

Tutorial 9: Simplex Explained

“Does an ‘explanation’ make it any less impressive?”

— Ludwig Wittgenstein

This tutorial explains how the simplex algorithm solves the LP model, i.e., how the simplex algorithm takes us from the mathematical model of the electricity market to the resulting prices and quantities.

Demonstrating the simplex algorithm

Objective and Constraints

As discussed in Tutorial 1: Explaining Prices, the simplex algorithm solves a Linear Programming (LP) model, which is a mathematical problem specified in terms of:

- *An objective* equation
- *Constraint* equations that must be obeyed while achieving the objective.

The objective of the electricity market model is to maximize the objective value, which is calculated as the difference between the value of the load that is supplied, and the cost of the generation offers that must be cleared in order to supply the load.

The constraints ensure that the mathematical model behaves in the same way as the physical electricity network, e.g., there is a constraint that models how power flows along a transmission circuit.

Small electricity market model

The details of the simplex algorithm can be demonstrated by means of a small electricity market model consisting of a generator and a load connected via a bus, as shown in Figure 119. Build the model by tapping the Bus, Gen, Load buttons on the Build toolbar.




Figure 119: The small model

Constraints for the small model

The constraints that model the physical reality of the small model are listed in Table 8.

Table 8: Constraints in the small model

Component	Purpose of associated constraint
Bus00	Node balance constraint ensures that the quantity of electricity that enters the bus is equal to quantity that leaves
Gen00	Limit the cleared offer quantity to no more than the quantity that was offered
Load00	Limit the cleared bid quantity to no more than the quantity that was bid

The constraints listed in Table 8 are modelled by the equations shown in Equation 25, Equation 26 and Equation 27.

$$genOffer_{Cleared} - loadBid_{Cleared} = 0$$

Equation 25: Node balance equation

$$loadBid_{Cleared} \leq loadBid_{Max}$$

Equation 26: Max cleared bids

$$genOffer_{Cleared} \leq genOffer_{Max}$$

Equation 27: Max cleared offers

Objective function for the small model

The objective function for the small model is represented by Equation 28.

Maximize:

$$\begin{aligned} & loadBid_{cleared} \times loadBid_{Price} \\ & - genOffer_{cleared} \times genOffer_{Price} \end{aligned}$$

Equation 28: Objective function

Simplex standard form

The simplex algorithm requires that the equations be expressed in a standard form:

- All equations expressed as $ax \leq b$ with $b \geq 0$, i.e., the RHS values are all ≥ 0
- Objective function expressed as a value to be maximized.

The electricity market model in standard form

Our objective function already has maximization as its requirement, which matches standard form.

All the constraints except for the node balance constraint meet the $ax \leq b$ requirement. The node balance constraint is an equality constraint, so we replace the $=$ with a \geq and a \leq .

The \leq is already in standard form. The \geq equation is not, so we multiply both sides by -1 to turn it into a \leq equation. We can do this because the RHS is zero; hence after multiplying it by -1 the RHS is still zero and we still meet the $b \geq 0$ requirement.

After this adjustment the complete set of constraints for the model, expressed in the form required by the simplex algorithm, are as shown in Equation 29.

$$-loadBid_{Cleared} + genOffer_{Cleared} \leq 0$$

$$loadBid_{Cleared} - genOffer_{Cleared} \leq 0$$

$$loadBid_{Cleared} \leq loadBid_{Max}$$

$$genOffer_{Cleared} \leq genOffer_{Max}$$

Equation 29: Equation constraints for the small model, expressed in standard form

Viewing the constraints

To view the constraints for the small model, first you need to solve the model by tapping the Solve button, which takes you to the Solve Settings display, make sure all the options are selected OFF and then tap the Solve Now button.

When the solver has completed, tap the Results button, then tap the Constraints row, which takes you to the constraints display shown in Figure 120.

← Results Constraints
<p>LOAD00 LOAD @ BUS00</p> <p>bus00_load00_bid00: BidBlockMax(LTE) constraint: Shadow Price: \$90.00 $+1.00000 \cdot \text{bus00_load00_bid00}_{\{\text{Cleared}\}} \leq 100.00000$</p>
<p>GEN00 GEN @ BUS00</p> <p>bus00_gen00_offer00: OfferBlockMax(LTE) constraint: Shadow Price: \$0.00 $+1.00000 \cdot \text{bus00_gen00_offer00}_{\{\text{Cleared}\}} \leq 250.00000$</p>
<p>BUS00 BUS</p> <p>bus00: NodeBalance(LTE) constraint: Shadow Price: \$0.00 $+1.00000 \cdot \text{bus00_gen00_offer00}_{\{\text{Cleared}\}} - 1.00000 \cdot \text{bus00_load00_bid00}_{\{\text{Cleared}\}} \leq 0.00000$</p> <p>bus00: NodeBalance(GTE) constraint: Shadow Price: \$70.00 $-1.00000 \cdot \text{bus00_gen00_offer00}_{\{\text{Cleared}\}} + 1.00000 \cdot \text{bus00_load00_bid00}_{\{\text{Cleared}\}} \leq 0.00000$</p>

Figure 120: Constraint equations for the small model

Parameters and Variables

The elements of the constraint equations are *variables* and *parameters*.

Parameters are the fixed amounts:

$$loadBid_{Max} = 100$$

$$genOffer_{Max} = 250$$

Variables are the values that will be determined by the simplex algorithm:

$$loadBid_{cleared}$$

$$genOffer_{cleared}$$

Electricity market equations represented in tableau form

The simplex algorithm is easier to follow if the equations that form the mathematical model are represented as a tableau.

In the tableau the constraints are the rows, the variables are the columns, and the number at the intersection of the rows and columns represents the factor that the variable has in the constraint.

The following sections describe the actions that the simplex algorithm takes in order to maximize the objective value. You can track these actions via the tableaux.

The tableaux for the small model are shown in Figure 121 and Figure 122 (which are two halves of one spreadsheet, presented as two figures because it is too wide for this page).

###TABLEAU[0001]### col[1] will enter via row[2] and col[4] will leave			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	-1	1
[2]_{NodeBalanceGTE} basic col[4]	0	1	-1
[3]bid00_{BidBlockMaxLTE} basic col[5]	100	1	0
[4]offer00_{OfferBlockMaxLTE} basic col[6]	250	0	1
[5]ObjectiveFn	0	-160	70
###TABLEAU[0002]### col[2] will enter via row[3] and col[5] will leave			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	0	0
[2]_{NodeBalanceGTE} basic col[1]	0	1	-1
[3]bid00_{BidBlockMaxLTE} basic col[5]	100	0	1
[4]offer00_{OfferBlockMaxLTE} basic col[6]	250	0	1
[5]ObjectiveFn	0	0	-90
###TABLEAU[0003]### FINAL			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	0	0
[2]_{NodeBalanceGTE} basic col[1]	100	1	0
[3]bid00_{BidBlockMaxLTE} basic col[2]	100	0	1
[4]offer00_{OfferBlockMaxLTE} basic col[6]	150	0	0
[5]ObjectiveFn	9000	0	0

Figure 121: Left side of the tableaux

[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1
0	0	0	0
[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	1	0	0
0	1	0	0
0	-1	1	0
0	0	0	1
0	160	0	0
[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	1	0	0
0	0	1	0
0	-1	1	0
0	1	-1	1
0	70	90	0

Figure 122: Right side of the tableaux

Tracking the simplex solve

The app allows you to track the actions of the simplex algorithm by saving the tableaux that the simplex algorithm uses. After the simplex algorithm has finished solving, the tableaux can be exported to a csv file, e.g., as shown in Figure 121 and Figure 122.

We are going to export the tableaux for the small model. Before we click the Solve Now button the app needs to know that we want to save the tableaux. The tableaux are not saved by default, because for larger models the size of the tableaux grows quite quickly and writing them out will slow

down the time that it takes to solve the model (or if they are very large, the app will run out of memory).

Choosing to save tableaux

Tapping the Solve button takes you to the Solve Options display, which includes the “Save Tableaux” option shown in Figure 123. The “Some” option will save the first, second and last tableaux. Select the “All” option, then solve.



Figure 123: "Save Tableaux" on the Solve Settings display

Accessing the saved tableaux

To view the saved tableaux, after the solve has completed you will need to export them as a csv file via email. Go to the Results display and then tap the “Import Export” button, indicated in Figure 124.



Figure 124: The Import Export button

The “Import Export” icon leads to the “Import Export” display shown in Figure 125. To export the tableaux, tap the “Email Results” button.

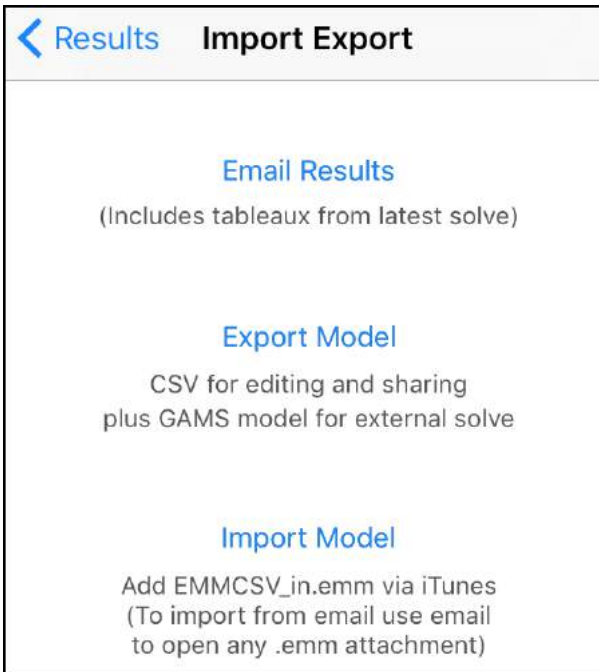


Figure 125: The Import Export display... use Email Results to export the tableaux

The “Email Results” email will contain several attachments... a screenshot of the model, a csv file containing the results of the latest solve and (assuming that there are saved tableaux) a csv file containing the saved tableaux from the solve. You can open the tableaux csv in a spreadsheet.

Layout of the tableau

The heading for each tableau describes the actions that the algorithm will apply to that tableau in order

to improve the objective value and thereby arrive at the next tableau, e.g., for Tableau00 the heading is “col[1] will enter via row[2] and col[4] will leave”. What this means is explained below, as we track the progress of the solve.

Each variable has its own column in the tableau. Each constraint has its own row. The label for each row includes the name of the constraint, the row number, and the column number of the *basic variable* (explained below) for that constraint.

Initial Tableau

The simplex algorithm works to maximise the objective value while meeting the requirements of the constraints. To do this the simplex algorithm converts the equations, which are supplied to the algorithm in the standard form described above, into a set of *equality* constraints that represent a feasible solution. These are the equations represented by the tableau.

Slack variables

Each tableau represents a valid solution that meets the requirements of the constraints. To create the initial tableau the first thing that the simplex algorithm does is to convert the \leq constraints to equality constraints (constraints that have an $=$

sign) by adding slack variables. Starting with the inequality constraint, e.g.,

$$-loadBid_{cleared} + genOffer_{cleared} \leq 0$$

...the simplex algorithm adds a slack variable, s_1 . The slack variable can take any non-negative value and this allows the inequality \leq in the original equation to be replaced with an equality, to create a new constraint:

$$-loadBid_{cleared} + genOffer_{cleared} + s_1 = 0$$

This meets the requirements of the original constraint, provided that the RHS is non-negative and all variables are ≥ 0 . These requirements will be enforced by the rules of the algorithm.

Adding slack variables

Adding slack variables produces an equivalent set of equality constraints:

$$-loadBid_{cleared} + genOffer_{cleared} + s_1 = 0$$

$$loadBid_{cleared} - genOffer_{cleared} + s_2 = 0$$

$$loadBid_{cleared} + s_3 = 100$$

$$genOffer_{cleared} + s_4 = 250$$

Basic Feasible Solution (BFS)

From the equations shown above, the simplex algorithm creates the initial feasible solution by

setting all the slack variables to the RHS value and all other variables to zero:

$$s_1 = 0$$

$$s_2 = 0$$

$$\text{loadBid}_{\text{Cleared}} = 0$$

$$s_3 = 100$$

$$\text{genOffer}_{\text{Cleared}} = 0$$

$$s_4 = 250$$

This is represented by the initial tableau from the csv, i.e., Tableau01, reproduced in Table 9.

Table 9: Constraints in the initial tableau

Basic	$\text{loadBid}_{\text{Clr d}}$ [1]	$\text{genOffer}_{\text{Clr d}}$ [2]	s_1 [3]	s_2 [4]	s_3 [5]	s_4 [6]	RHS
[3]	-1	1	1				0
[4]	1	-1		1			0
[5]	1				1		100
[6]		1				1	250

Basic variables

In this first Basic Feasible Solution (BFS) all the variables have their *value* set to zero except for the slack variables. The solution is feasible because

each constraint only has one non-zero variable, which is the slack variable.

The variables that are set to zero are referred to as the non-basic variables. The remaining variables, currently represented by the slack variables, constitute the basis of the solution and are referred to as the basic variables.

One basic variable per constraint

The basic variables are the variables that can be non-zero. Each constraint has its own basic variable, this basic variable has a factor of 1.0 in the constraint where it is the basic variable, and a factor of zero in all other constraints.

In the tableaux file, each row records the column that contains the basic variable for that row's constraint, e.g., for row [1] the basic column is column [3].

Note that while the non-basic variables are assigned a *value* of zero, they can still have a non-zero *factor*. For example, the LoadBidCleared variable has non-zero factors in three of the equations, but it is non-basic so the LoadBidCleared variable has a *value* of zero by definition of it being non-basic.

The objective function including slack variables

While this initial solution is feasible, the objective value as calculated by Equation 30 is zero because it was defined in terms of the original variables, which are currently all non-basic and therefore have values of zero. The slack variables have been added to the objective value, with factors of zero.

Equation 30: The objective, incorporating the slack variables

Maximize:

$$\begin{aligned} &150 \times loadBid_{cleared} \\ &-100 \times genOffer_{cleared} \\ &+ 0 \times s_1 + 0 \times S_2 + 0 \times s_3 + 0 \times S_4 \end{aligned}$$

Improving the objective value

Starting with the initial Basic Feasible Solution (BFS), the simplex algorithm moves to another BFS that improves the objective value. Improving the objective value is achieved by identifying the non-basic variable whose increase would most improve the objective value, and then making it a basic variable, by exchanging it with one of the existing basic variables.

Keeping track of the objective value

To keep track of the objective value, the objective function is re-written as a constraint. Using z to represent the objective value, the objective function as a constraint is shown in Equation 31.

$$z - loadBid_{cleared} \times loadBid_{price} + genOffer_{cleared} \times genOffer_{price} = 0$$

Equation 31: Writing the objective function as a constraint with the objective value as variable z

Table 10 shows the initial tableau with the objective function added as a constraint.

Table 10: Constraints and objective as a tableau

z	$loadBid_{clrd}$	$genOffer_{clrd}$	s_1	s_2	s_3	s_4	RHS
	-1	1				1	0
	1	-1			1		0
	1			1			100
		1	1				250
1	-160	70					0

The algorithm progresses by making non-basic variables basic to improve the objective, while obeying the rule that there is one basic variable per constraint and the basic variable has a factor of zero in all other rows. This is achieved by a series of row

operations. This will maintain the relationship between the variables as enforced by the original set of constraints. The row operations will also update the objective.

It turns out that we don't need to display the z column, because other than z there are no other values in this column so row operations won't change the z factor, and we won't be scaling the objective row, hence the factor of z will always be one.

The initial tableau can now be presented as shown in Table 11, which lines up with the csv export shown in Figure 121 and Figure 122.

Table 11: The initial tableau, incorporating the objective

Basic	$loadBid_{clrd}$ [1]	$genOffer_{clrd}$ [2]	s_1 [3]	s_2 [4]	s_3 [5]	s_4 [6]	RHS
[3]	-1	1	1				0
[4]	1	-1		1			0
[5]	1				1		100
[6]		1				1	250
	-160	70					0

The basic variables will also have factors of zero in the objective function constraint. As the other variables in the objective constraint are non-basic

and have *values* of zero, the only variable that is not set to zero and has a non-zero factor is the objective value z , which is therefore equal to the RHS.

Currently the RHS of the objective row is zero and therefore the objective value is zero. The algorithm will now work to improve the objective value.

Reduced costs

The only variables in the objective function constraint that have non-zero *factors* are the non-basic variables.

The non-basic variables have *values* of zero. If one of these variables became non-zero then the objective value z , (hidden off to the left with its factor always one) would need to change, in order to keep the constraint balanced.

For example, if the $loadBid_{clrd}$ variable, which has a factor of -150 in the objective function constraint, were to take on a value of 1.0 then the objective value RHS (and therefore z) would need to take on a value of 150 in order to keep the objective constraint balanced.

Likewise, if the $genOffer_{clrd}$ variable with a factor of 70 were to take on a value of 1.0 then the objective value would need to be -70.

Hence, the *factors* in the objective function constraint indicate which variable is the best to increase in order to improve the objective value. These factors are referred to as the *reduced costs*.

In this case only the $loadBid_{clrd}$ variable can improve the objective value. Because it has a reduced cost of -150, each unit of increase by the $loadBid_{clrd}$ variable improves the objective value by 150.

The entering variable and the leaving variable

To increase the objective value, the simplex algorithm chooses the variable with the most negative reduced cost, in this case $loadBid_{clrd}$, and then increases the value of that variable as much as possible. Because the only variables with non-zero reduced costs are the non-basic variables (assigned a value of zero), in order for the value of this variable to be increased it needs to become a basic variable.

However, the tableau is feasible by virtue of the fact that there is only one basic variable for each constraint, so in order for the non-basic variable to become basic, i.e., to *enter the basis*, one of the basic variables will need to become non-basic, i.e., *leave the basis*.

For the non-basic variable to enter, the tableau will need to be manipulated so that canonical form is maintained, i.e., so that the newly basic variable has a factor of 1.0 in the entering constraint and zero in all other constraints, including the objective function constraint.

Determining the entering constraint

The entering constraint is selected by finding the row that allows the factor of the entering variable to be set to 1.0 in that row and zero in all other rows.

What drives this selection is not what happens in the entering row, because there it is a simple division to obtain a factor of one, but rather what happens to the other rows in order to set the entering variable's factor to zero in those rows.

The first step to making the entering variable basic is to divide the entering row by the factor of the entering variable in that row, so that its factor becomes 1.0. Then the row operations will subtract a multiple of the entering row from the other rows so that the entering variable's factor becomes zero in those rows.

Row operations example on a dummy tableau

To demonstrate how and why the algorithm determines the entering row, we will use the

dummy tableau shown in Table 12. We are using a dummy tableau because the numbers in our actual model don't lend themselves to explaining the issue.

The entering row is dictated by the requirement that variables must be non-negative (remember that this requirement arose when the original inequality constraints were replaced with slack variables and equality constraints).

The algorithm cannot select an entering row that would cause any of the basic variables to be assigned a negative value (the *non*-basic variables have values of zero regardless of the equations, so we don't have to worry about them).

Table 12: Dummy tableau to demonstrate entering and leaving variables

	Col#1	2	3	4	5	RHS
	<i>var1</i>	<i>var2</i>	<i>var3</i>	<i>var4</i>	<i>var5</i>	
Row#1	-7	1				1
2	5		1			20
3	10			1		1
4	40		1		1	8
5	-150	70				0

In the dummy example the entering variable is in column one (because it has the most negative reduced cost). When selecting the entering row we obviously can't use the row with the -7 factor because to make the factor 1 would result in a RHS of $-1/7$, and then the value of the basic variable would be $-1/7$, which cannot happen because all variables must be ≥ 0 .

If row 2 was selected, then in order for the entering variable to take on a factor of 1.0 in that row, the row would be divided by 5.0, and the RHS becomes 4.0. Row 2 and 3 would then be as follows:

	Col#1	2	3	4	5	RHS
Row#2	1		1/5			4
Row#3	10			1		1

Row 2 is OK, but now to give the entering variable a factor of zero in row 3 we need to subtract $10 \times$ row 2 from row 3. This will result in a RHS value of -39 for row 3 and then the basic variable for row 3 would become negative... so we can't do this.

Therefore, in order to prevent a negative RHS for any of the other rows, the entering row must be the row with the smallest positive ratio of RHS to entering variable factor... so that when any other row subtracts the entering row (in order to zero the

entering variable in that row), the RHS of that row will remain positive.

In our dummy example, the row with the smallest positive ratio of RHS to entering factor is row 3, where the ration is $1/10$. After the entering variable has entered via row 3, the tableau is shown in Table 13.

Table 13: Dummy tableau after pivot

	Col#1	2	3	4	5	RHS
Row#1	0	1		0.7		1.7
#2	0		1	-0.5		19.5
#3	1			0.1		0.1
#4	0		1	-4	1	4
Obj	0	70		15		15

The new tableau after the pivot

Back to our small model, the column with the smallest ratio of RHS to entering factor is row 2, with a ratio of zero. This is the same as the ratio for row 1... for consistency if more than one row has the same ratio then we always use the last row that we find, i.e., row 2 in this case (it does not matter how we make the decision as long as we are consistent).

After row operations have been performed to allow $loadBid_{clrd}$ to become a basic variable, the tableau appears as shown in Table 15.

The entering variable is $loadBid_{clrd}$. The leaving variable is the variable that was previously basic in the entering row. In this case the leaving variable is the slack variable s_2 , which becomes a non -basic variable.

Now that $loadBid_{clrd}$ is a basic variable it is no longer confined to a value of zero. However, the equation where $loadBid_{clrd}$ is basic has a RHS of zero, therefore the $loadBid_{clrd}$ variable is assigned a value of zero by the equation.

We can also see from the RHS of the objective constraint that the objective value is still zero.

Table 14: Initial tableau (again)

Basic	$loadBid_{clrd}$ [1]	$genOffer_{clrd}$ [2]	s_1 [3]	s_2 [4]	s_3 [5]	s_4 [6]	RHS
[3]	-1	1	1				0
[4]	1	-1		1			0
[5]	1				1		100
[6]		1				1	250
	-160	70					0

Table 15: Second tableau

Basic	$loadBid_{clrd}$ [1]	$genOffer_{clrd}$ [2]	s_1 [3]	s_2 [4]	s_3 [5]	s_4 [6]	RHS
[3]			1	1			0
[1]	1	-1		1			0
[5]		1		-1	1		100
[6]		1				1	250
		-90		160			0

Updating the objective value

Bringing the entering variable into the basis included a row operation to set its reduced cost to zero. The reduced cost of zero indicates that the variable has been increased as much as possible and cannot improve the objective value any further. The row update that set the reduced cost to zero also updated the objective value.

The reduced cost of the entering variable was -150, hence to zero the reduced cost, the row operation added 150 times the entering row to the objective row. This added 150 times the RHS of the entering row to the RHS of the objective row. Because the RHS of the entering row is the value of the entering variable, this has the effect of adding the entering

variable to the objective value... which was the intent of the iteration in the first place.

Adjusting the reduced costs

The first step of the row operations was to scale the entering row so that the entering variable has a factor of one. This also scaled all other factors in that row.

Apart from the entering variable, the only other variables with non-zero factors in the entering row are non-basic variables. After scaling, the entering variable has a factor of one and the factor of each non-basic variable in the row now indicates how much a change in their value would force the entering variable to change (if they were to become non-zero).

After we scaled the entering row to make the entering factor one, we added that row, multiplied by the reduced cost of the entering variable, to the objective function, this was to zero the reduced cost of the entering variable. This updated the reduced cost of all the non-basic variables based on how much freedom they can provide the entering variable, and the value of the entering variable.

Looking at tableau 2 in Figure 126 and Figure 127 (the tableaux csv again) we can see that the update

of the objective function has changed the reduced cost of the cleared offer from positive to negative, because its potential to allow the cleared bid to increase has overcome the cost of the offer.

###TABLEAU[0001]### col[1] will enter via row[2] and col[4] will leave			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	-1	1
[2]_{NodeBalanceGTE} basic col[4]	0	1	-1
[3]bid00_{BidBlockMaxLTE} basic col[5]	100	1	0
[4]offer00_{OfferBlockMaxLTE} basic col[6]	250	0	1
[5]ObjectiveFn	0	-160	70
###TABLEAU[0002]### col[2] will enter via row[3] and col[5] will leave			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	0	0
[2]_{NodeBalanceGTE} basic col[1]	0	1	-1
[3]bid00_{BidBlockMaxLTE} basic col[5]	100	0	1
[4]offer00_{OfferBlockMaxLTE} basic col[6]	250	0	1
[5]ObjectiveFn	0	0	-90
###TABLEAU[0003]### FINAL			
Variables->	RHS	[1]bid00_{Cleared}	[2]offer00_{Cleared}
[1]_{NodeBalanceLTE} basic col[3]	0	0	0
[2]_{NodeBalanceGTE} basic col[1]	100	1	0
[3]bid00_{BidBlockMaxLTE} basic col[2]	100	0	1
[4]offer00_{OfferBlockMaxLTE} basic col[6]	150	0	0
[5]ObjectiveFn	9000	0	0

Figure 126: Left side of the tableaux csv

[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1
0	0	0	0
[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	1	0	0
0	1	0	0
0	-1	1	0
0	0	0	1
0	160	0	0
[3]slack{NodeBalanceLTE}	[4]slack{NodeBalanceGTE}	[5]slack_bid00_{BidBlockMaxLTE}	[6]slack_offer00_{OfferBlockMaxLTE}
1	1	0	0
0	0	1	0
0	-1	1	0
0	1	-1	1
0	70	90	0

Figure 127: Right side of the tableaux csv

Why the objective value is still zero

Now that the cleared bid is a basic variable it can take a non-zero value but because the RHS of its basic constraint is zero the cleared bid is constrained to zero, and therefore the objective value is still zero.

The next entering variable

Now that the cleared bid variable is basic it is constrained by the RHS of its constraint. The cleared bid can only increase if enabled to do so by another variable in this constraint. The other variable would need to have a negative factor in the constraint... increasing the value of this other

variable would allow the cleared bid variable to increase.

We know the contents of the node balance constraint in this model, i.e., cleared bids out of the bus must be matched by cleared offers in, so it is no surprise that the variable that will allow the cleared bid variable to increase is the cleared offer variable.

The reduced cost for cleared offers is now $-\$90$, i.e., the difference between the $\$70$ cost of the cleared offer (i.e., its original reduced cost) and the $\$160/\text{MW}$ benefit resulting from the increase in cleared bids that the cleared offer would enable.

This difference was calculated when the row operations zeroed the reduced cost of the cleared bid, passing the per MW benefit of the cleared bid along to the cleared offer, offsetting the existing cost of the offer.

Now that the cleared offer variable is the variable with the most negative reduced cost, it will be the next entering variable. The next (and, in this small model, final) step is for the cleared offer variable to enter the basis. This will follow the same procedure as above... find the entering row as the row which will allow the necessary row operations without making any basic variables negative, and then perform row operations in order to allow the

entering variable to become basic (by meeting the requirement that it has a factor of 1 in the entering row and zero in all others).

A closer look at the entering row

The entering variable becomes the basic variable in the row with the lowest ratio of RHS to factor, because choosing any other row would create a negative RHS.

Looking at the numbers, we can relate this to the reality of our model. One of the constraints in tableau02 is the offer max constraint, which in the initial tableau prevents the offer from clearing more than its offered quantity. After the first pivot, this constraint now acts to limit the cleared offer to be no more than the cleared bid.

The row operations of the pivot effectively moved the relationship between the bids and offers from the node balance constraint to the bid max constraint.

This “new” constraint, which limits the cleared offer to be no more than the cleared bid (which is the bid max), becomes the entering constraint for the cleared offer. We can see that if we had allowed the cleared offer to enter via the other constraint, i.e., the offer max constraint, then the cleared offer

quantity would be more than the cleared bid quantity, breaking the node balance.

The algorithm can't see this, but it selects the bid max row as the entering row for the cleared offer because it is the constraint with the lowest ratio of factor to RHS, which means the same thing as our observation... it is the constraint that will bind first when the entering variable is increased that determines the entering row.

Solve complete

When the cleared offer enters the basis, the RHS of the bid max constraint sets the value for the cleared offer. Because there are now no negative reduced costs there is no way to further improve the objective, hence the solve is complete.

"The tremendous power of the simplex method is a constant surprise to me" - George Dantzig (inventor of the simplex algorithm)

Extracting the results from the final tableau

Quantities from RHS values

In the final tableau shown in Table 16 the cleared offer and cleared bid variables are both basic, i.e., able to be non-zero.

The cleared bid (column 1) is basic in row 2, with a RHS of 100, the cleared offer (column 2) is basic in row 3, also with a RHS of 100. These RHS values provide the values of the basic variables.

Table 16: The final tableau

	$loadBid_{clrd}$	$genOffer_{clrd}$	s_1	s_2	s_3	s_4	RHS
Basic	[1]	[2]	[3]	[4]	[5]	[6]	
[3]			1	1			
[1]	1				1		100
[2]		1		-1	1		100
[6]				1	-1	1	150
				70	90		9000

Shadow price of constraints

As discussed in Tutorial 1: Explaining Prices, the shadow price of a constraint is the rate of change of the objective value due to relaxing that constraint. The shadow price therefore indicates the value of the quantity that the constraint is constraining.

The bus price, i.e., the value of power at a bus, is the shadow price of the node balance constraint... the node balance constraint constrains the flow of power into and out of the bus, therefore the value of relaxing the node balance constraint tells us the value of electrical power at the bus.

Shadow price from final tableau

The shadow prices can be extracted from the final tableau.

The values in the final tableau represent the optimum objective value and also meet the requirements of the original constraints. For example, given the original set of constraints (shown in Equation 32) we can plug in the values from the final tableau and the equations will solve.

Equation 32: Constraints of the linear programme (again)

$$-loadBid_{cleared} + genOffer_{cleared} + s_1 = 0$$

$$loadBid_{cleared} - genOffer_{cleared} + s_2 = 0$$

$$loadBid_{cleared} + s_3 = 10$$

$$genOffer_{cleared} + s_4 = 20$$

Using the values from the final tableau, the second node balance constraint is shown in Equation 33.

$$\begin{array}{ccccccc} loadBid_{cleared} & - & genOffer_{cleared} & + & s_2 & = & 0 \\ 100 & - & 100 & + & 0 & = & 0 \end{array}$$

Equation 33: Node balance with final tableau values

We can see from Equation 33 that decreasing the value of slack variable s_2 relaxes the node balance constraint, and therefore the impact on the

objective value of decreasing s_2 in the final tableau will provide the shadow price of the node balance constraint in the final tableau, and therefore the bus price.

To find the impact on the objective value of decreasing s_2 remember that the progress of the algorithm was driven by looking for negative reduced costs because they indicate the rate at which the objective value will improve if we increase the associated variable.

From this it follows that a positive reduced cost indicates the rate at which the objective will *decrease* if we increased the variable. And also, the observation that helps us here, a positive reduced cost indicates the rate at which the objective value will increase if we could *decrease* the variable.

The positive reduced cost of s_2 provides the rate of objective value increase that would result from decreasing s_2 . Decreasing s_2 relaxes the node balance constraint that originally contained s_2 . Therefore, the positive reduced cost of s_2 provides the improvement to the objective value due to relaxing the node balance constraint.

From the final tableau we can see that the reduced cost of s_2 is 70, which is therefore the shadow price of the node balance constraint.

The above is true of any of the constraints. The positive reduced cost in the final tableau of the constraint's original slack variable will provide the shadow price of the constraint.

Viewing the simplex iterations

As well as exporting the tableaux to see what the simplex algorithm did, you can also view details of the simplex solve from within the app.

Iteration details

On the Results display select the Iterations row, this will take you to the iteration details as displayed in Figure 128 (these are the iterations from the small model that we have just studied).

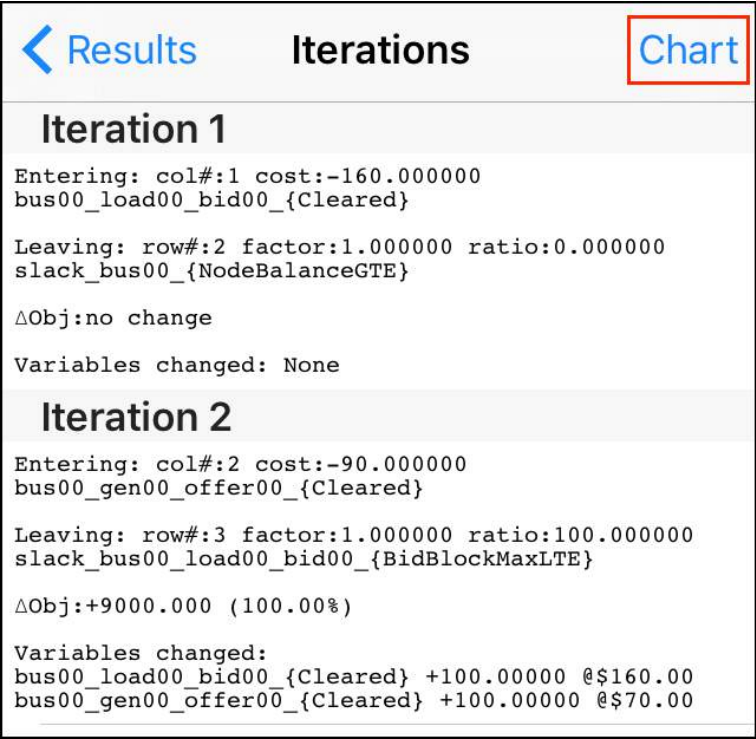


Figure 128: Iteration details display

Graphical view of iterations

The simplex iterations can also be viewed as a chart by tapping the Chart button on the Iterations display (this chart is also displayed while the solver is solving). The chart shows iterations vs objective, as displayed in Figure 129. This chart shown is from the Hawkes Bay 033 sample case which was presented in Tutorial 8: Actual Market Data.




Figure 129: Iterations vs objective (shown here for the Hawkes Bay 033 model)




Figure 130: Zoom in to view individual iterations

Use the pinch gesture to zoom in and view the individual iteration points, as shown in Figure 130. Tapping on an iteration point will present the

simplex actions that produced this result, see the example in Figure 131.

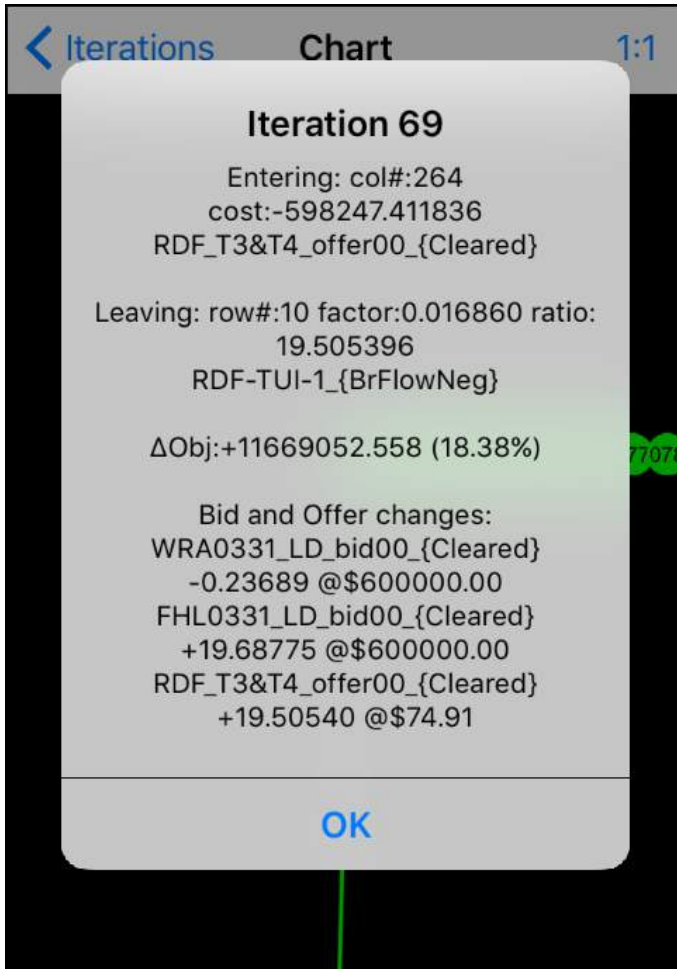


Figure 131: Tap iteration point to view simplex actions

Summary

In this section we built and solved a small electricity model then exported its tableaux and tracked in detail how the simplex algorithm arrived at the result. Along the way we explained how and why the simplex algorithm makes the decisions that it does, and why it works.

When the result was complete, we saw how the prices and quantities are extracted from the final tableau.

We also saw how to view the simplex iterations in the app, either via a description of the iterations, or graphically as a chart of objective value vs iteration count.