Welcome to DMD Recitation 5! ☺

- Due electronically on Tuesday April 21, 2015
  - ALSTOM case – work individually, 2-page memo with appendix!
  - Exercises 7.2 (a-c), 7.7 (a-b) – complete individually!
  - Submit the PDF files on Stellar.
- To reduce background noise, please mute your phone/ computer!
- Please feel free to raise your hand or chat through WebEx if you have any questions or comments!

Outline

- Linear optimization
  - Model formulation
  - Geometric interpretation
  - Shadow price
  - Sensitivity analysis
- Solver tutorial
  - Setting up in Excel
  - Running Solver
  - Interpreting the output

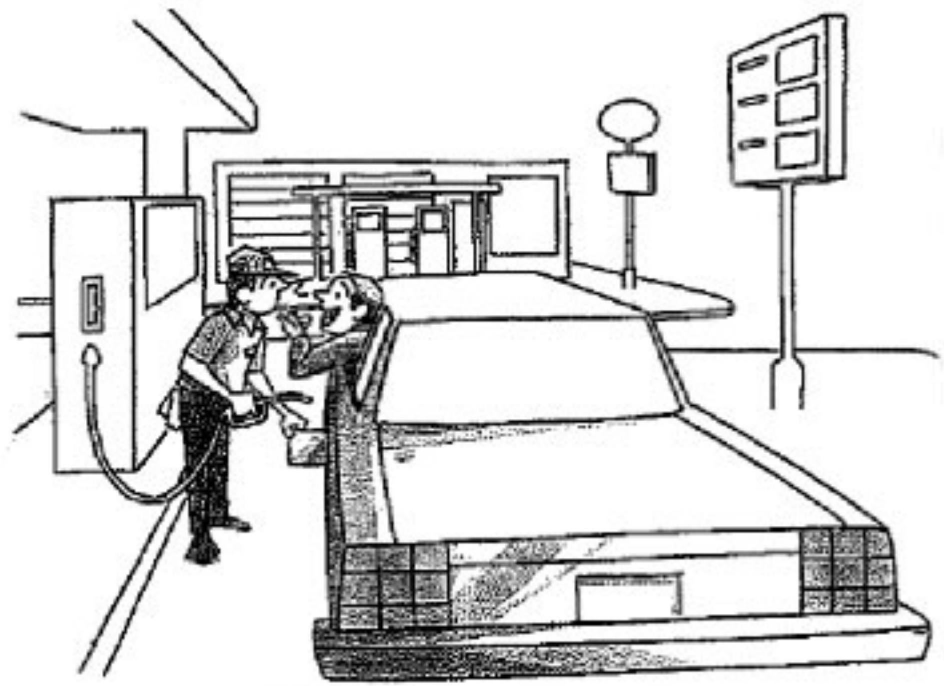

"Just enough to get me across the street to the cheaper station."

Setting up a model

1. **Define the decision variables**
   - What can you change? Ex: prices, inventory levels
2. **Write the objective function**
   - What are you trying to optimize (maximize/minimize)?
   - Written in terms of the decision variables
   - Ex: maximize profit = revenue – cost
3. **Write the constraints**
   - What is holding you back?
   - Written in terms of the decision variables
   - Ex: advertising budget is $1M, maximum machine output is 5 units/hr
4. **Restrict the decision variables**
   - Nonnegative? (coming soon: Integer? Binary?)

Linear optimization

What does "linear" mean?

- The objective function and all constraints are linear functions of the decision variables

- Example: 2 decision variables, x and y

	Linear	Nonlinear
Objective function	$x + y$ $\mathrm{sqrt}(2)y - \pi$	$5xy$ $x^2 + y$ $\mathrm{sqrt}(x)$
Constraints	$x + y \leq 10$ $x - 20y \geq 10x + 3$ $x = 10$	$xy \leq 23$

Example: Diet problem

- Imagine you want to set food policy
- You want to find a diet that
  - Satisfies the nutrient requirements of infants
  - Is as affordable as possible
- Assume you have 2 foods available: A and B

Diet data

- Each child needs daily
  - 615 calories
  - 10 grams protein
  - 20.5 grams fat
  - 10 mg iron

- The two foods provide the following amounts of nutrition:

	A	B
Calories (per 100g)	149	875
Protein (g per 100g)	4.8	0.1
Fat (g per 100g)	0.2	98.9
Iron (mg per 100g)	6.2	0.1
Cost (per 100g)	9¢	12¢

1. What is the problem we're trying to model and solve?

- What is the cheapest combination of the 2 foods that meets the daily nutritional needs?
- Alternative problem: what is the most nutritious diet that does not surpass a given budget?
  - What would be the objective function?
  - Need a single number to capture the "nutritiousness" of a diet

2. Define the decision variable

- Decision variables are the quantities that are under your control
  - You change them in order to optimize some objective

- What can we change?
  - Let A = amount of food A in the diet (measured in 100g)
  - Let B = amount of food B in the diet (measured in 100g)

- Don't forget to specify the units in the definition of decision variables
  - Make sure all variables and calculations use consistent units

3. Write the objective function

- The objective function defines the quantity you are trying to maximize or minimize
  - It is defined in terms of the decision variables
- What quantity are we trying to optimize?
  - Cost of the diet plan: 9A + 12B
- Are we maximizing or minimizing?
  - Minimize cost
- Therefore, our objective is **min 9A + 12B**

4. Identify and define the constraints

- Without constraints, optimization problems would be easy

- Constraints make the problem realistic and interesting

- Constraints are limitations on what values the decision variables can take on

4. Constraints: calories

- Each child needs 615 calories
- How do we represent this as a constraint?
- Our decision variables are A and B
  - A provides 149 cal/100g
  - B provides 875 cal/100g
- Total calories provided by the decision variables must be at least 615
  - Constraint: $149A + 875B \geq 615$
- Remember: keep track of the units of the decision variables

4. Constraints: protein, fat, iron

- Each child also needs
  - 10 grams protein
  - 20.5 grams fat
  - 10 mg iron

- Protein: $4.8A + 0.1B \geq 10$
- Fat: $0.2A + 98.9B \geq 20.5$
- Iron: $6.2A + 0.1B \geq 10$

5. Restrict the decision variables

- Are there any other restrictions on the decision variables?
- Can they be negative?
  - Probably not…
- Add nonnegativity constraints:
  - $A \geq 0$
  - $B \geq 0$
- Do they have to be integer valued?
  - Maybe… but we will discuss this later in the class

Full LP formulation

min	9A	+	12B		
s.t.	149A	+	875B	≥	615
	4.8A	+	0.1B	≥	10
	0.2A	+	98.9B	≥	20.5
	6.2A	+	0.1B	≥	10
	A			≥	0
			B	≥	0

Everything is linear with respect to the decision variables: A, B

Constraint LHS: linear function of the decision variables

Constraint RHS: constants

Geometry of diet problem

- With 2 decision variables, we can visualize the problem
  - x-axis: A
  - y-axis: B
- The x-y plane is the space of possible decisions (A,B)
- Ex: red dot is (2.4, 0.5)
  - A=2.4, B=0.5

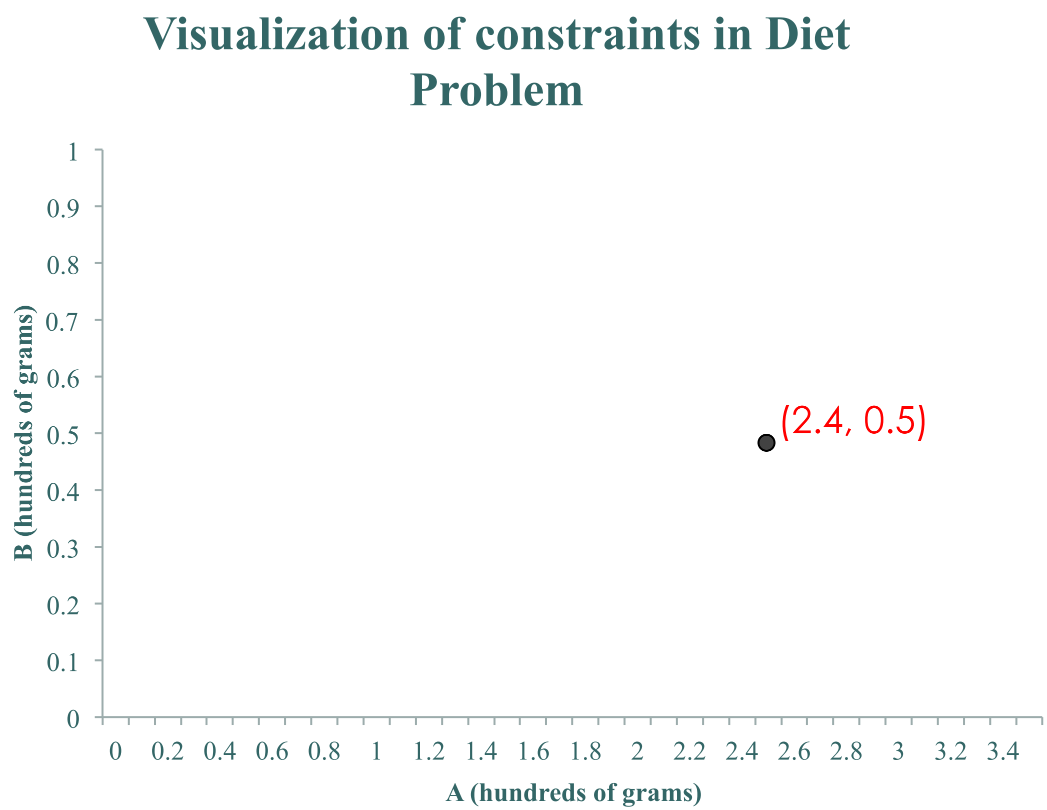

Geometry of diet problem

- Not all (A,B) pairs are feasible
- Constraints restrict the space of possible decisions
- Ex: calorie constraint
  - 149A + 875B ≥ 615
- Plot the line
  - 149A + 875B = 615
- Determine which side/half of the line is feasible
  - Easy test: plug in (0,0) and see whether the constraint is satisfied

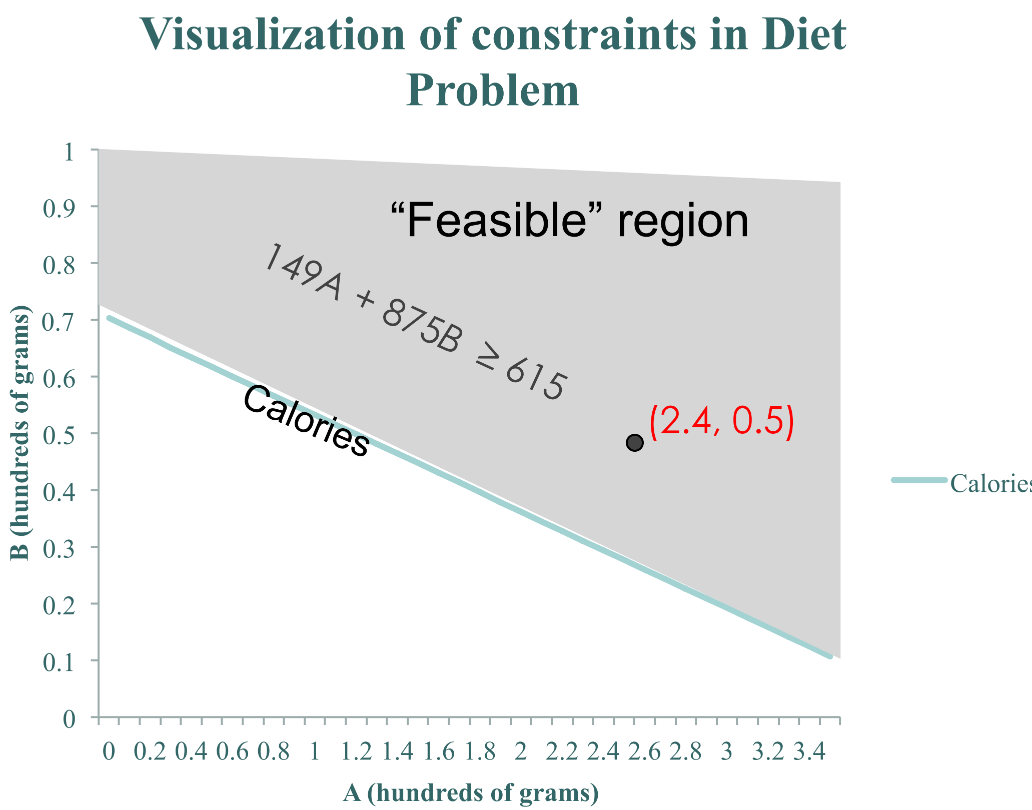

Visualization of constraints

- By graphing all additional constraints, we obtain the complete feasible region for the problem
- Any point in the feasible region will satisfy all of the constraints (i.e., meet the nutrition requirements). But which one is best?

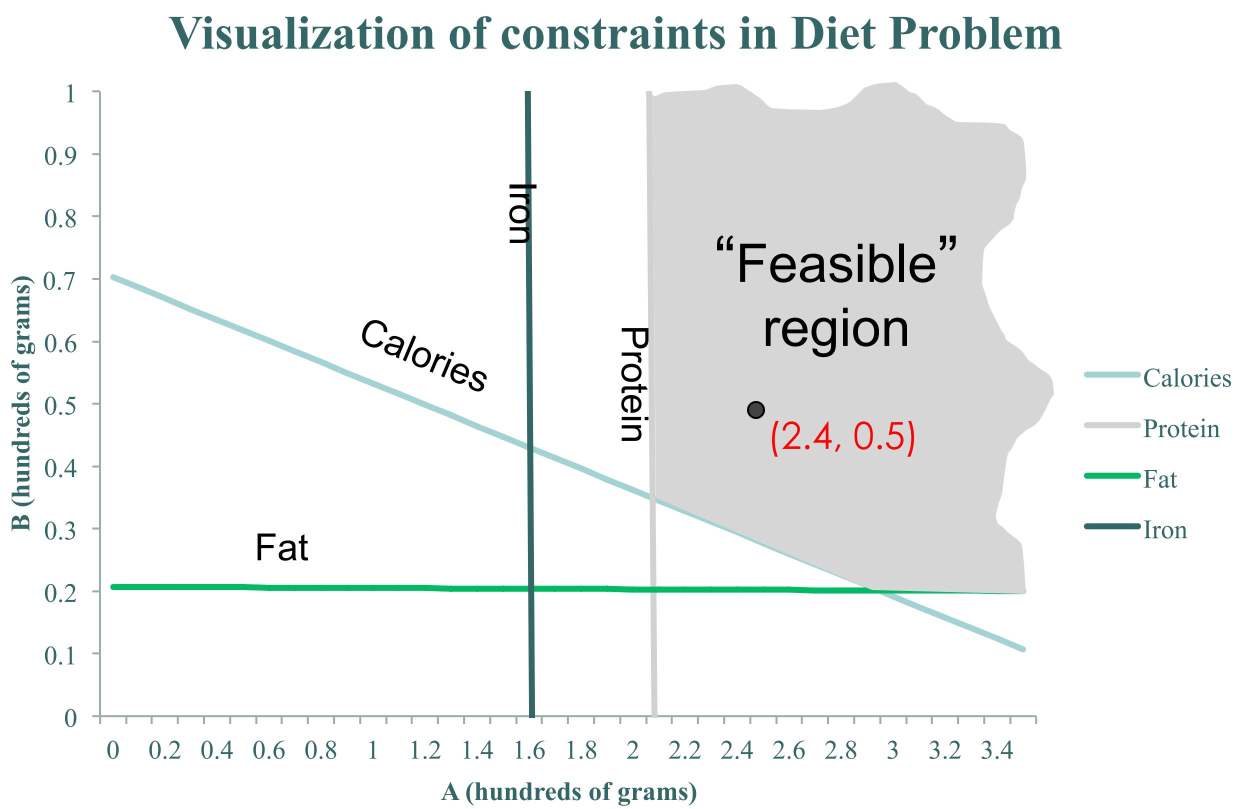

Incorporating the objective function

- Recall: the objective function is the cost of the diet: **min 9A + 12B**
  - This is not an equation – there is no right-hand-side
- Pick your favorite number, call it c
  - 9A + 12B = *c* is an equation of a line in the A-B plane
  - We can plot it, together with the feasible region
  - Every point on this line has the same cost *c*
  - Since we're minimizing, we like small c's
  - Points on the line 9A + 12B = -10000 are very desirable
  - But are they feasible?
- We still need to satisfy the constraints
  - Therefore, the line 9A + 12B = c must intersect the feasible region at at least 1 point
  - Points in the intersection are feasible and have a cost of c
  - We want to find the smallest c such that there is still an intersection

Plotting the objective function

- Example: choose c=27.6 => 9A + 12B = 27.6
- Intersects the feasible region
- (2.4, 0.5) is a feasible decision that achieves a cost of 27.6

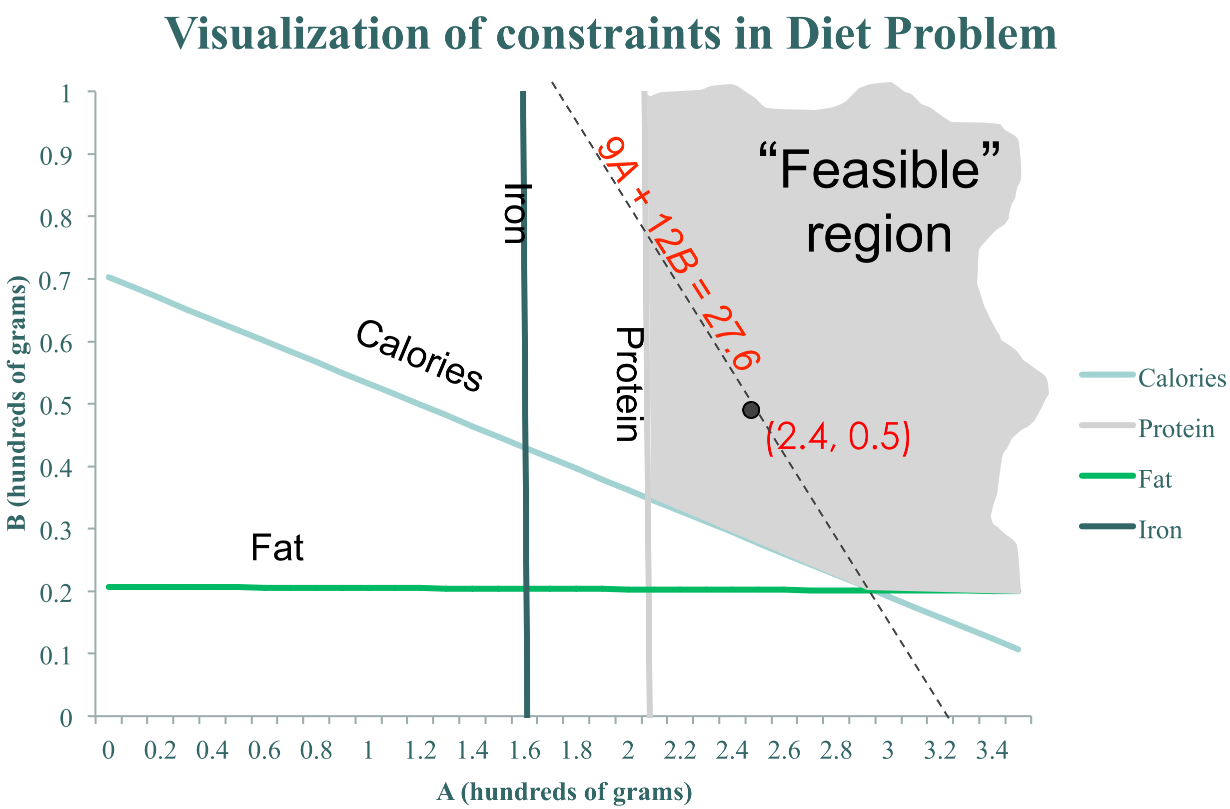

Plotting the objective function

- Example: choose c=18 => 9A + 12B = 18
- Does not intersect the feasible region
- Cost of 18 is not achievable given the constraints

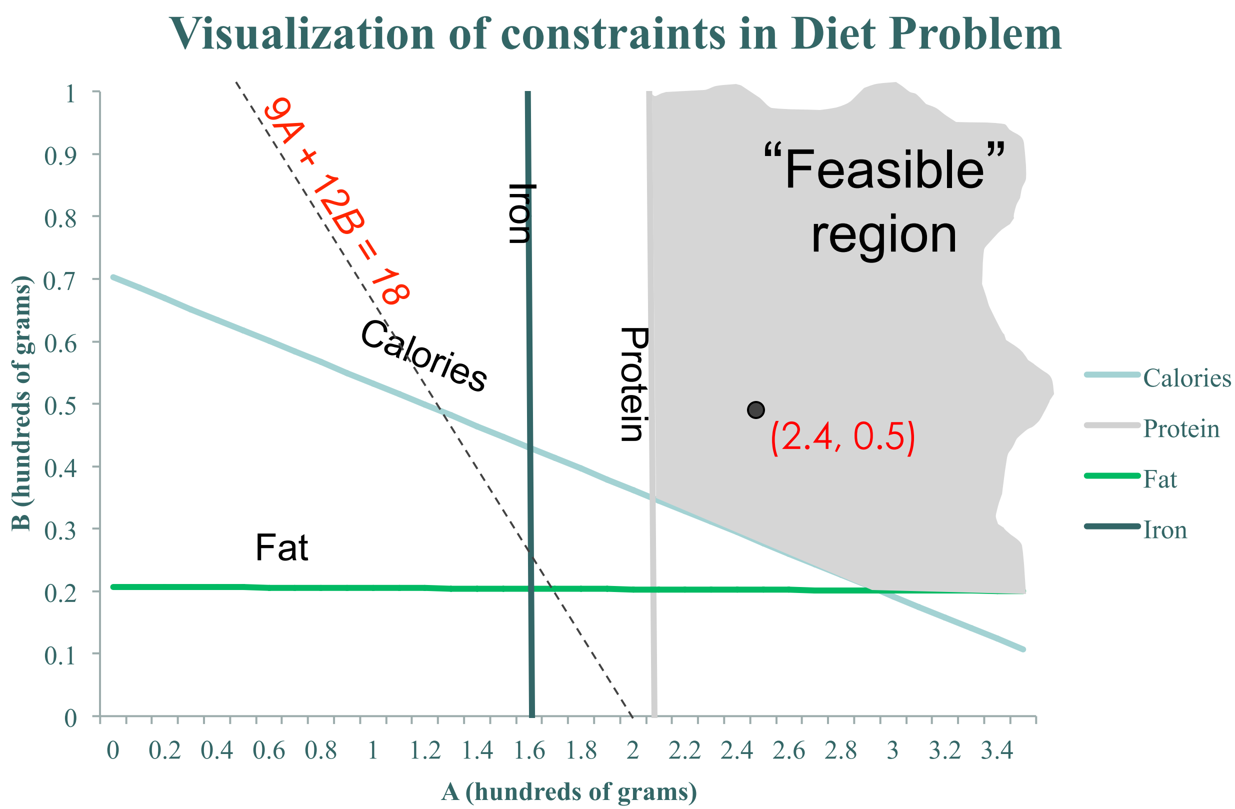

Graphical interpretation of optimizing the objective function

- Each choice of cost c gives a line
- Changing the target cost c is equivalent to sliding the line up and down
- We want to slide the line as far as possible while still maintaining feasibility

Finding the optimal solution (graphically)

- The point at the star is (2.08, 0.35)
- This is 208 grams of A and 35 grams of B
  - Remember the units
- The cost is 2.08*9 + 0.35*12 = 23 cents
- Can we do any better?
  - No, this is the optimal solution
  - If we slide the line any further, there would be no intersection

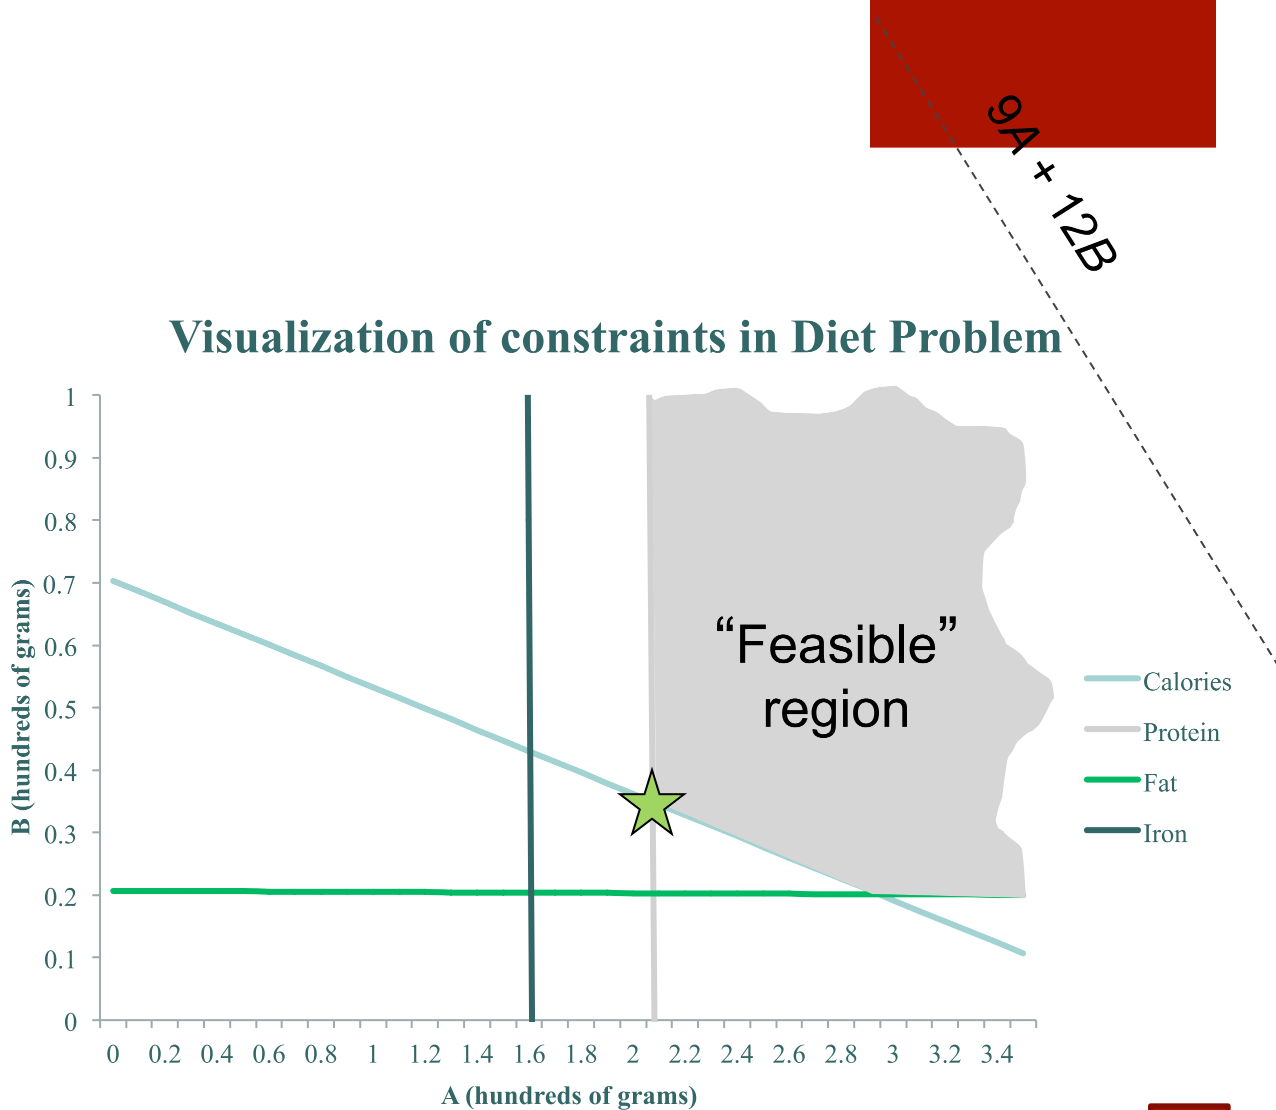

Are we done?

- Your analysis should not stop at finding an optimal solution.
- It's important to look at sensitivity analysis.
- Why?
  - Gives us more insight into the problem and the solution we found
  - Intuitively, constraints are "costly" because they reduce the space of feasible solutions; sensitivity analysis tells us precisely how costly each constraint is
  - Our data is not always exact – who knows if 10g is the right amount of protein?
  - With sensitivity analysis, we can answer questions like "how much more would it cost if we required 2g more protein?"

Analyzing the optimal solution

- Which constraints matter (i.e., are binding/active at the optimal solution)?
  - Protein
  - Calories
- Which constraints are not binding/ active?
  - Iron
  - Fat

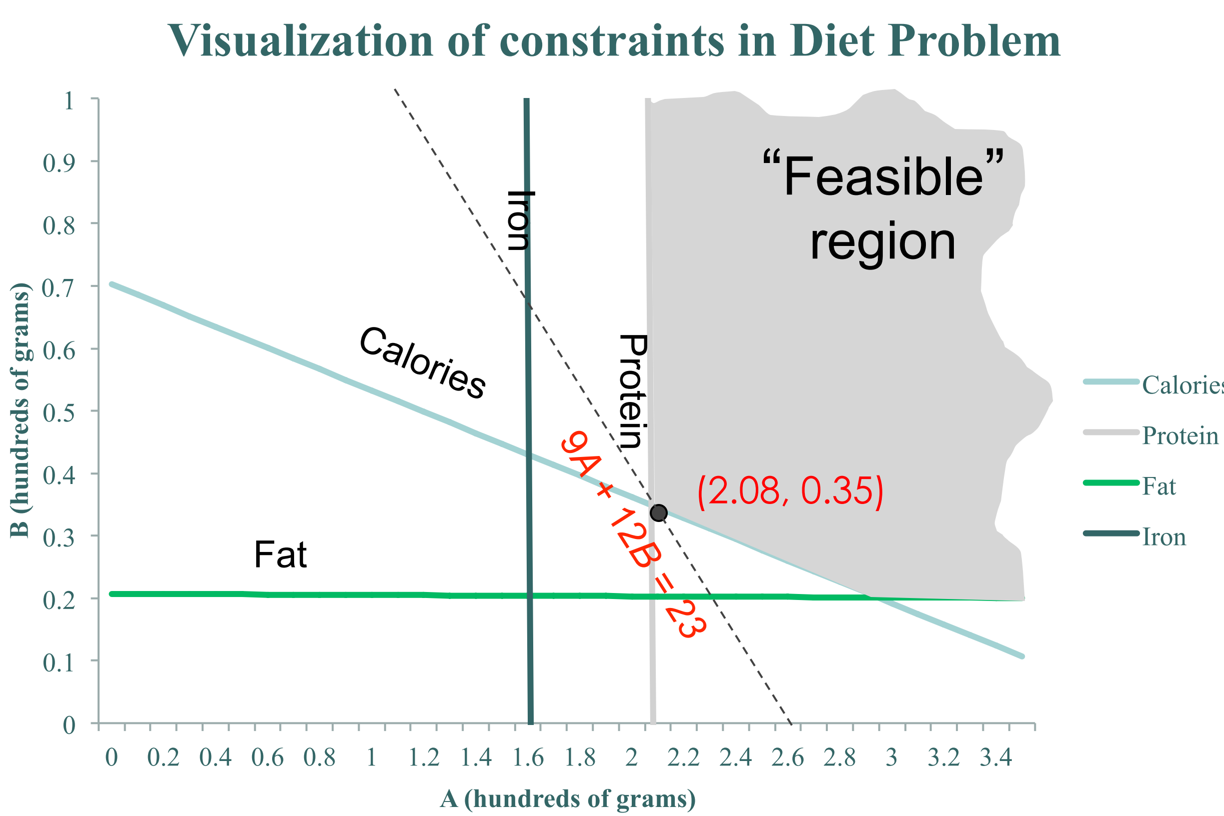

Sensitivity analysis

- Changing the constraints => changing the feasible region
- For nonbinding constraints, changes to the RHS (within a certain range) have no effect on the optimal solution
  - These constraints aren't the ones that are holding you back, so there is some leeway
- For binding constraints, even miniscule changes to the RHS can change the optimal solution
  - The optimal solution just barely satisfies these constraints, so if they are changed even slightly, the original optimal solution may no longer be feasible

Shadow prices

- Question: how much more will it cost if we require 1 more gram of protein?
  - Solution 1: update the model and resolve
  - Solution 2: use shadow price information
- Question: what is the change in the optimal objective value when the RHS of a constraint is increased by 1 unit?
  - This quantity is called the constraint's **"shadow price"**.
- If the shadow price of the protein constraint is 1.45 (cents), we immediately know that an optimal diet requiring 1 more gram of protein would cost 1.45 more cents
- Shadow prices are associated with a particular optimal solution.
- **Shadow prices are valid only when changes in the RHS are within a given range**
  - If you are interested in changes outside of this range, you will have to use Solution 1: update the model and resolve

Sensitivity analysis

- **Non-binding**
  - Changes won't incur any change within a given range
  - No effect of the constraint on the objective value
  - Shadow price = ZERO
- **Binding**
  - Changes will incur changes!
  - Effect of the constraint on the objective value dictated by shadow prices.

Variable Cells

Cell	Name	Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
E29	How much to get of each of A	2.07605553	0	9	567	6.956571429
F29	How much to get of each of B	0.349334544	0	12	40.85234899	11.8125

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
S5	Calories	615	0.013548063	615	86885	127.5225406
S6	Protein (g)	10	1.454445533	10	4.15730611	2.248368792
S7	Fat (g)	34.96439751	0	20.5	14.46439751	1E+30
S8	Iron (mg)	12.90647774	0	10	2.906477742	1E+30

Changing constraints

- Moving the fat and iron constraint lines have no effect (up to a certain point)
- But moving the protein constraint to the right changes the feasible region
  - The optimal solution also changes

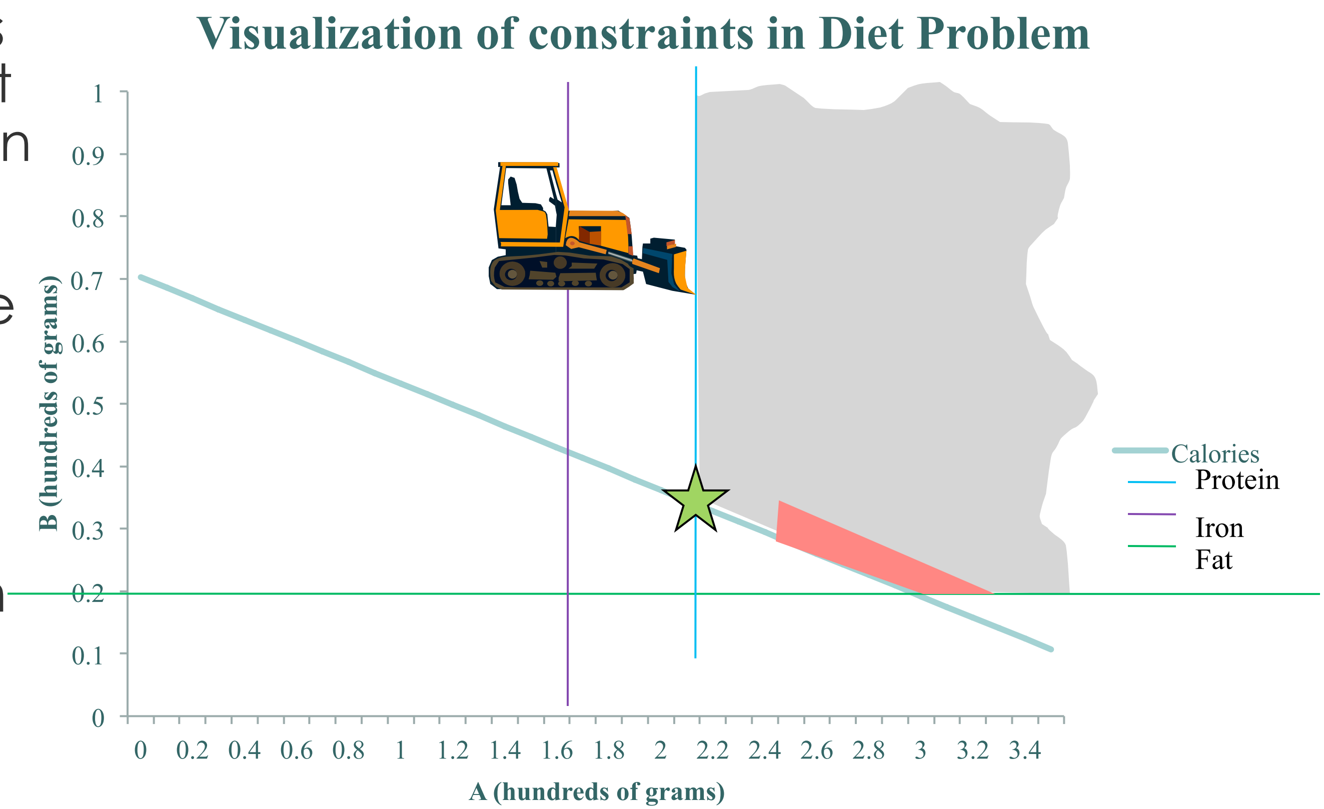

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
S5	Calories	615	0.013548063	615	86885	127.5225406
S6	Protein (g)	10	1.454445533	10	4.15730611	2.248368792
S7	Fat (g)	34.96439751	0	20.5	14.46439751	1E+30
S8	Iron (mg)	12.90647774	0	10	2.906477742	1E+30

Nonbinding constraint

- The iron constraint is nonbinding.
- Its shadow price is 0.
- $6.2A + 0.1B \geq 10$
- We can increase the RHS from 10 to 12.9 and the solution will not change
- … or decrease to -∞

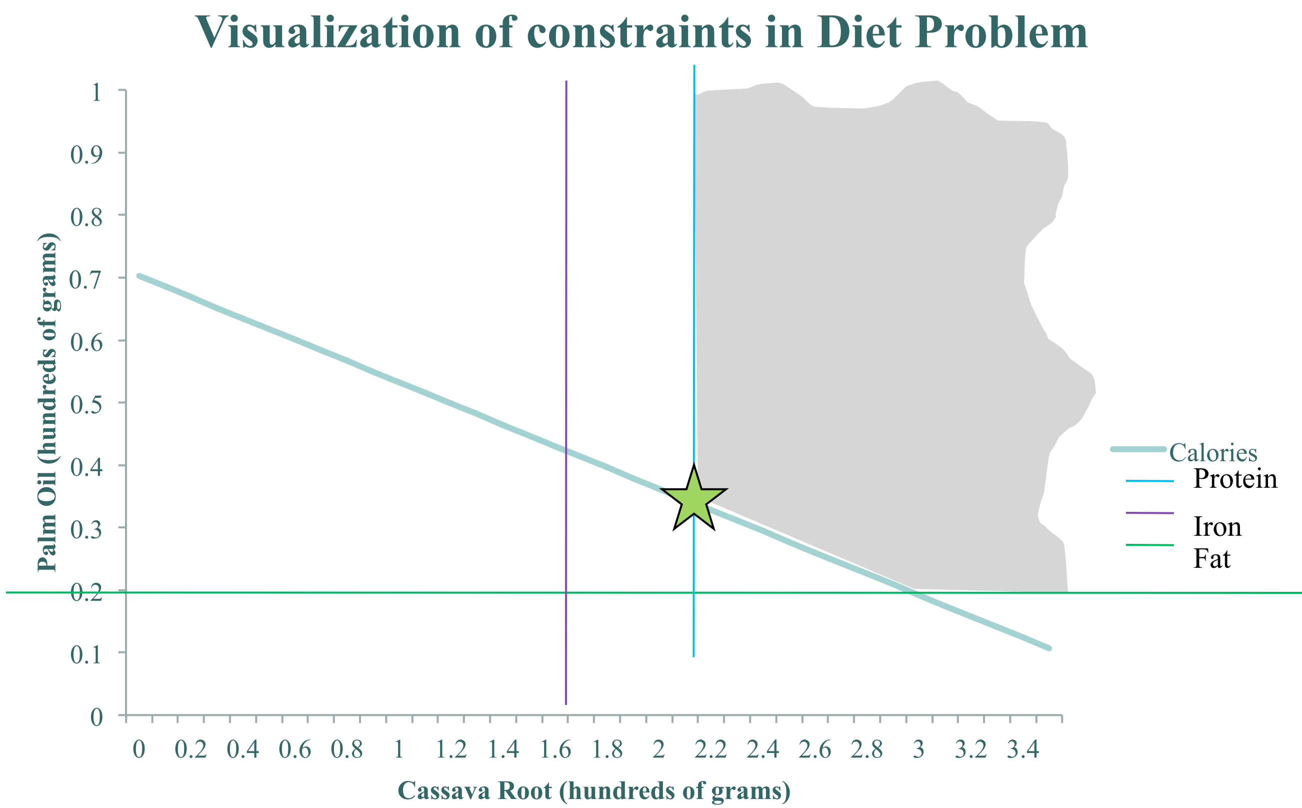

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
S5	Calories	615	0.013548063	615	86885	127.5225406
S6	Protein (g)	10	1.454445533	10	4.15730611	2.248368792
S7	Fat (g)	34.96439751	0	20.5	14.46439751	1E+30
S8	Iron (mg)	12.90647774	0	10	2.906477742	1E+30

Binding constraint

- The protein constraint is binding
- Its shadow price is 1.45 cents
  - Allowable increase: 4.16
  - Allowable decrease: 2.25
- $4.8A + 0.1B \geq 10$
- If we change the RHS to 12, the optimal cost would increase by 2*1.45 = 2.9 cents

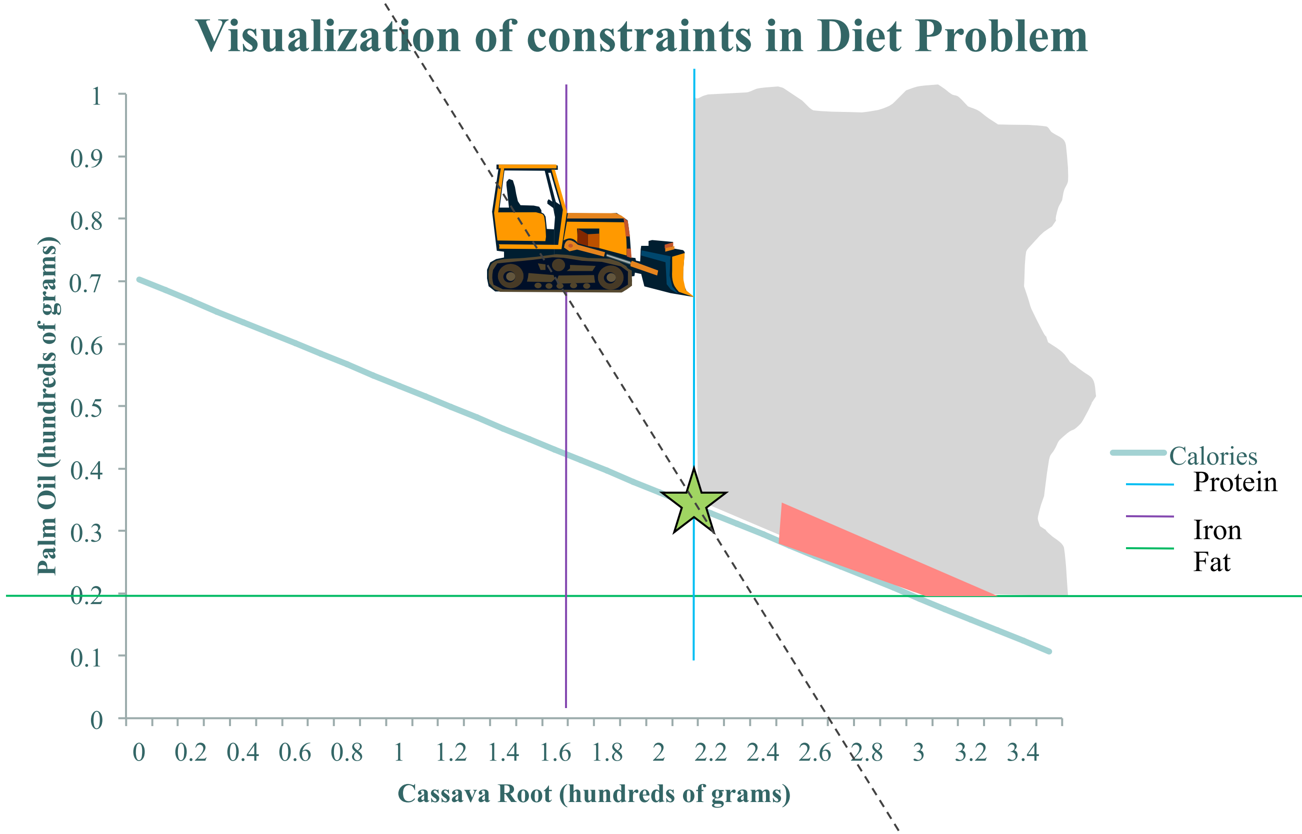

Sensitivity questions

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
S5	Calories	615	0.013548063	615	86885	127.5225406
S6	Protein (g)	10	1.454445533	10	4.15730611	2.248368792
S7	Fat (g)	34.96439751	0	20.5	14.46439751	1E+30
S8	Iron (mg)	12.90647774	0	10	2.906477742	1E+30

- How much more will it cost to feed an infant if the fat requirement is raised to 30 grams?
  - Shadow price is 0, change is within allowable increase, so it will cost nothing additional
- How much less will it cost to feed an infant if the caloric requirement is reduced by 100 calories?
  - Shadow price is 0.014, change is within allowable decrease, so 100*0.014=1.4 fewer cents
- How much more will it cost to feed an infant if the iron requirement is increased by 5 mg?
  - Can't tell immediately: 5 is outside the allowable increase of 2.9, so we would have to modify the constraint and rerun Solver

Higher dimensions

- What if there was a third food C?
  - Add another decision variable
  - New data: nutritional content and cost of C
  - Modify the constraints to include nutritional contributions from C
  - Modify the objective function to include cost of C
- Visually,
  - There would be 3 axes and we would plot in 3 dimensions
  - Constraints would be "hyperplanes" instead of lines
  - Feasible region would be a 3-dimensional "polyhedron"
- Could add 4^{th}, 5^{th}, …

Another example

Resource	Product 1	Product 2	Product 3	Product 4
Raw Material	2	3	4	7
Hours of Labor	3	4	5	6
Sales Price (tens of thousands of $)	4	6	7	8

- Maximizing sales revenue for a manufacturer
- 4,600 units of raw material and 5,000 hours of labors are available.
- Exactly 950 total units must be produced. Also, at least 400 units of product 4 must be produced.

Another example

Variable Cells

Cell	Name	Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
D15	Product 1 (units)	0	-1	4	1	1E+30
D16	Product 2 (units)	400	0	6	0.666666667	0.5
D17	Product 3 (units)	150	0	7	1	0.5
D18	Product 4 (units)	400	0	8	2	1E+30

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
D26	raw materials	4600	1	4600	250	150
D27	labor hours	4750	0	5000	1E+30	250
D28	must produce	950	3	950	50	100
D29	P4, at least	400	-2	400	37.5	125

Optimal solution:
Max revenue: $66.5M
$X_1 = 0$
$X_2 = 400$
$X_3 = 150$
$X_4 = 400$

- Suppose that a total of 980 units must be produced. What is the new optimal solution (including the new total sales revenue)?
  - Extra 30 units are still within the allowable range (30 < 50)
  - (Increment) x (shadow price) = 30 x ($30000) = $900000.
  - The new revenue is then $66.5M + $0.9M = $67.4M
  - (The solution will be changed too: (0, 520, 60, 400)).

Another example

Variable Cells

Cell	Name	Final Value	Reduced Cost	Objective Coefficient	Allowable Increase	Allowable Decrease
D15	Product 1 (units)	0	-1	4	1	1E+30
D16	Product 2 (units)	400	0	6	0.666666667	0.5
D17	Product 3 (units)	150	0	7	1	0.5
D18	Product 4 (units)	400	0	8	2	1E+30

Constraints

Cell	Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
D26	raw materials	4600	1	4600	250	150
D27	labor hours	4750	0	5000	1E+30	250
D28	must produce	950	3	950	50	100
D29	P4, at least	400	-2	400	37.5	125

Optimal solution:
Max revenue: $66.5M
$X_1 = 0$
$X_2 = 400$
$X_3 = 150$
$X_4 = 400$

- Suppose that 4500 units of raw material are available. What is the new optimal solution to the LP (including the new total sales revenue)?
  - 4500: (Increment) x (shadow price) = -100 x ($10000) = -$1M.
  - The new revenue is then $66.5M - $1M = $65.5M
  - What if only 4400 units are available?

Wrap up! ☺

- Tonight, we covered linear optimization and Excel Solver.
- Due electronically on Tuesday April 21, 2015
  - ALSTOM case – work individually, 2-page memo with appendix!
  - Exercises 7.2 (a-c), 7.7 (a-b) – complete individually!
  - Submit the PDF files on Stellar.
- Please email us any questions or comments/suggestions for improving the recitation.
- Office Hours now!