Module 11: Linear Programming

Learning Goal for Module 11

Module 11 presents the linear programming model, which serves as the foundation for operations research. It also presents the basic linear optimization model in microeconomics. This module will help you refine your ability to set up and solve elementary linear optimizations.

By the end of this Module, you will:

- Recognize an optimization problem
- Be able to specify the objective function, technology matrix, and constraints
- Create the Solver process and identify the optima for the objective function, as well as the value of the change variables associated with this optimization point, and impose constraints on the solution

1. **Introduction**

"Analytics" encompasses two fields of study: 1) It is part of *operations research*, a field that has existed since 1940, and 2) More recently, it meant the statistical and machine analysis (learning) of large constantly updated datasets. Machine analysis is a generic term for advanced data processing that combines statistical analysis with a feedback process to improve the validity and reliability of the model parameters. *Operations research (OR)* is a set of mathematical models intended to find the "best" solution for production, consumption, and distribution of goods and services in complex systems.

"Best" means:

- Maximum output, revenue, profit…
- Minimum cost (resources, labour, time…)
- Reaching a defined target (output, cost…)
- Decision processes to select from among multiple outcomes

Module 11: Linear Programming

Precursors to the development of formal methods to allocate scarce resources included:

- *Joseph Fourier* (French mathematician and physicist) developed early partial solutions for the linear program (LP) problem.
- *George Kantorovich* created the first linear programming solution in 1937 as a method to reduce costs to the army while increasing costs to the enemy.
- *George Dantzig* developed the simplex method (1947), the first efficient method to solve LP problems.

What is an LP problem?

How can I allocate 70 people (each with different skills) to 70 jobs with different requirements to maximize output?

This is a complex problem to solve directly, but it becomes solvable when specified as an LP model.

This module introduces a fundamental OR technique in the form of the LP model.

1. Structure of the LP model

Examples illustrate the types of problems linear programming can solve:

- ***Example 1***: In the 1990s, Merrill-Lynch needed to adapt to competition from discount brokerages. This meant finding efficiencies and developing new services and price structures. It needed to process five million client records (over 250 million trades) to identify the price-service combinations that could increase revenues.

- ***Example 2****:* Swift & Company, a protein producing business (beef and related products) created 45 linked LP programs to enable its sales reps to manage more than 8,000 customers, optimize the shifts at its plants, and to manage production and shipping schedules. This process realized increased sales of $12 million.
- ***Example 3****:* Memorial Sloan Kettering, the premier private cancer treatment centre in the US, used OR techniques to optimize the treatment of prostate cancer using a form of LP modelling known as mixed integer programming.

Linear program
- A mathematic model represents optimization problem.
- Optimization means *maximizing* (sales), *minimizing* (costs), or *meeting* a goal
- Linear equations represent all aspects of the mathematical problem.
- *Program* refers to a structured search for the solution or a *plan* (if it exists) to find the best solution.

A linear program formulates all problems with linear equations.
Linear programming problems usually have a solution (unless they are *over-constrained*).

The starting point is to define the objective function: What are we are trying to maximize, minimize, or meet?

Example: A pizza business tries to *maximize* sales. With one pizza type (Q_1), sales are simply

$$\text{Sales} = P_1 Q_1$$

With two types of pizza, sales are

$$P_1 Q_1 + P_2 Q_2$$

A student wishing to maximize their GPA allocates study time among the courses that will produce the highest marginal increase in GPA for an additional hour of studying, subject to (st) the constraint that only 168 hours exist in a week, and that biological maintenance requires a specific minimum time each week.

1.1. The objective function

The objective function typically involves one of three goals – maximization, minimization, and equality to a specific value, often 0. For example, maximizing revenues is easy once we know the elasticity of demand. Maximizing profits is much harder, since profits are a function of prices times the quantities produced and costs. Costs depend on the cost of inputs and the technology needed to create the good or service. We never have unlimited capacity to change quantities, or produce at any level, or charge whatever price we wish.

Maximizing or minimizing something is always subject to a constraint, otherwise profits would have no limit. Even technology places limits on the capacity to produce any product, and despite those who believe they multi-task, a physician usually only sees one patient at a time.

With one product and one input, the determination of a maximum by pencil and paper calculations is easy, but as we add more products to the mix, identifying the combination of prices and quantities that supports a unique maximization requires trial and error, becoming a tedious exercise. For example, a pizza restaurant may vary the prices to identify a "sweet spot" in sales. Often the pizza owners will adopt a multidimensional strategy, such as varying the prices by topping, reflecting both the cost of the topping and its popularity (demand elasticity again), adjusting the quality of ingredients, and changing the size.

The original application of LP was to optimize diets for an army where the goal was to minimize costs by producing a varied diet subject to the condition that each soldier needed a minimum number of calories per day.

1.2. Constraints are everywhere

The mathematical expression of the LP model has **linear** constraints, such as:

$$Q_1 + Q_2 = Q_T$$
$$Q_1 > a$$
$$Q_2 \leq b$$

We face constraints everywhere

- The day before a final exam, you face the constraint of having only 24 hours to study.
- Purchasers may set a maximum budget (constraint) and/or insist on a red car.
- A manufacturer faces limits the quantity of labour time available from a worker or how much raw material accessible in a day.

Consumers have budget and time constraints, while producers have input (cost and availability) and margin constraints (profits must be positive).

The constraints define the "feasible" region for a solution. Most often, maximization problems, such as maximizing revenues, have "less than" (≤) constraints, while minimization problems (costs) typically have "greater than" (≥) constraints. On rare occasions, LP problems in economics have equality constraints.

Video: Graphing a linear equation

Video: Sliders and spin bars

Example: Maximize the total revenue from two products that use two inputs, X_1 and X_2, with fixed prices, and the constraints on X1 and X2 are:

$$2X_1 \leq 12$$
$$3X_2 \leq 15$$
$$2X_1 + 2X_2 \leq 9$$

The linear constraints each appear as an inequality shaded area, with negative values excluded.

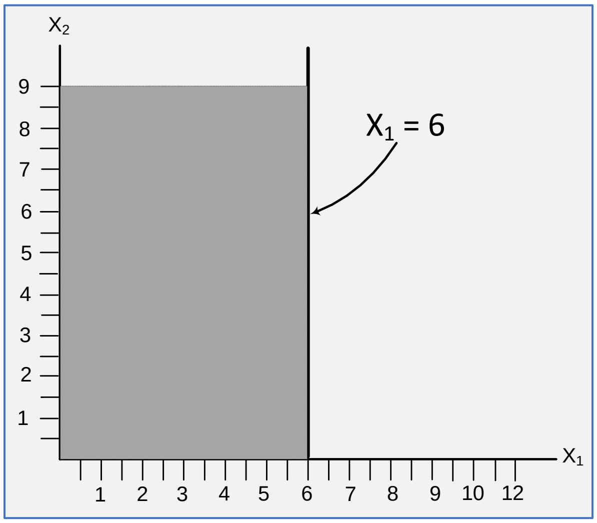

Figure 2:Constraint 1

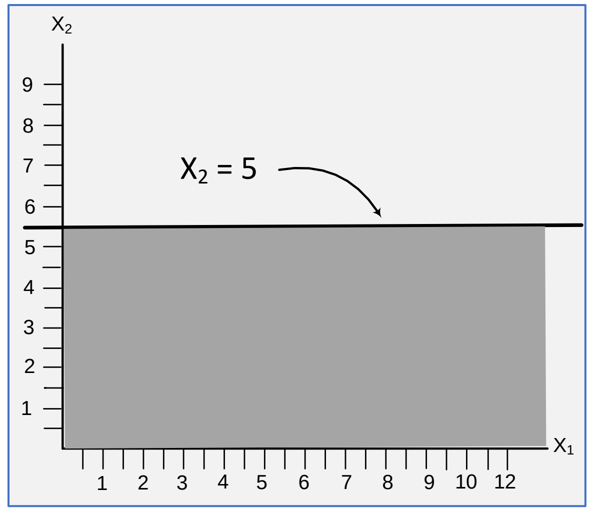

Figure 1: Constraint 2

Dividing through by the lowest common denominator creates the mathematical equivalent.

Video: [Working with constraints]

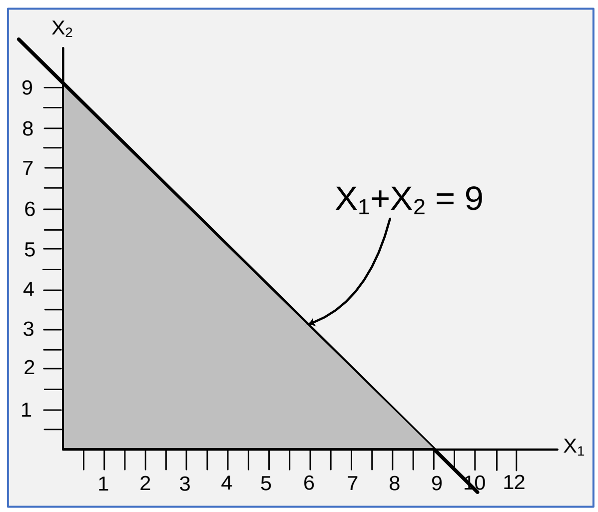

Figure 3: Constraint 2

Combining the three constraints forms the feasible region (dark line). Any solution within the feasible region will satisfy the optimization. Most often, solutions lie along one of the faces of the feasible region or at a "point."

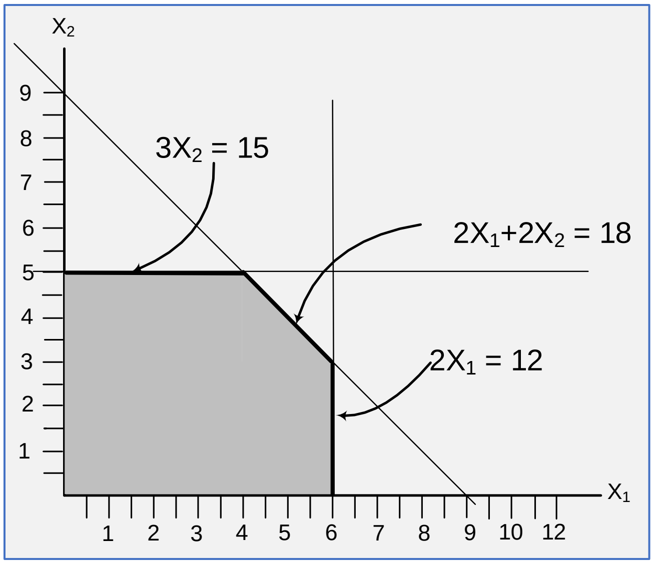

Figure 4: All constraints combined

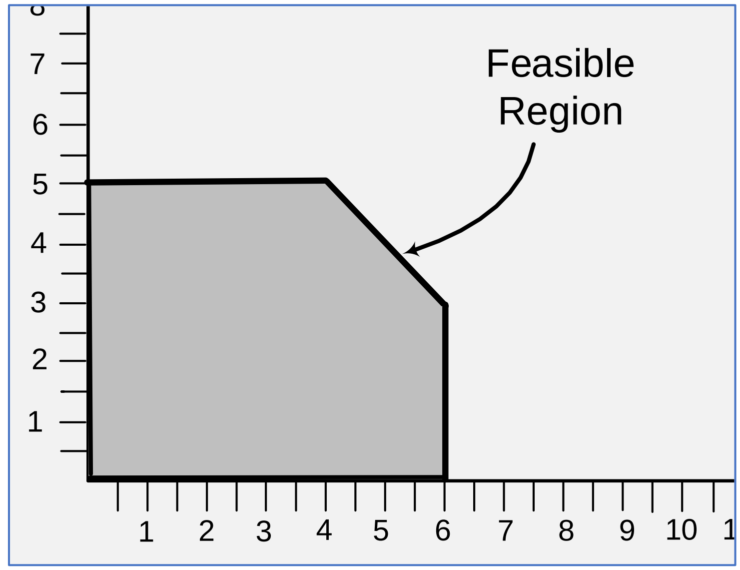

Figure 5 Feasible region

1.3. Adding the objective function completes the LP problem

Now imagine the objective function were

$$3X_1 + 5X_2 = A,$$

which is purposely open, with A adjusting to shift the function parallel to itself. The goal for a maximization problem is to change (increase) A until the objective function touches an "outmost" point on the feasible region. This is the solution to the LP problem.

1.4. Irrelevant and binding constraints

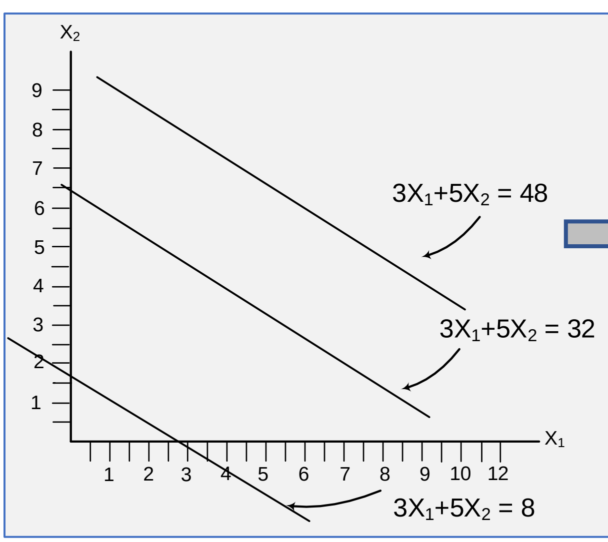

Figure 6: Shifting a single constraint

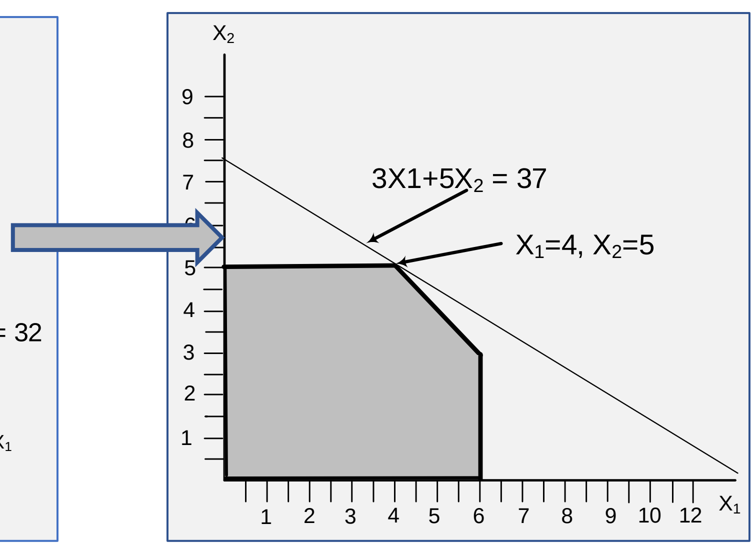

Figure 7: Creating a corner solution t

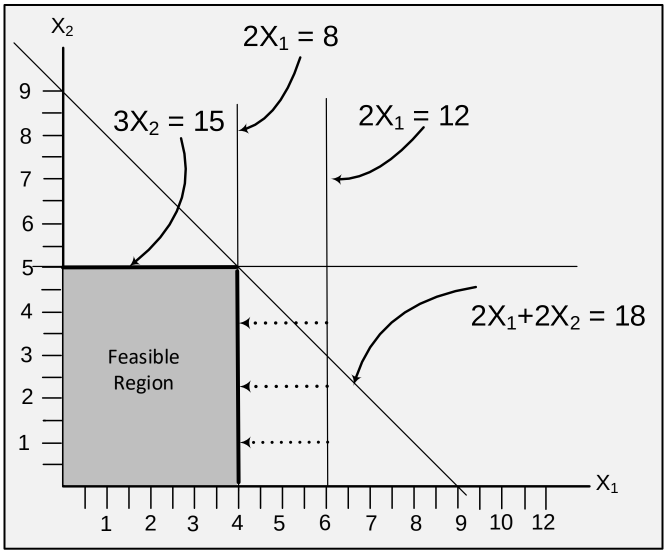

Figure 8: Irrelevant constraint

Using ABC Manufacturing, if we change the production time available for Plant 1 from 12 to 8, this shifts the constraint to the left to the line $2X_1=8$. At this point, the constraint $2X_1+2X_2=18$ becomes *irrelevant.* The optimum stays the same.

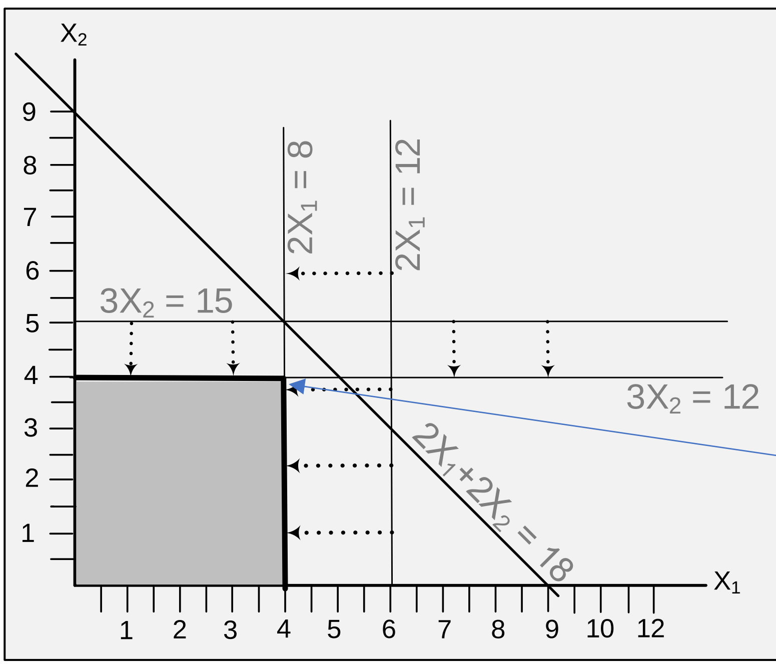

Figure 9: Binding constraint

But imagine that for Plant 2, availability dropped to 12 hours from 15. The line for that constraint now becomes $3X_2=12$ or $X_2 = 4$. It is easy to see that the maximum for the objective function also declines. The second constraint becomes more *binding* and dominates the third constraint ($2X_2+2X_2=18$). The optimum also contracts to 4 and 4 for X_1 and X_2.

2. Using Solver to locate solutions to the LP Problem

The LP problem will appear as a "word" puzzle.

Two examples

Example: ABC Door Manufacturing

ABC produces exterior doors – solid wood (Product 1) and steel-clad composite doors (Product 2) at three plants.

- Plant 1 produces only wood doors.
- Plant 2 produces only steel doors.
- Plant 3 produces hardware, frames, and fabricates the doors.

ABC production technology has the following attributes:

- solid wood doors need the output of Plant 1 and Plant 3, but not of Plant 2; and
- steel-clad composite doors need the output of Plants 2 and 3, but not 1.

Management wishes to maximize profit in the context of the following constraints:

- The total number of hours of production time at each plant are 12, 15, and 18, respectively.
- The hours needed for each product are two hours per Door 1 at Plant 1, three hours per Door 2 at Plant 2, and two hours for both doors at Plant 3.
- Profit per batch of product is $3,000 for Door 1 and $5,000 for Door 2.

The technology matrix appears as shown in Figure 10. This becomes the following LP problem:

$$\begin{aligned} \text{Max } P &= 3X_1 + 5X_2 \\ \text{st} \quad 2X_1 &\leq 12 \\ 3X_2 &\leq 15 \\ 2X_1 + 2X_2 &\leq 18 \end{aligned}$$

Plant	Production time per batch (hours)		Available time per week (hours)
	Product		
	1	2	
1	2	0	12
2	0	3	15
3	2	2	18
Profit/Batch	$3,000	$5,000	

Figure 10: Technology and profit matrix

> Profits are measures in $thousands, which is why the objective function appears as
>
> Max P = $3X_1+5X_2$

The ABC Manufacturing problem is easy to specify in Solver:

1. Create two columns for the X_1 and X_2 values.
2. Write the objective function in a cell below the column (or anywhere on the sheet).
3. Now add the right-side values in three cells.
4. Activate Solver.
5. Enter the cell that locates the objective function formula.
6. Enter the constraints.
7. Solve.
8. If a solution exists, Solver will present the value of the objective function.

The key to setting up an LP problem is to construct the spreadsheet with embedded equations. The blue regions designate data, and the green are the change cells (values Solver iterates) to calculate the objective function in yellow. (See Figure 11.)

Pay attention to cells D7:D9, which construct the data for the constraints, and cell F12 that calculates the objective function. Cells B12 and C12 have initial values that can be 0.

The spreadsheet shows the results of the Solver process and shows that profit is maximized (37) when we produce four batches of wood doors and five batches of steel doors.

See ABC Manufacturing.xlsx

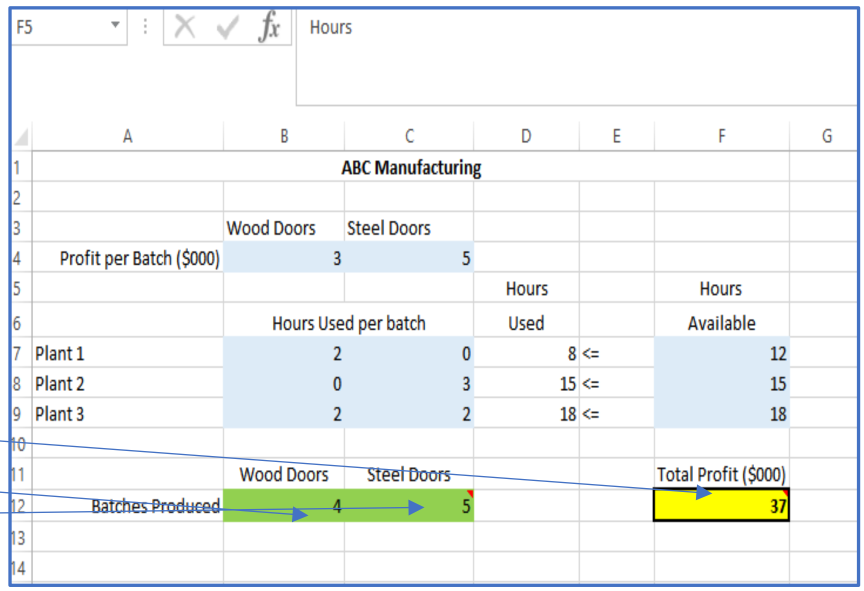

F5 Hours

	A	B	C	D	E	F	G
1			ABC Manufacturing				
2							
3		Wood Doors	Steel Doors				
4	Profit per Batch ($000)	3	5				
5				Hours		Hours	
6		Hours Used per batch		Used		Available	
7	Plant 1	2	0	8	<=	12	
8	Plant 2	0	3	15	<=	15	
9	Plant 3	2	2	18	<=	18	
10							
11		Wood Doors	Steel Doors			Total Profit ($000)	
12	Batches Produced	4	5			37	
13							
14							

Figure 11: Solver setup for ABC manufacturing

Video: ABC Manufacturing 1

Video: ABC Manufacturing 2

If no solution exists with a linear model (objective function and constraints are linear relationships), then it is possible that the problem is “over-constrained” or that the objective function and feasible region run parallel.

- *Example*: Optimizing a diet for space travel. One of the original applications of LP was allocation of food rations for an army. Here, imagine we are using freeze dried food for space travel. Four foods exist, with five attributes. The goal is to supply a minimum of 2,100 calories per day, with each daily ration providing at least 100 grams of protein, 200 grams of carbs, and 10 grams of fat, all at the lowest cost.

Ration Type	A	B	C	D
Calories	700	1500	850	1200
Protein	50	150	75	20
Carbs	100	75	200	300
Fat	5	7	3	12
Cost($/gram)	1	1.50	.75	2.50

The problem set up in Excel appears as follows – see Space Travel.xlsx

	A	B	C	D	E	F	G	H	I
1				Space Travel Diet					
2		Daily ration							
3			Food						
4		Beer	Chips	Popcorn	Donuts				
5	Consumption ('00 grams)								
6	Unit Cost ($/'00 gram)	$0.50	$1.50	$0.75	$2.20				
7	Total cost/day	$0.00							
8									
9	Calories (/100 grams)	700	1500	850	1200	0	>=	2100	
10	Protein (/100 grams)	25	25	74	20	0	>=	100	
11	Carbs (/100 grams)	100	300	200	350	0	>=	500	
12	Fat (/100 grams)	2	25	3	20	0	>=	30	
13									

Figure 12: Set-up for space diet

The technology matrix (cells B9:E12) reflects the nutritional composition of the four foods. The prices of the foods appear in cells (B6:E6) and setting these to blank – Solver will fill this in as part of the solution. The objective function appears in B7 as a SUMPRODUCT. The constraints appear in cells (H9:H12); the diet must meet these requirements. Cells F9:F12 show the nutrients delivered by the four foods. The goal is to calculate quantities of each food that minimizes cost, while equalling or exceeding the constraint values.

3. Creating the objective function

First, create the cost model. By adjusting the quantities of each food and using the prices, this part of the model calculates the cost of the diet (Cell B7). It is this cost that must be minimized. Cells B5:E5 will contain the quantities of each food that meet the conditions of the of the technology. This price of the foods appears in cells B6:E6.

The total cost (B7) is the price of beer (Pb) times the quantity of beer (Qb), plus the price of chips times the quantity of chips (Pc*Qc), etc.

Total cost = Pb*Qb+Pc*Qc+Pp*Qp+Pd*Qd = SUMPRODUCT(B5:E5,B6:E6)

This total cost appears in cell B7 and is the "objective function." Note that this cell contains the equivalent of the formula above, comprised of the fixed values for the prices and the adjustable values of the quantities. Note also that the quantities B5:E5 and B7 are 0 at the start and then change once Solver presents the solution.

Technology matrix

The technology matrix (B9:E12) contains the nutritional content of each food (reading down) or the distribution of each nutrient by food (reading across). Looking at the composition description in the problem, the derivation of each row and column in the technology matrix should be apparent.

Constraints

The constraint matrix is H9:H12; this is the ideal space diet. Looking at the technology matrix, no food has sufficient of any nutrient to meet the diet. Each constraint reflects the total contributed by the foods selected in meeting a specific nutrient. For example, cell F11 shows how the foods selected meet the calorie content. Cells F12, F13, and F14 show how the total diet meets each other nutrient requirement. Note the use of SUMPRODUCT and absolute cell references for cells B5:E5 (which supports quick copying and pasting). See

Constraints and Objective Function.xlsx

The next step is the Solver setup; pay attention to the construction of the constraints.

The Solver screen appears at the right. Note the specification of minimization and the direction of the inequalities. Also note that the cells specifying the objective functions and the left-hand side of the constraints are all formulas. The change cells are the quantities to be determined.

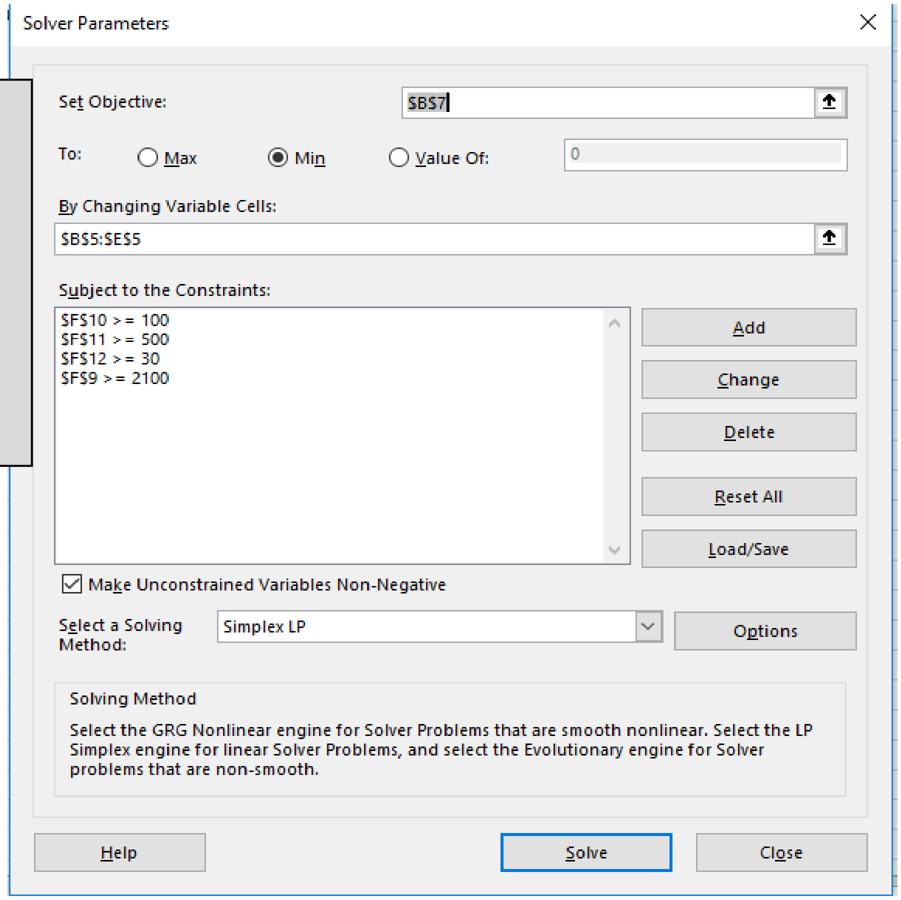

Figure 13: Solver set-up for space diet

Here is the solution. With these prices and technology matrix, we can meet/exceed the space diet with 1.08g of chips and .99g of popcorn, at least for cost. Beer and donuts will not make the trip. Constraints 2 and 3 dominate constraints 1 and 4. The cost minimizing diet offers more calories and more carbs than the minimum specifications. These astronauts might land with a thud after several months in space!

	A	B	C	D	E	F	G	H
4		Beer	Chips	Popcorn	Donuts			
5	Consumption ('00 grams)	0	1.08169014	0.98591549	0			
6	Unit Cost ($/'00 gram)	$0.50	$1.50	$0.75	$2.20			
7	Total cost/day	$2.36						
8								
9	Calories (/100 grams)	700	1500	850	1200	2460.563	>=	2100
10	Protein (/100 grams)	25	25	74	20	100	>=	100
11	Carbs (/100 grams)	100	300	200	350	521.6901	>=	500
12	Fat (/100 grams)	2	25	3	20	30	>=	30
13								

Figure 14: Creating the template for space diet

4. Minimization

Minimization problems (minimizing costs) will typically have constraints specified as $\geq$.

$$\begin{aligned} &\text{Min: } a_1x_1 + a_2x_2 \\ &\text{ST } b_1x_1 + b_2x_2 \geq B \\ &\quad c_1x_1 + c_2x_2 \geq C \\ &\quad d_1x_1 + d_2x_2 \geq D \end{aligned}$$

The feasible region appears unbounded in the northeast direction and may or may not be bounded by the two axes.

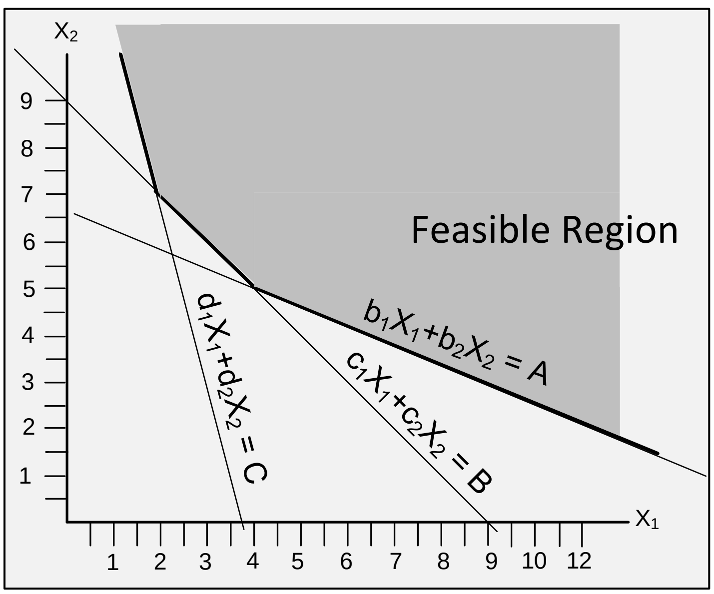

Figure 15: Feasible region for minimization problem

5. Summary

Optimization is integral to economic and business analytics. Solver offers a powerful method for locating minima/maxima/equality conditions. The essential skill is first to recognize the nature of the problem by specifying the technology matrix that supports the specification on the objective function and constraints. Next set up Solver to define the objective function, the change/driver variables that must appear in the objective function, and any constraints. This requires practice.

Annex: Key Excel functions and formulas

Function	Example	Explanation
=SUMPRODUCT(array1,array2, array3…)	=SUMPRODUCT(A1:A10,C1:L1)	This multiplies the vector A1:A10 by the vector C1:L1. The two vectors must have equal dimensions. Note: Like all Excel functions, =SUMPRODUCT supports layers of sophistication. See help (F1) for more examples.