding entries in the right-hand side of the original optimal tableau. (Do this for the objective row too, to get the new optimal objective value.) Now you can easily read off the values of the basic variables in the optimal solution to the changed LP.
- If the constraint is **not tight** in the optimal solution:
  - For a $\leq$ constraint, the corresponding *slack variable* will be positive. The value of this slack variable is the **maximum decrease** for the right-hand side of the constrant. There is **no maximum increase** (i.e., there is no limit to the amount the right-hand side of the constraint can be increased).
  - For a $\geq$ constraint, the corresponding *surplus variable* will be positive. The value of this surplus variable is the **maximum increase** for the right-hand side of the constrant. There is **no maximum decrease** (i.e., there is no limit to the amount the right-hand side of the constraint can be decreased).
  - Within the limits identified above, any change to the right-hand side of the constraint will have *no effect* to the objective value or the optimal solution (except the value of the corresponding slack or surplus variable).

If the right-hand side of a constraint is changed by the maximum amount indicated by sensitivity analysis, the optimal solution for the changed LP will be **degenerate**. (One of the numbers in the *right-hand column* of the optimal tableau will be 0.)

Objective function coefficient sensitivity analysis

When the proposed change to the LP is to a *coefficient of a variable in the objective function.* (Suppose the variable is x_i.)

- If the variable x_i is **basic** in the optimal solution:
  1. Look at the row corresponding to x_i. Call this row the "working row."
  2. Form quotients by dividing the entries in the *objective row* of the optimal tableau by the entries in the working row, but *ignore* the following columns:
     - the basic columns, including the column for x_i;
     - any columns for *artificial variables*; and
     - the right-hand column.

     [Mnemonic: Use the *objective row* for *objective function coefficient sensitivity analysis.*]
  3. To find the **maximum increase**: Consider the quotients from step 2 that use *negative* entries from the working row. Flip the signs of those quotients to make them positive, and then choose the smallest.
     - If the working row has no negative entries, then there is no maximum increase, i.e., there is no limit to the amount the coefficient can be increased.
  4. To find the **maximum decrease**: Consider the quotients from step 2 that use *positive* entries from the working row. Choose the smallest such quotient.
     - If the working row has no positive entries, then there is no maximum decrease, i.e., there is no limit to the amount the coefficient can be decreased.
  5. Within the limits identified in steps 3 and 4, every unit *increase* in the objective function coefficient of x_i will increase the value of the objective function by the optimal value of x_i. Every unit *decrease* in the coefficient will have the opposite effect. The optimal *solution* will not change.
     - For example, if the optimal value of x_3 is 57 and its objective function coefficient is increased by 1, then the optimal objective value will increase by 57 (because every unit of x_3 is now making $1 more profit). If its objective function coefficient is increased by 2, then the optimal objective value will increase by $2 \times 57 = 114$.
- If the variable x_i is **nonbasic** in the optimal solution:
  - The variable x_i does not have a positive value because it is not profitable enough. The entry in the objective row in the x_i column, which must be nonnegative (because the tableau is optimal), indicates the amount that the objective value would *decrease* for each unit increase to the value of x_i. So the per-unit profit for x_i (i.e., the objective function coefficient of x_i) would need to *increase* by this amount in order to make x_i profitable enough to produce. Until the per-unit profit increases this much, x_i will still not be profitable enough, so the optimal solution will not change.
  - Therefore, the **maximum increase** to the objective function coefficient of x_i (before the optimal solution changes) is given by the value in the objective row in the x_i column. There is **no maximum decrease**, i.e., there is no limit to the amount the coefficient can be decreased.
  - Any change to the objective function coefficient of x_i that does not exceed the maximum increase identified above will have *no effect* on the optimal objective value or the optimal solution.

If an objective function coefficient is changed by the maximum amount indicated by sensitivity analysis, the optimal solution for the changed LP will be **non-unique**. (One of the numbers in the *objective row* of the optimal tableau, in a nonbasic column, will be 0.)