15.053/8 February 26, 2013

Sensitivity analysis and shadow prices

special thanks to Ella,
Cathy, McGraph, Nooz,
Stan and Tom

Quotes of the Day

“If the facts don't fit the theory, change the facts.”

-- Albert Einstein (attributed)

“What's the use of a good quotation if you can't change it?”

-- Doctor Who

“What’s the use of a good quotation if you can’t change it, and then claim credit for it?”

-- Professor Orlin

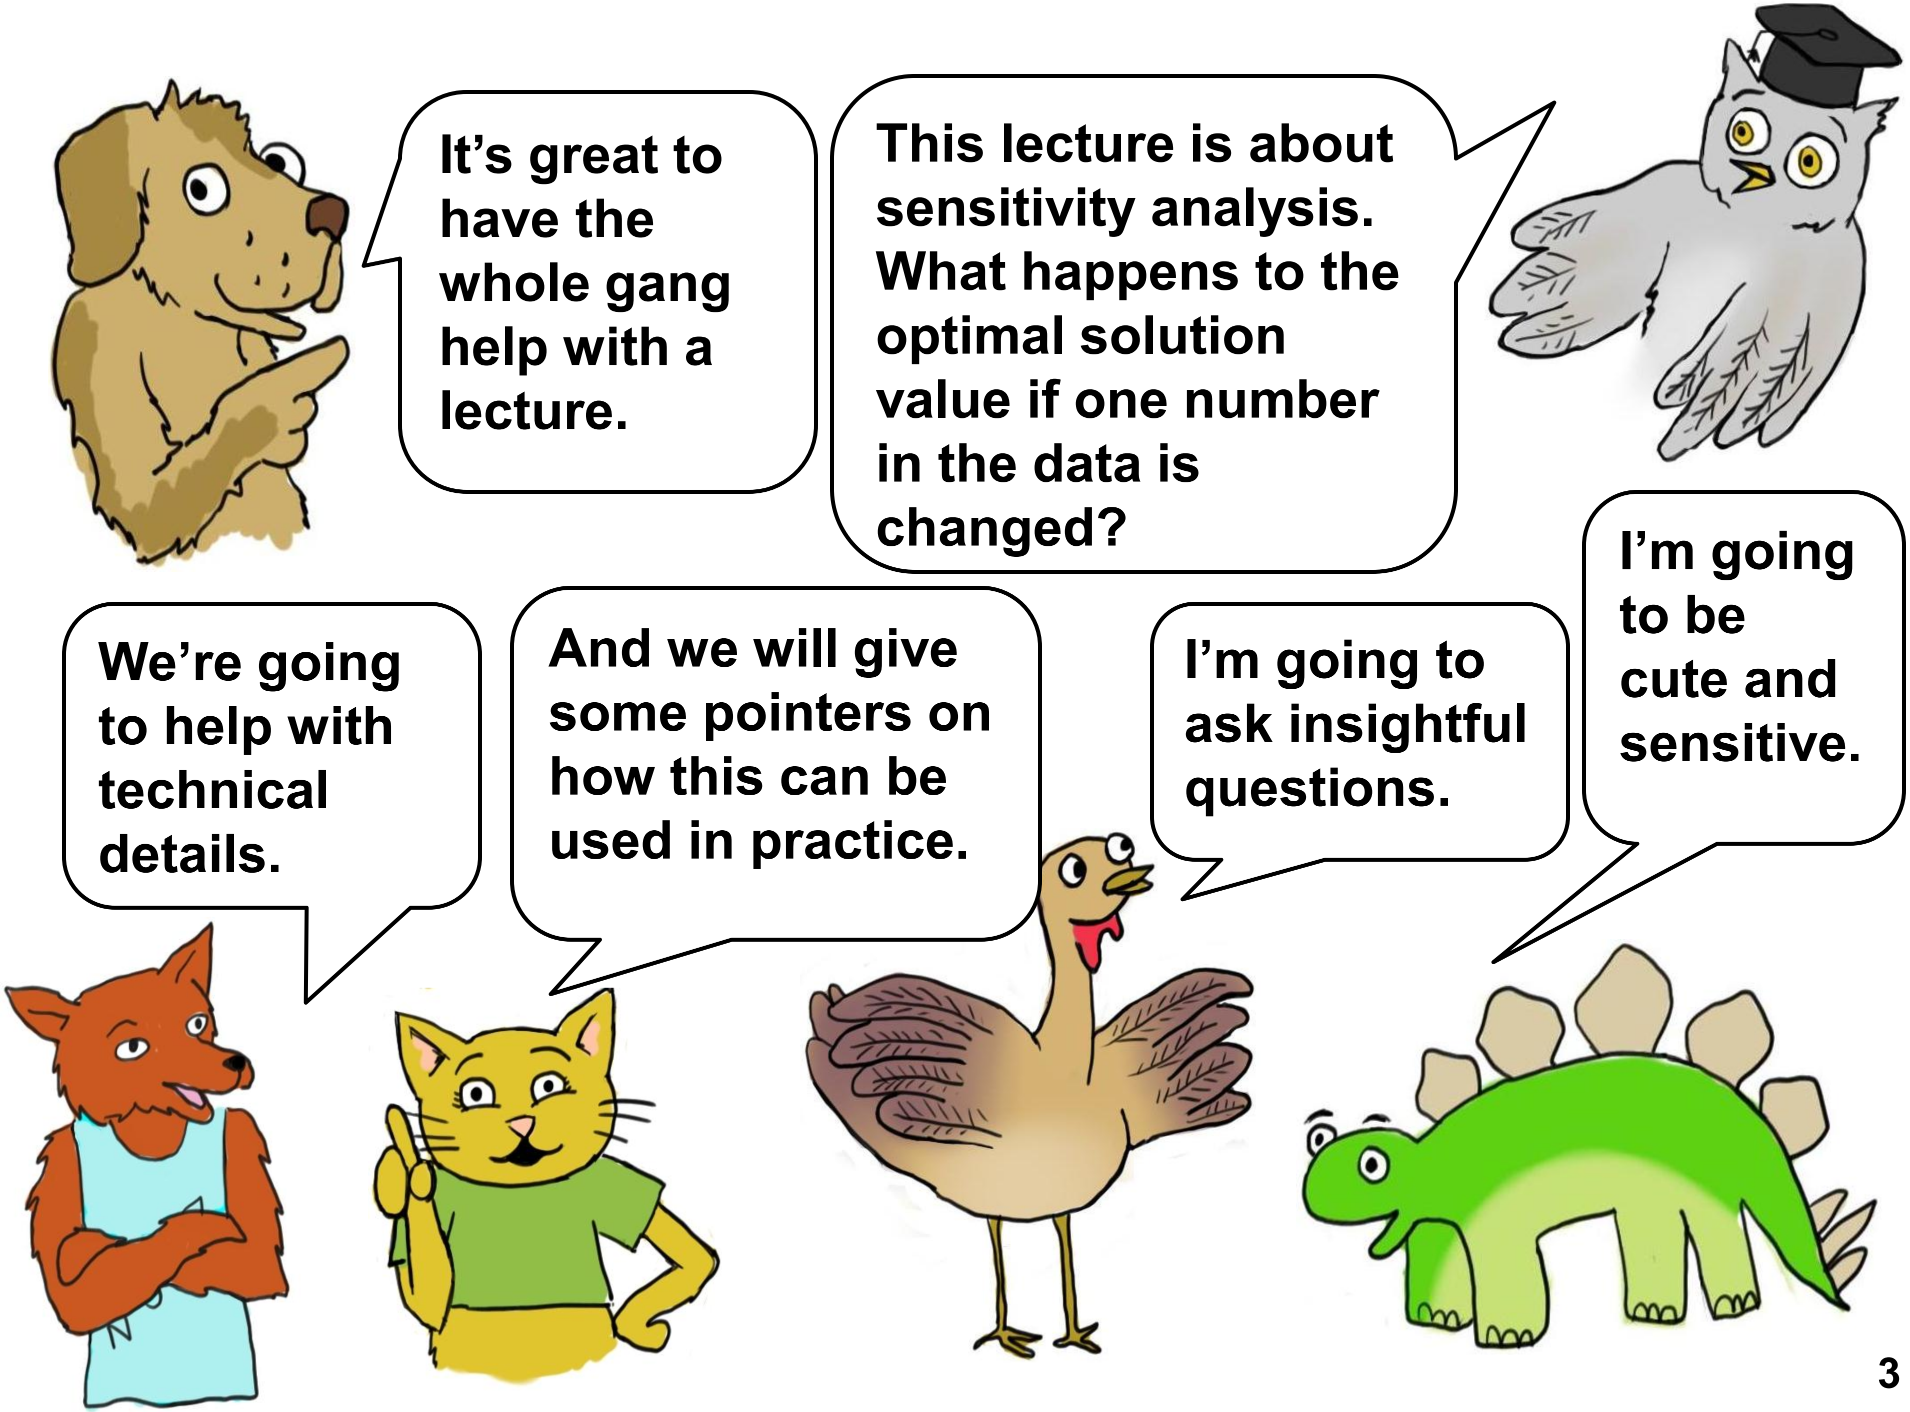
It's great to have the whole gang help with a lecture.
This lecture is about sensitivity analysis. What happens to the optimal solution value if one number in the data is changed?
We're going to help with technical details.
And we will give some pointers on how this can be used in practice.
I'm going to ask insightful questions.
I'm going to be cute and sensitive.

MIT Computer Corporation (mc^2)

motto: transforming Mass., energizing the world

The following is a fictional case. It is based on the DEC case, developed by Rob Freund and Brian Shannahan in 1988.

It demonstrates the use of linear programming for planning the "ramping" of new computer hardware.

Background

MIT Computer Corp (mc^2) announced a new family of tablet computers and e-readers in the second quarter of 2010. Shipments began in the 3rd quarter of 2010. The tablet computers and e-readers had the following code names:

- **Aardvark.** A high-end, general purpose tablet computer, with touch screen and with large memory disk space.
- **Bison.** A medium-end, general purpose tablet computer with touch screen
- **Cougar.** A general purpose tablet computer requiring a tablet pen
- **Deer.** A high-end e-reader with many additional functionalities
- **Emu.** An e-reader.

The Aardvark required newly developed high speed memory, which was in limited supply.

All of the Bisons, half of the Deers, and 20% of the Aardvarks required a new type of disk drive, which was in limited supply.

	A	B	C	D	E	Amount (in 1000s)
High Cap Mem. Chips	2	0	0	0	0	40
Low Cap Mem. Chips	0	2	2	2	1	240
Avg. # of new disk drives	.2	1	0	.5	0	20
List Price (in $1000s)	1.2	.8	.6	.6	.3	

Item	Demand (in 1000s)
A	18
C	3
Tablets	38
e-readers	32

A Number of Aardvarks manufactured (in 1000s)

B Number of Bisons manufactured (in 1000s)

C, D, E …

$$
\begin{array}{llllllll}
\text{Max} & 1.2A + .8B + .6C + .6D + .3E \\
\text{s.t} & 2A & & & & & \le & 40 \\
& & 2B + & 2C + & 2D + & E & \le & 240 \\
& .2A + & B + & & .5D & & \le & 20 \\
& A & & & & & \le & 18 \\
& & & C & & & \le & 3 \\
& A + & B + & C & & & \le & 38 \\
& & & & D + & E & \le & 32 \\
& & & A, B, C, D, E \ge 0
\end{array}
$$

Original Spreadsheet: the solution

	A	B	C	D	E
Decision Variables	**18**	**16.4**	**3**	**0**	**32**

Profit $46.12 in $ millions

Constraints				
High Cap Memory	36	≤	40	
Low Cap Memory	70.8	≤	240	
New Drives	20	≤	20	binding
Max for A	18	≤	18	binding
Max for C	3	≤	3	binding
Max for tablets	37.4	≤	38	
Max for e-readers	32	≤	32	binding

General rule: if there are small changes in the data, the optimal set of basic variables does not change.

A	B	C	D	E	S1	S2	S3	S4	S5	S6	S7	RHS
1.2	0.8	0.6	0.6	0.3								0
2	0	0	0	0	1							40
0	2	2	2	1		1						240
0.2	1	0	0.5	0			1					20
1	0	0	0	0				1				18
0	0	1	0	0					1			3
1	1	1	0	0						1		38
0	0	0	1	1							1	32

First tableau

A	B	C	D	E	S1	S2	S3	S4	S5	S6	S7	RHS
			-0.1				-0.8	-1.04	-0.6		-0.3	-46.12
1			0				0	1	0		0	18
	1		0.5				1	-0.2	0		0	16.4
		1	0				0	0	1		0	3
			1	1			0	0	0		1	32
			0		1		0	-2	0		0	4
			0			1	-2	0.4	-2		-1	169.2
			-0.5				-1	-0.8	-1	1	0	0.6

Optimal and final tableau

The structure of the solution stays the same with small changes of data

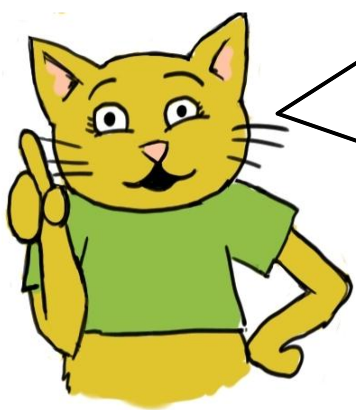

The basic feasible solution is obtained by solving linear systems of equations. With small changes in data, we solve almost the same system.

- If the RHS changes, the solution and the optimal objective change linearly.
- If the cost coefficients change, the optimal solution stays the same.
- The sensitivity report puts lots of this information in a useful format.

Sensitivity Report (SR) Part 1

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
A	18	0	1.2	1.00E+30	1.04
B	16.4	0	0.8	5.2	0.2
C	3	0	0.6	1.00E+30	0.6
D	0	-0.1	0.6	0.1	1.00E+30
E	32	0	0.3	1.00E+30	0.1

Sensitivity Report (SR) Part 2

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
High Cap Memory	36	0	40	1.00E+30	4
Low Cap Memory	70.8	0	240	1.00E+30	169.2
New Drives	20	0.8	20	0.6	16.4
Max for A	18	1.04	18	0.75	18
Max for C	3	0.6	3	0.6	3
Max for tablets	37.4	0	38	1.00E+30	0.6
Max for e-readers	32	0.3	32	169.2	32

What is the optimal solution?

The optimal solution is in the original spreadsheet.

	A	B	C	D	E
Decision Variables	**18**	**16.4**	**3**	**0**	**32**
Profit	$46.12	in $ millions			

It is also in SR part 1, in the column labeled "final value."
Note: the SR report does not have the profit.

Troubleshooting tip #1

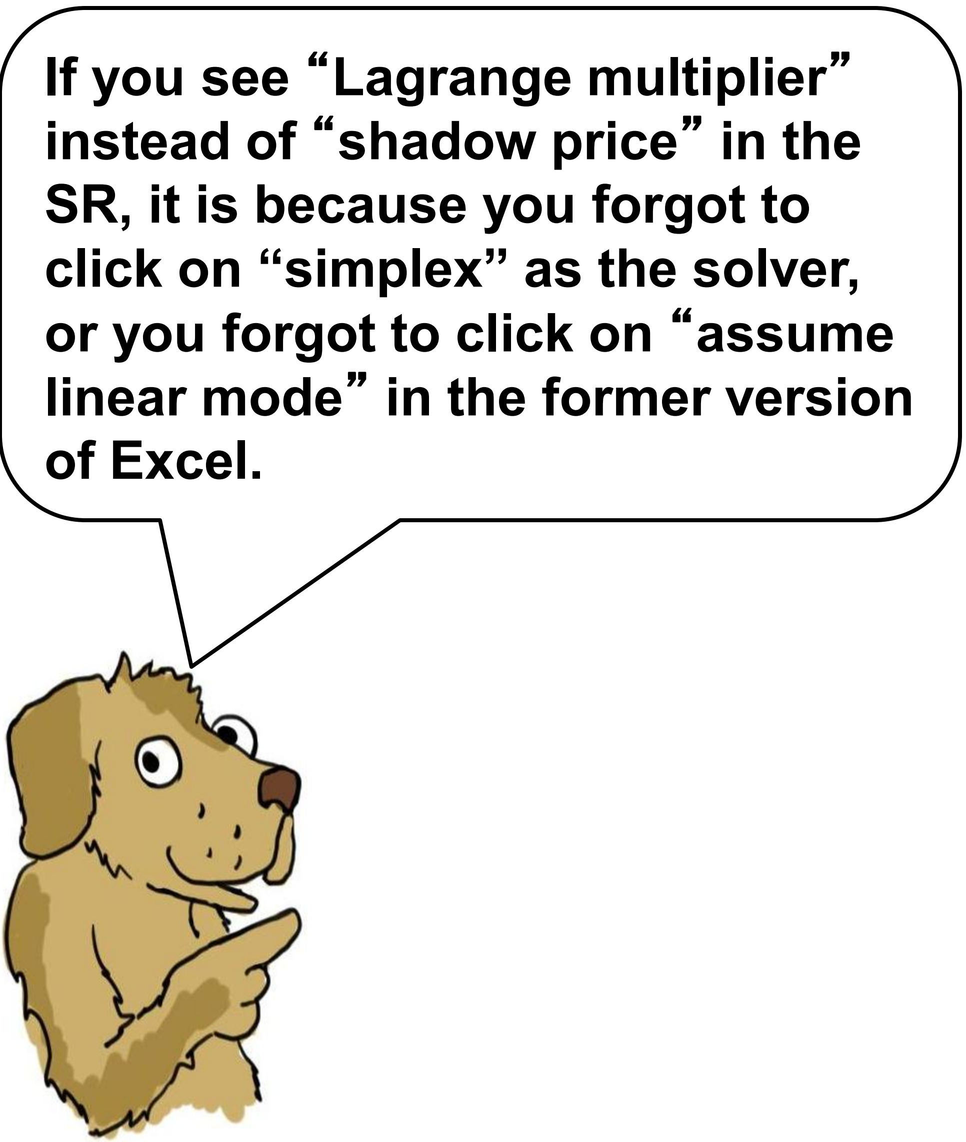

Shadow Price	Lagrange Multiplier
0	0
0	0
0.8	0.8
1.04	1.04
0.6	0.6
0	0
0.3	0.3

Changes that we will consider

1. Change the cost coefficient of a variable
2. Change the RHS of a constraint
   - Changing the initial conditions
   - Purchasing resources at a cost
3. Introducing a new product
   - Reduced costs

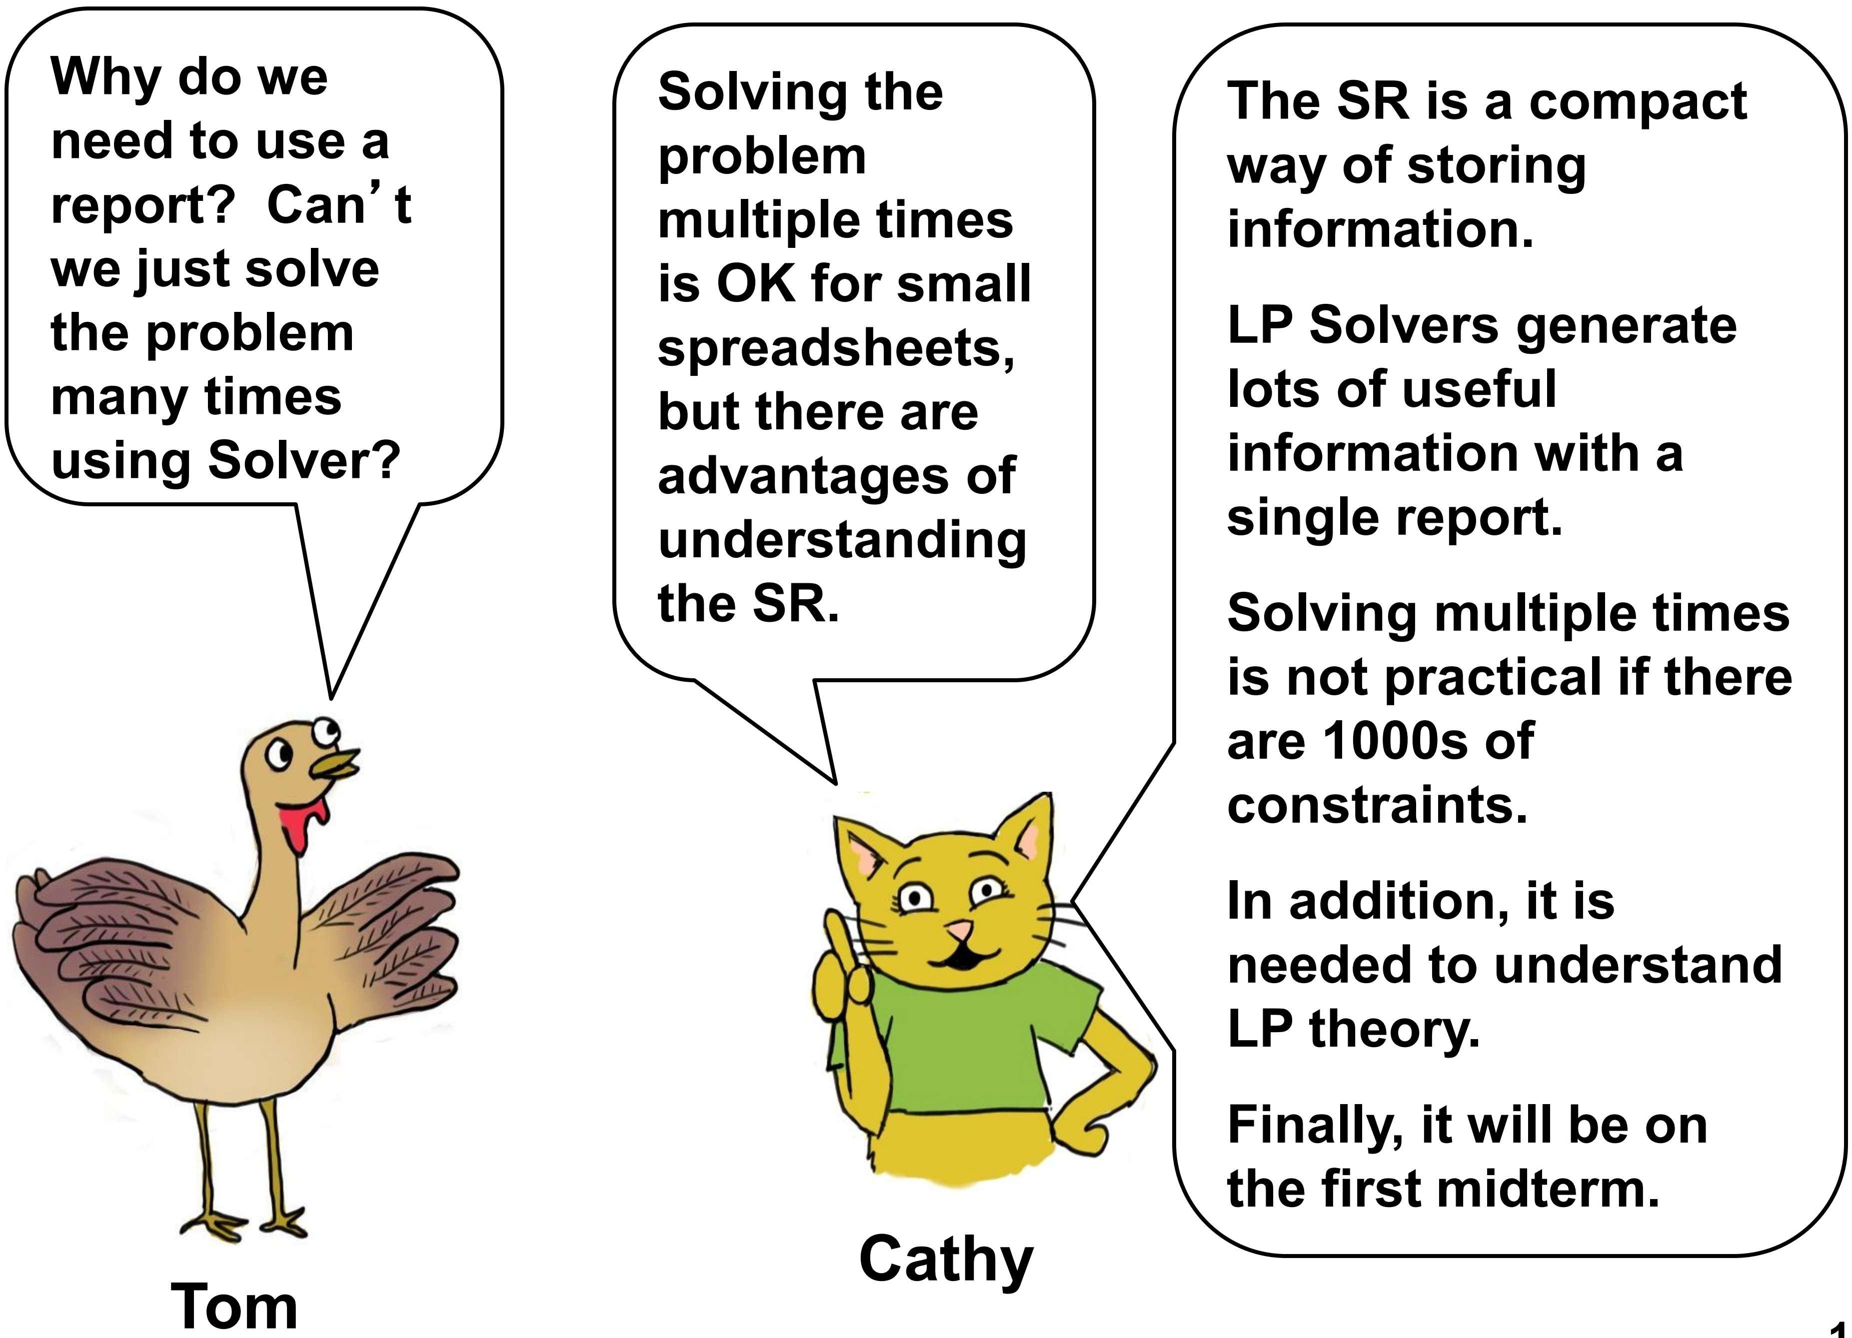
Why do we need to use a report? Can’t we just solve the problem many times using Solver?
Tom
Solving the problem multiple times is OK for small spreadsheets, but there are advantages of understanding the SR.
The SR is a compact way of storing information.
LP Solvers generate lots of useful information with a single report.
Solving multiple times is not practical if there are 1000s of constraints.
In addition, it is needed to understand LP theory.
Finally, it will be on the first midterm.
Cathy

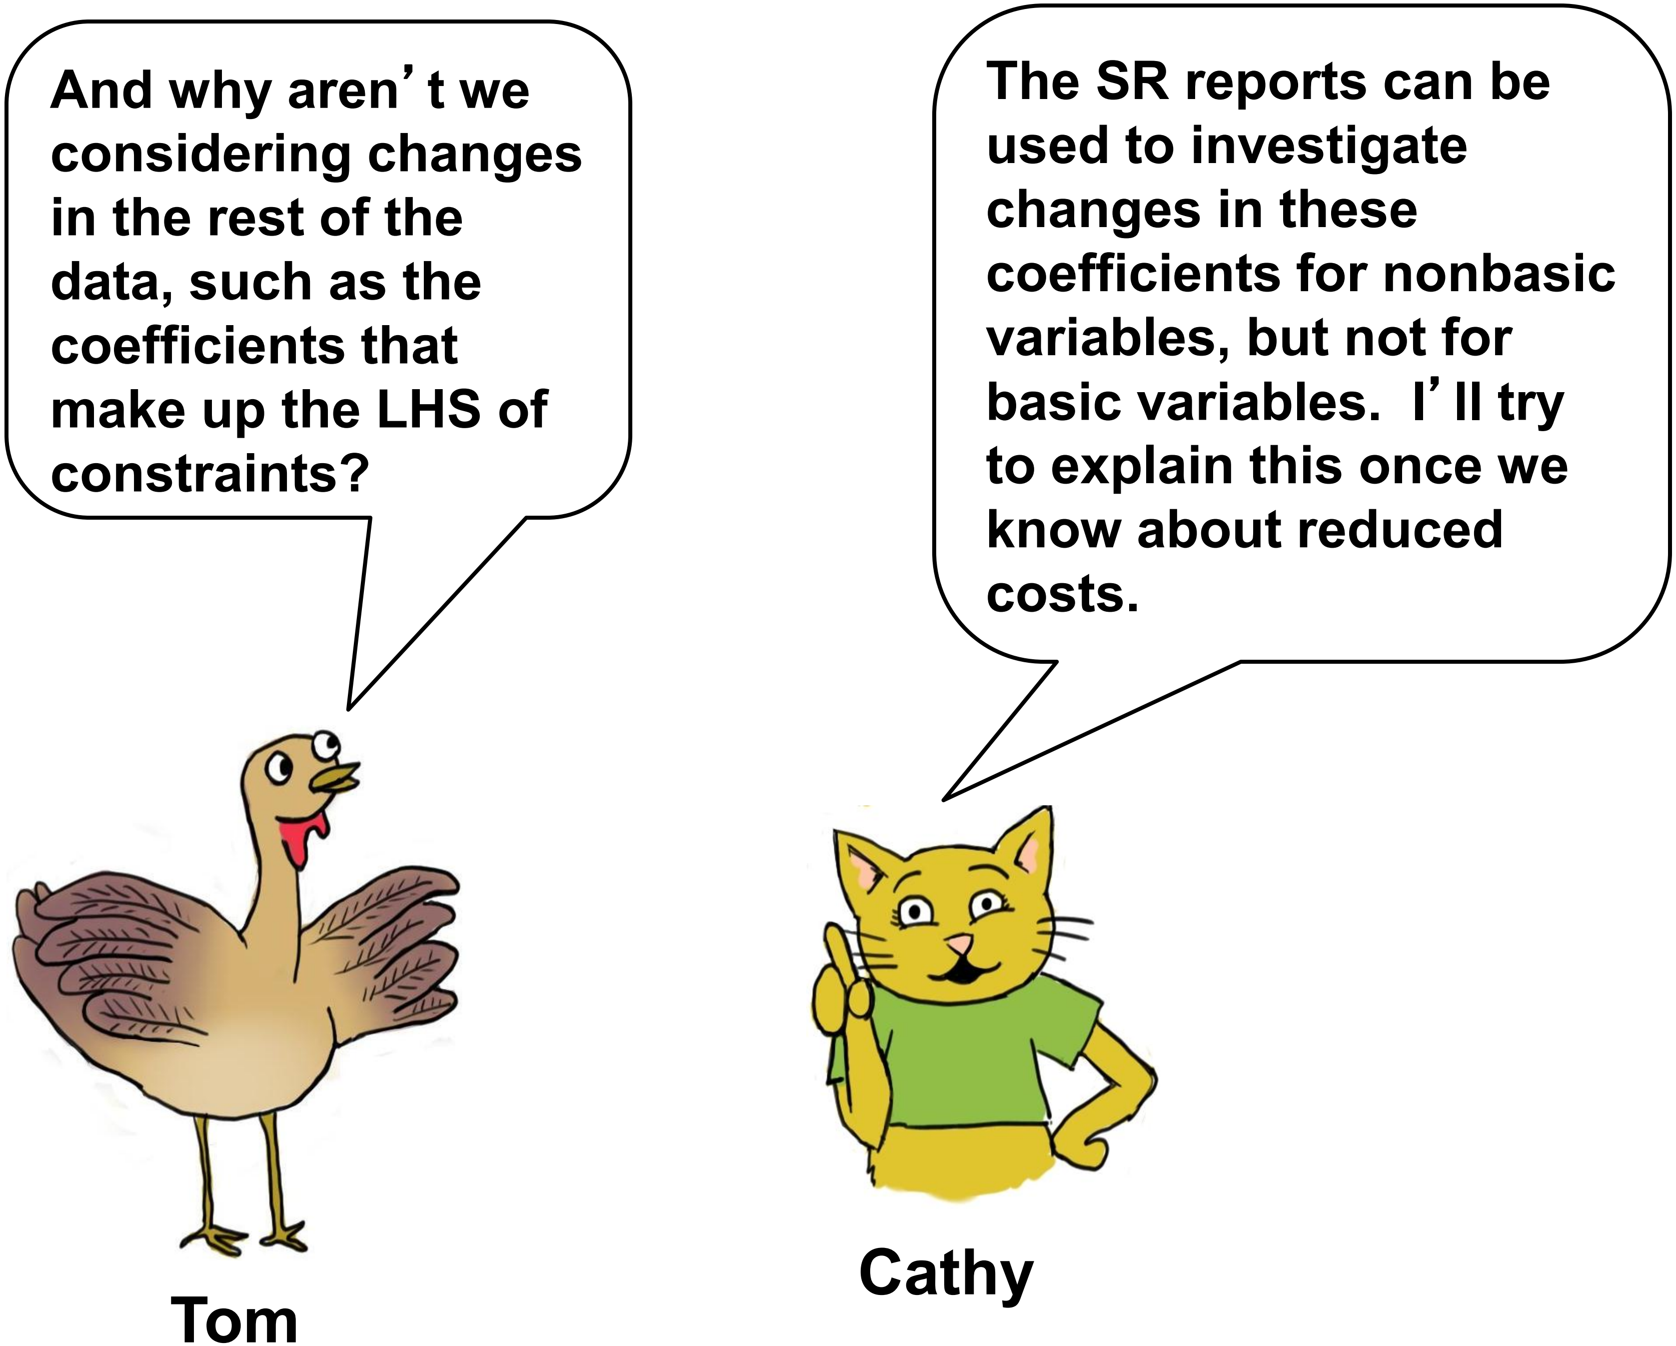
And why aren’t we considering changes in the rest of the data, such as the coefficients that make up the LHS of constraints?
The SR reports can be used to investigate changes in these coefficients for nonbasic variables, but not for basic variables. I’ll try to explain this once we know about reduced costs.
Tom
Cathy

Changing the cost coefficient of a basic variable

mc^2 is uncertain as to whether the Aardvark is priced to high. They are considering lowering the price from \$1200 to \$1000. What will be the impact on the total revenue?

In practice, lowering the price should result in an increase in demand. But here we assume demand is unchanged. In this sensitivity analysis, we change only one number in the data at a time, and assume all other data is unchanged.

The analysis

- For very small changes in the cost coefficients, the optimal solution is unchanged.
- Check the allowable increase and decrease of the cost coefficient to see if the solution changes.
- If the optimal solution is unchanged, then you can compute the new objective value.

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
A	18	0	1.2	1.00E+30	1.04

QUESTION FOR STUDENTS

Suppose that the list price of the Aardvark is changed from $1,200 to $11,200. What is the best answer below?

1. **The sensitivity report says that the optimal solution will not change.**
2. **The objective value will increase by $180 million.**
3. **The model becomes very inaccurate since it assumes that demand for Aardvarks does not change.**
4. **All of the above.**

Currently, there are no Deer being produced. What would be the list price at which deer would be produced?

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
D	0	-0.1	0.6	0.1	1.00E+30

1. $100
2. $500
3. $700
4. Deer would never be produced

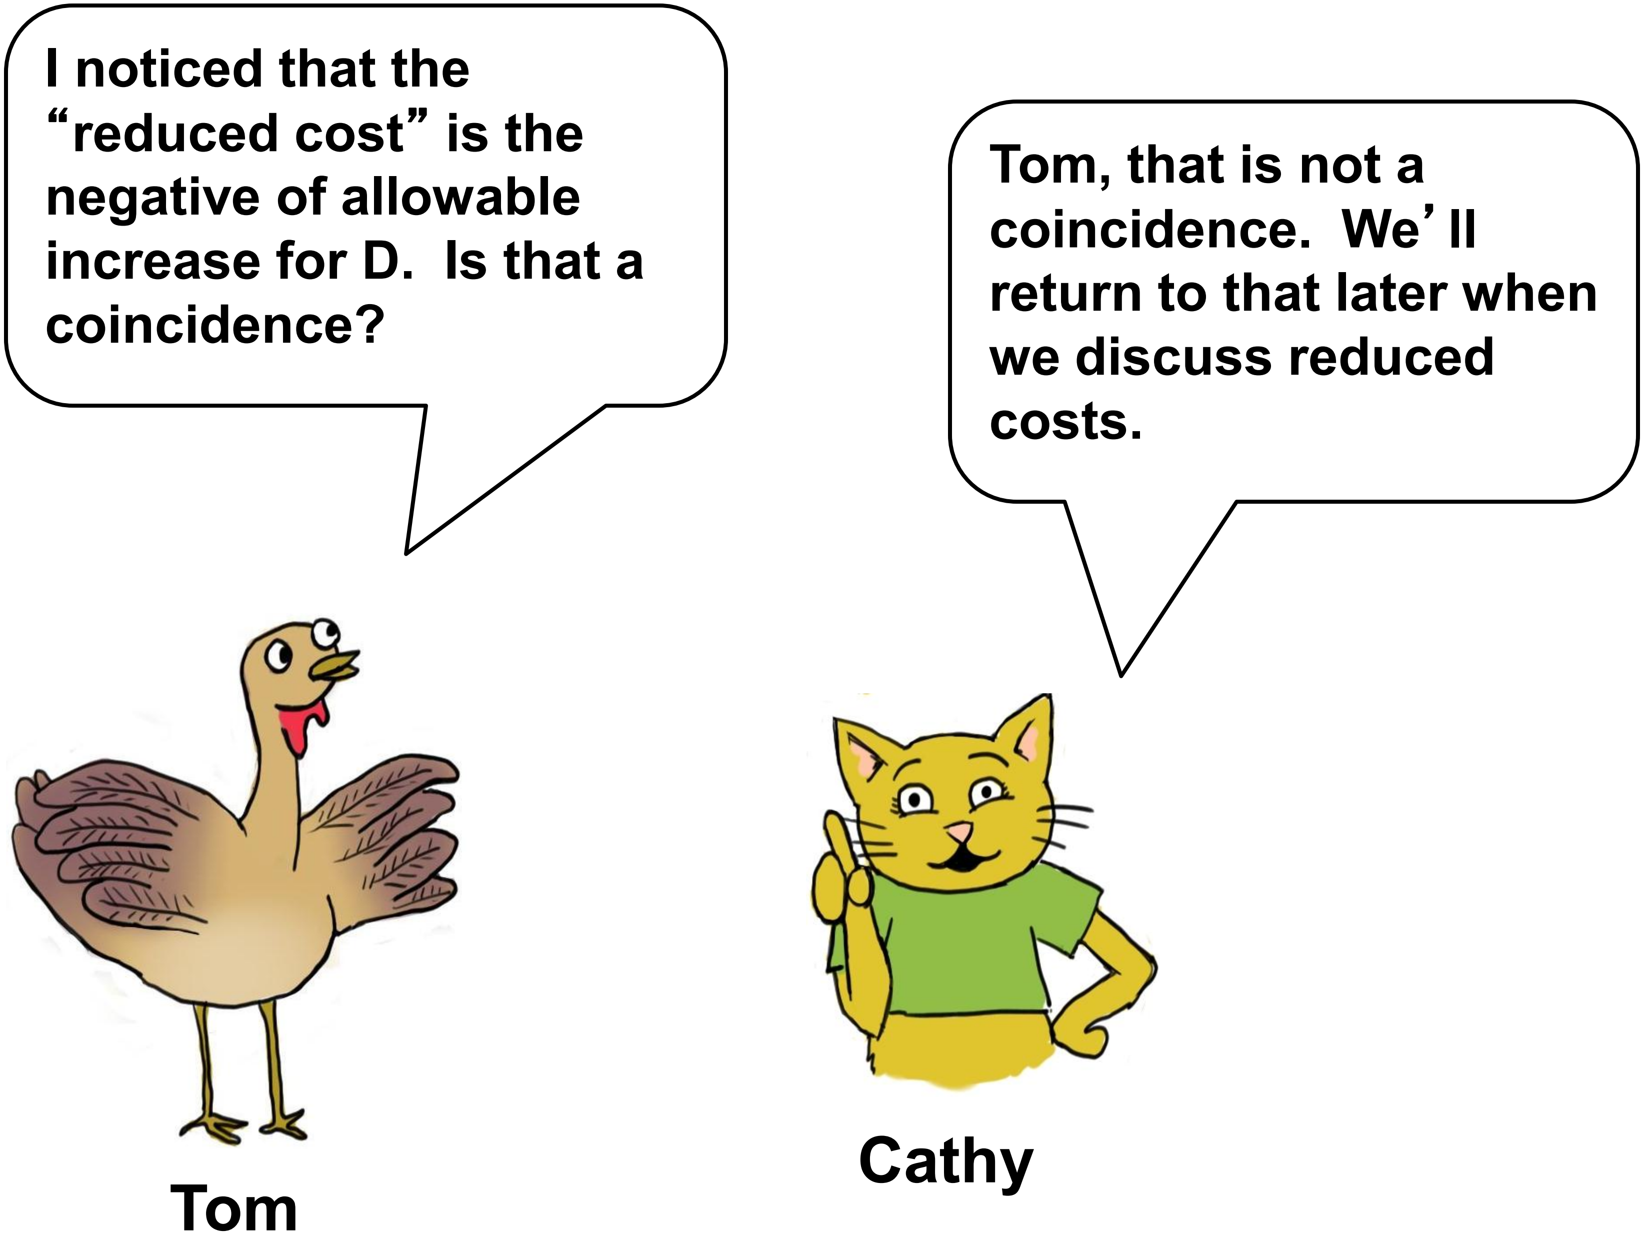
I noticed that the "reduced cost" is the negative of allowable increase for D. Is that a coincidence?
Tom, that is not a coincidence. We'll return to that later when we discuss reduced costs.
Tom
Cathy

Changing the RHS of a constraint

- mc^2 expects to receive 20,000 new drives. However, the drives are manufactured in a country experiencing labor strikes. What would be the impact on the optimal solution value if only 15,000 new drives were available?

We want to find out how the optimal plan would change if the number of drives were 15,000.

If we learned about labor strikes after starting to implement a plan, things would be far worse since it is costly to change a plan after it is implemented.

Shadow Prices

- **Definition:**
  - The shadow price of a constraint of a linear program is the increase in the optimal objective value per unit increase in the RHS of the constraint.

- **VERY IMPORTANT:**
  - The Shadow Price of the i-th constraint is ONLY valid within the RHS range of the i-th constraint.

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
New Drives	20	0.8	20	0.6	16.4

On the shadow price for new drives

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
New Drives	20	0.8	20	0.6	16.4

Shadow Price = .8

Shadow Price is valid if the RHS number ND of new drives satisfies:

$$3.6 \le ND \le 20.6$$

Nooz

The SR report tells how the objective value changes, but does not say what the new solution is.

It also doesn't tell you what happens if the RHS change is not in the allowable range.

Question. What is the "increase" in the optimal objective value if the number of disk drives is reduced to 15,000?

1. It cannot be determined from the data
2. $4,000
3. -$4,000
4. $4,000,000
5. -$4,000,000
6. $12,000,000

On the demand for tablet computers

- The total demand for tablets (A, B, C) is currently 38,000. What would be the value of increasing the demand to 40,000, possibly via a marketing campaign?

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
Max for tablets	37.4	0	38	1.00E+30	0.6

If an inequality constraint does not hold with equality in the optimum solution, the constraint is non-binding. The shadow price of each non-binding constraint is 0.

Question: mc^2 is considering an advertizing campaign that will increase the demand of Aardvarks to 18,500. The cost of the campaign is $400,000. Is it worthwhile?

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
Max for A	18	1.04	18	0.75	18

1. Yes
2. No
3. Cannot be determined from the available data.

Question: mc^2 is considering an advertizing campaign that will increase the demand of Aardvarks to 20,000. The cost of the campaign is $1,000,000. Is it worthwhile?

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
Max for A	18	1.04	18	0.75	18

1. Yes
2. No
3. Cannot be determined from the available data.

The real change in Z as the bound on A increases

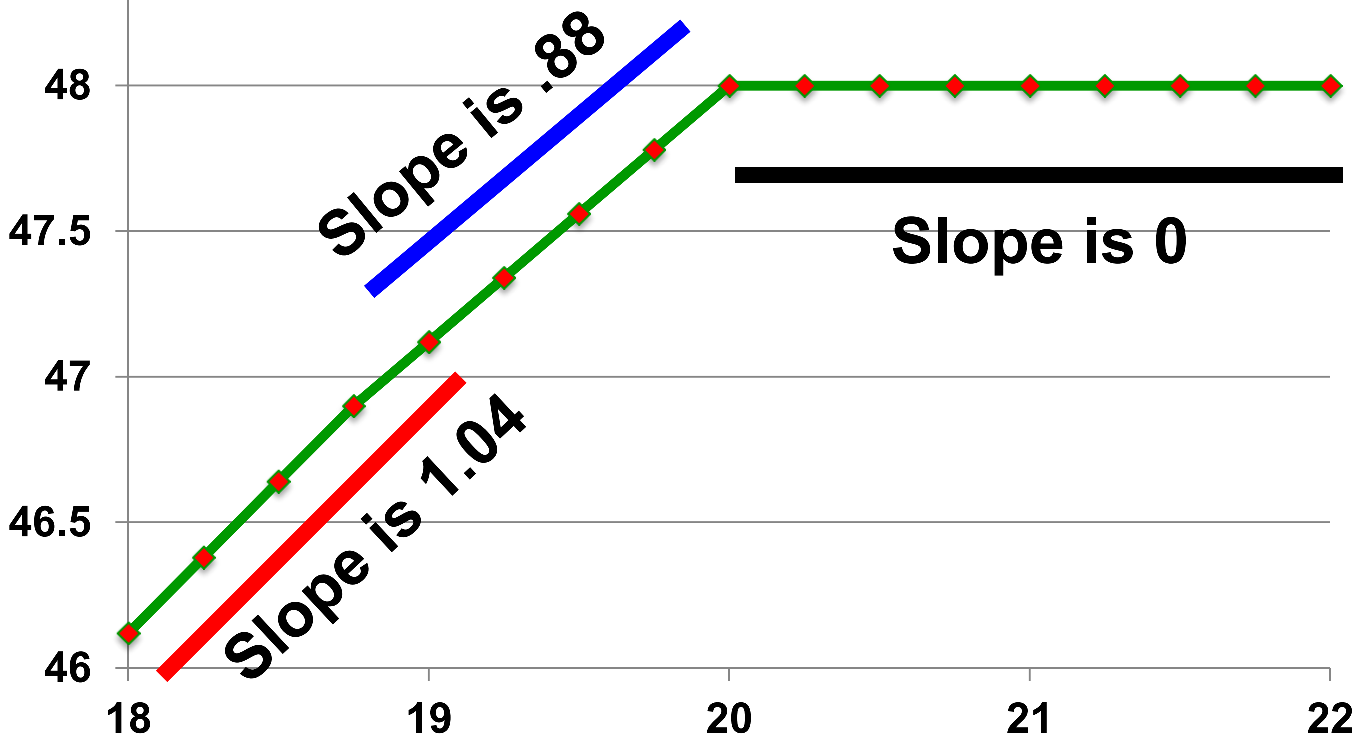

Midclass Break

Simultaneous Changes in the RHS

- One of the forecasters at mc^2 is pessimistic about the demand forecasts. One can't rely on the demand for A and C to be as high as predicted. It is much safer to use estimates of 15 and 2.
- What will be the impact on the optimum objective value if the maximum value of A is reduced to 15 and the maximum value of C is reduced to 2?

Name	Final Value	Shadow Price	Constraint R.H. Side	Allowable Increase	Allowable Decrease
Max for A	18	1.04	18	0.75	18
Max for C	3	0.6	3	0.6	3

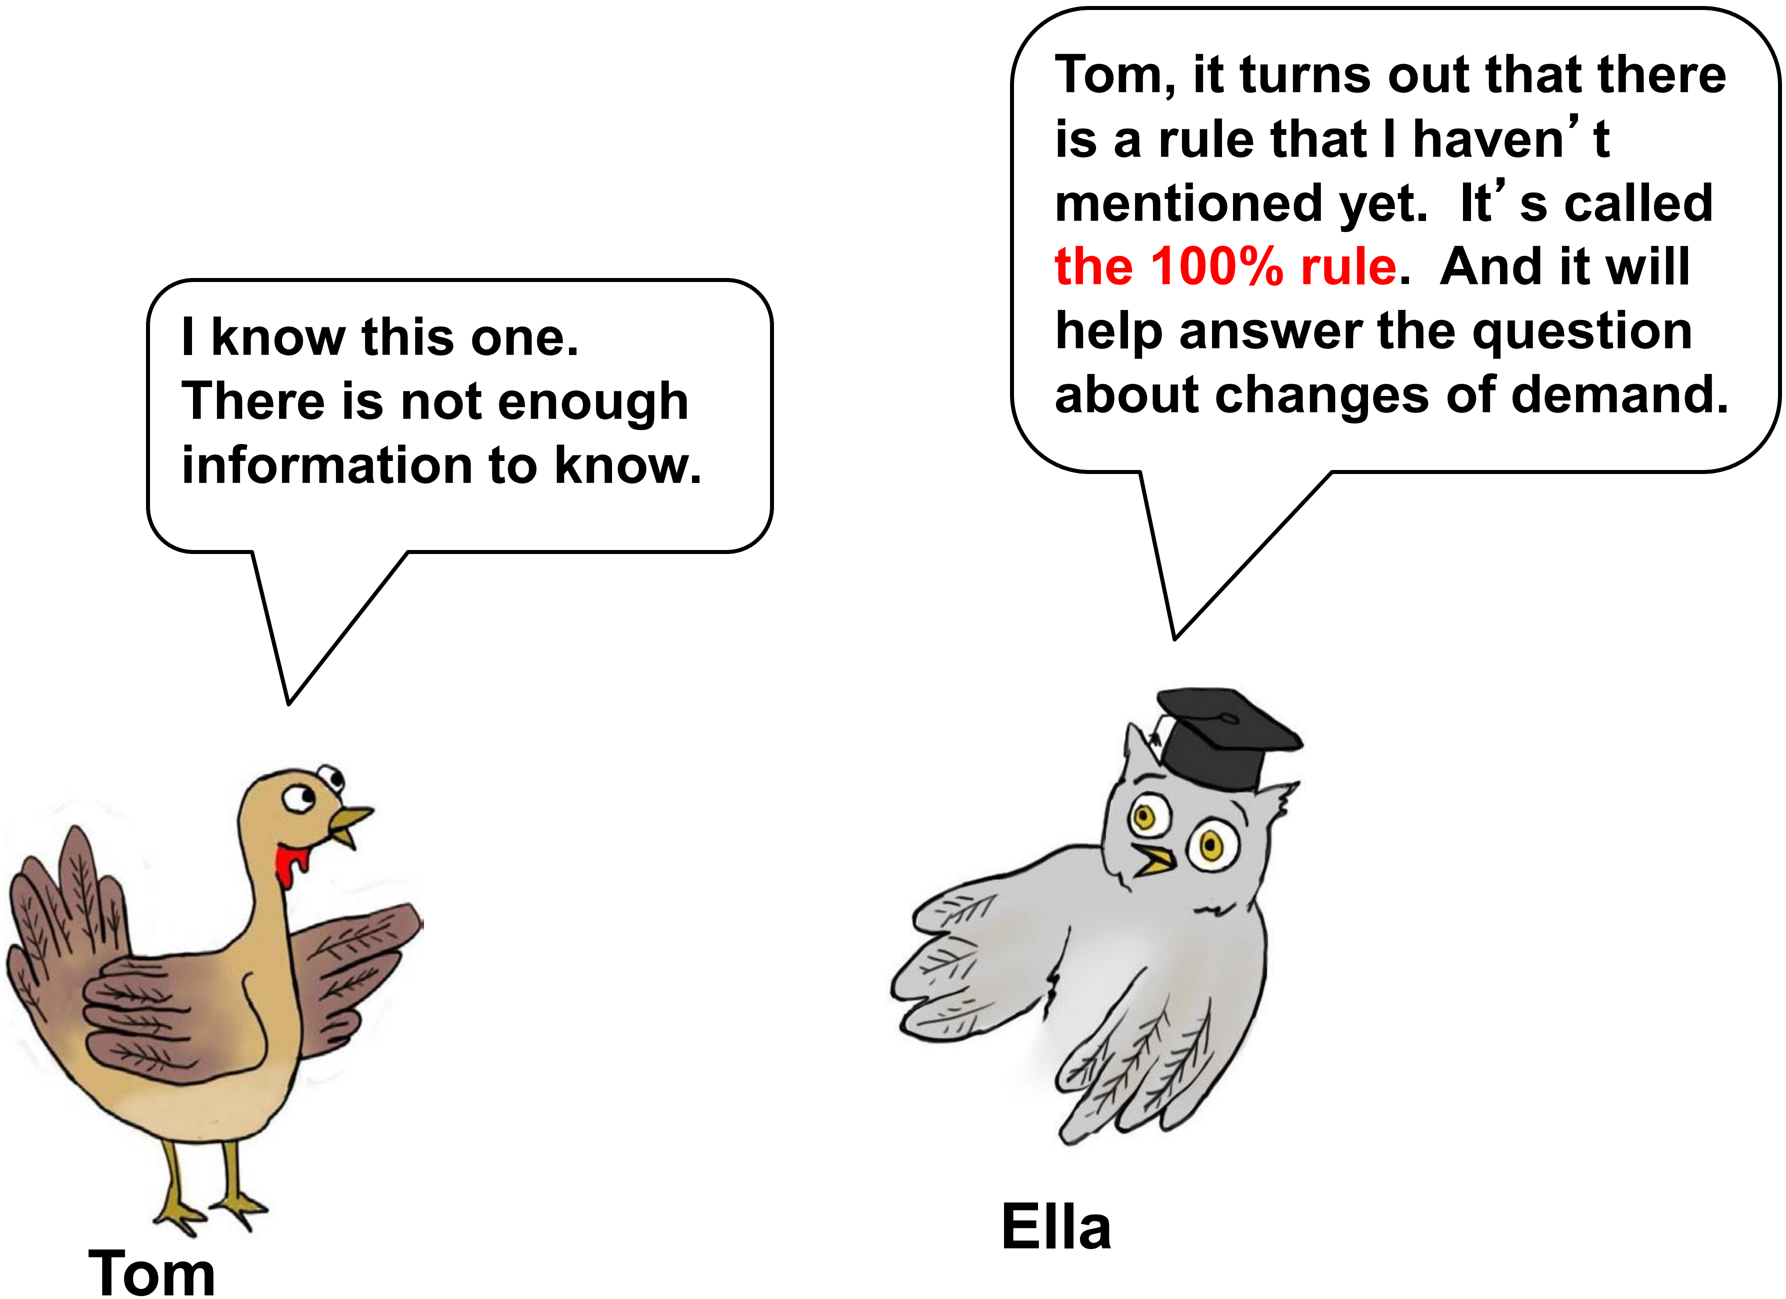
I know this one. There is not enough information to know.
Tom, it turns out that there is a rule that I haven’t mentioned yet. It’s called the 100% rule. And it will help answer the question about changes of demand.
Tom
Ella

A rule for two changes in RHS

Name	Shadow Price	R.H.S	Allowable Increase	Allowable Decrease
Max for A	1.04	18	0.75	18
Max for C	0.6	3	0.6	3

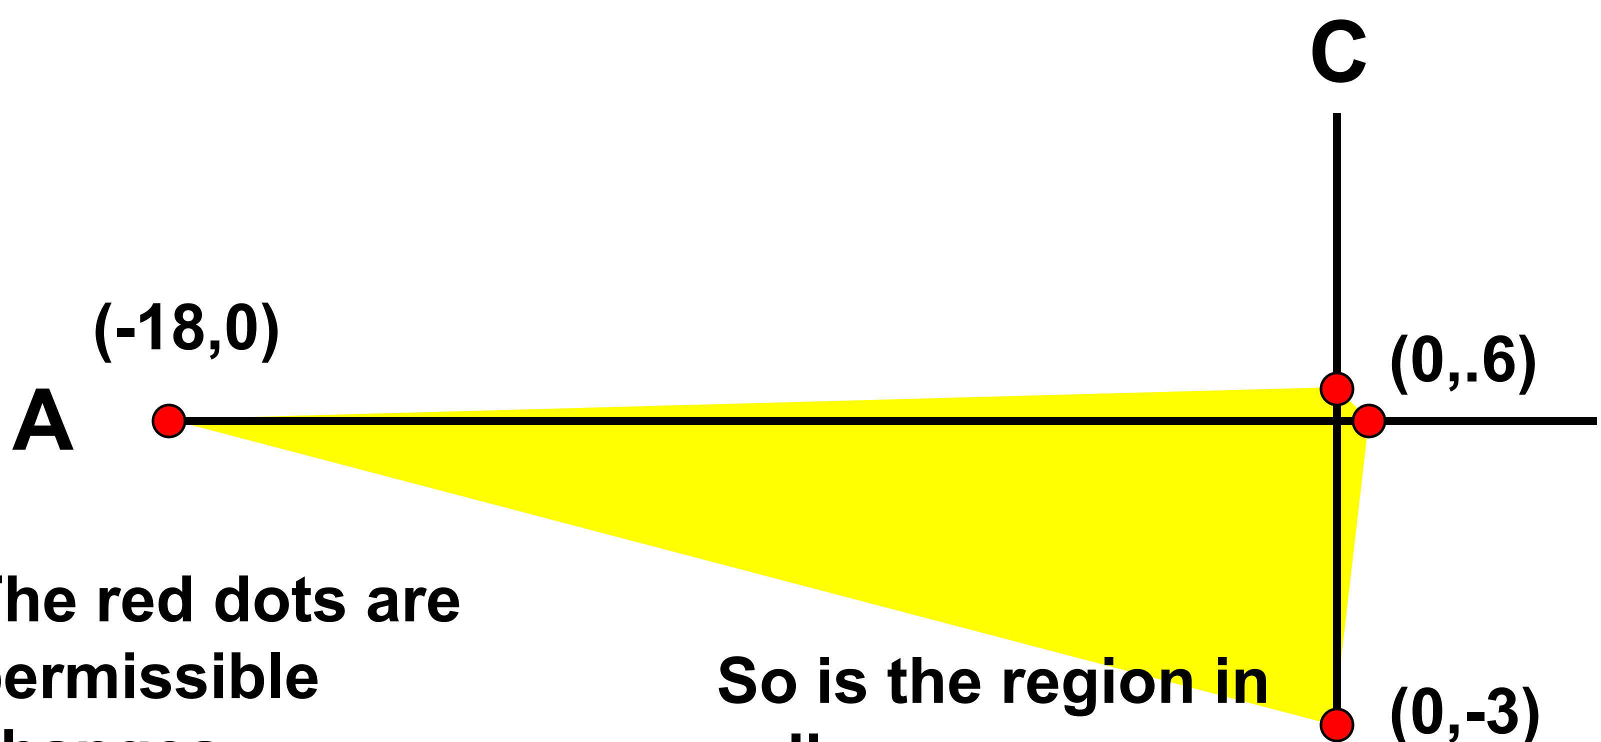

The red dots are permissible changes.

So is the region in yellow.

The 100% for two changes in RHS

Name	Shadow Price	R.H.S	Allowable Increase	Allowable Decrease
Max for A	1.04	18	0.75	18
Max for C	0.6	3	0.6	3

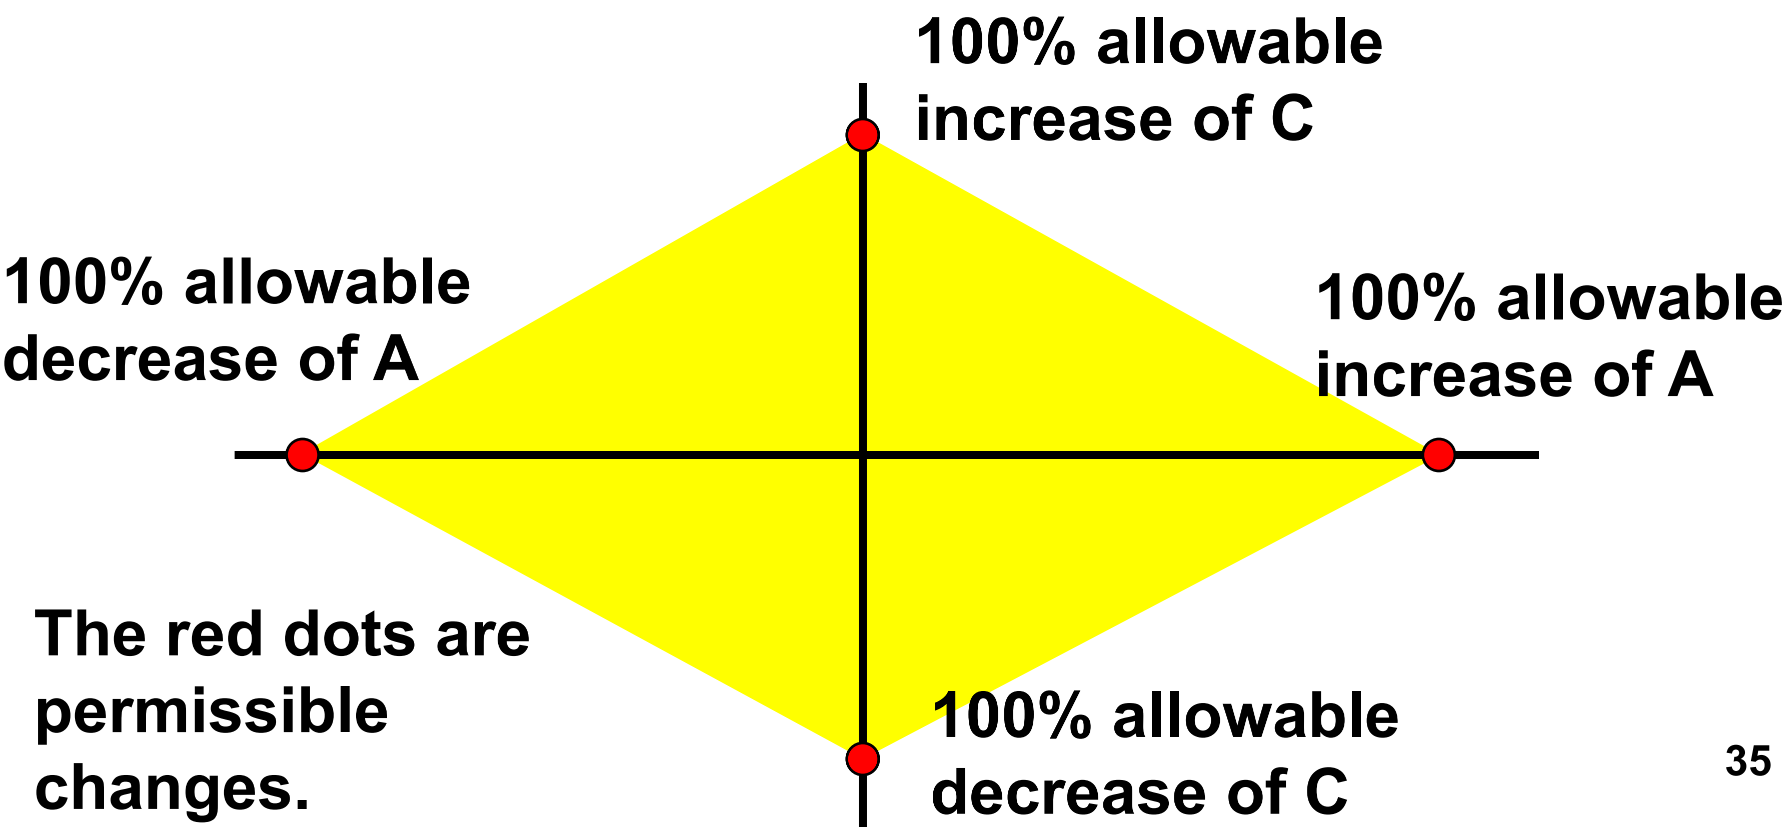

The red dots are permissible changes.

The 100% rule illustrated

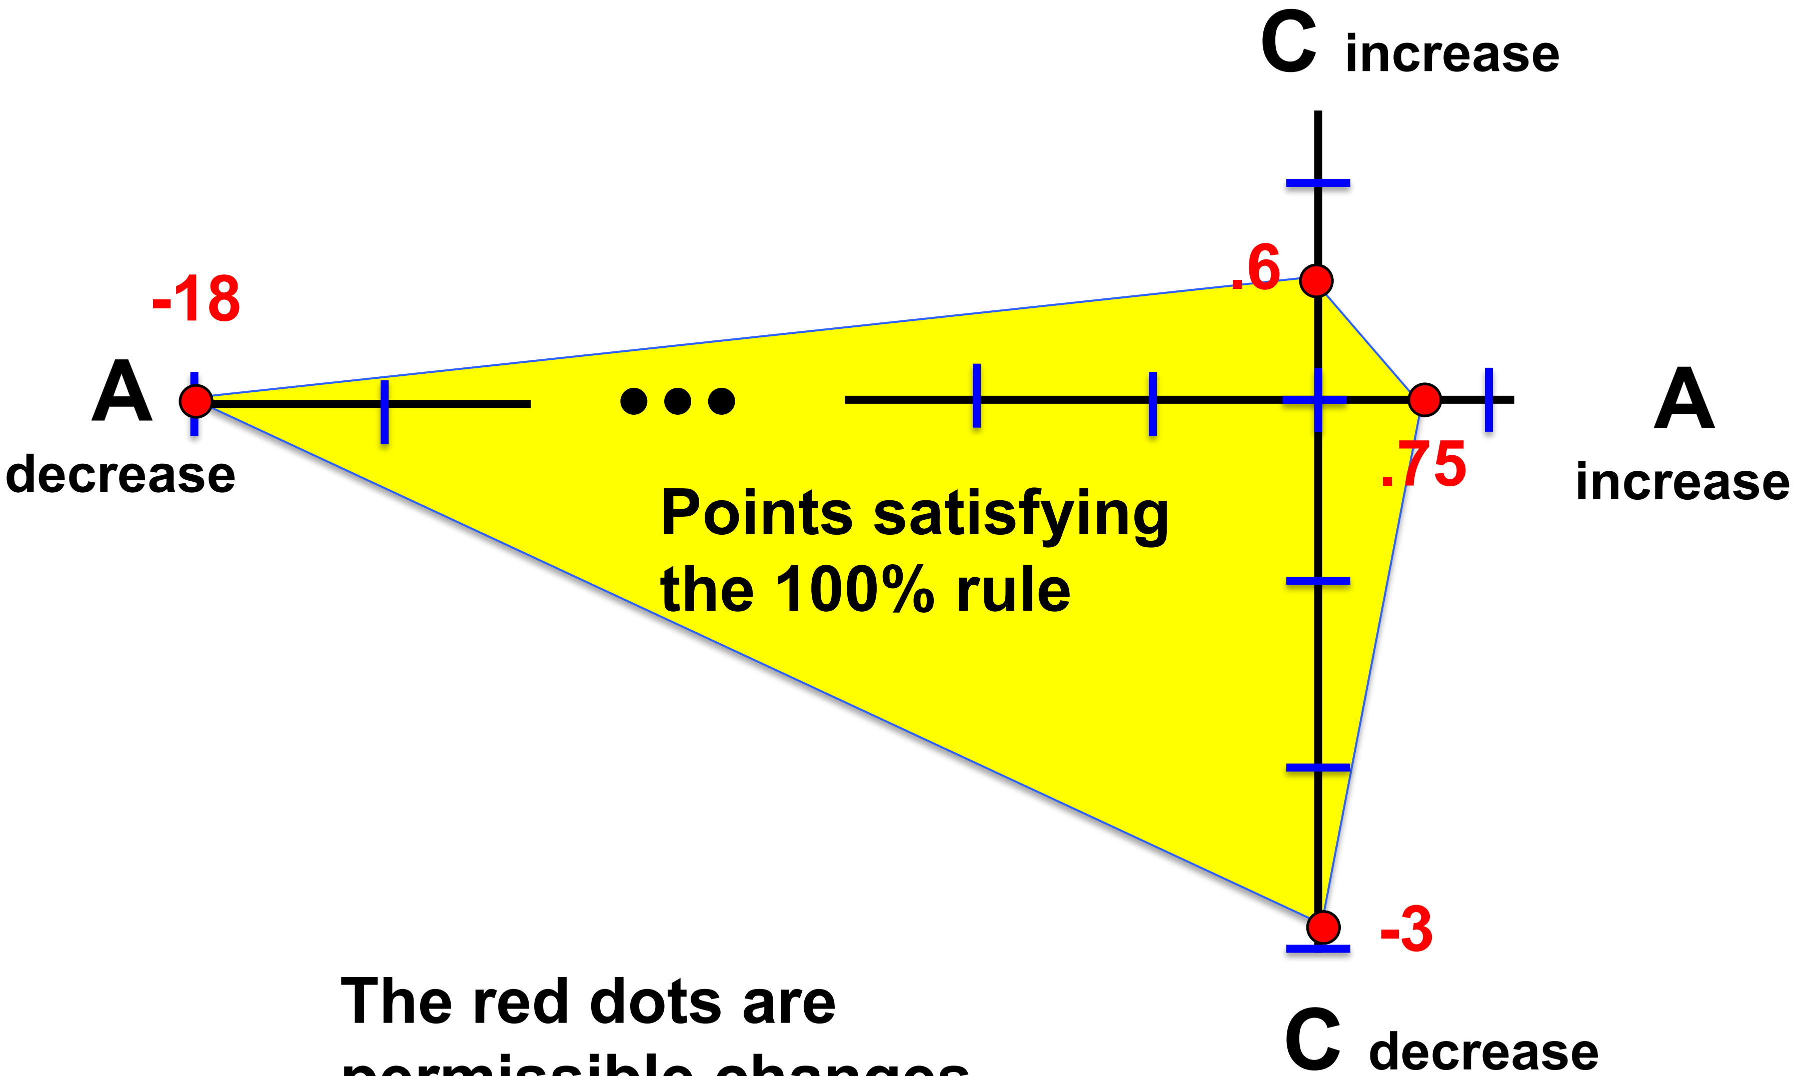

The red dots are permissible changes.

The actual increases in A and C so that the basis does not change.

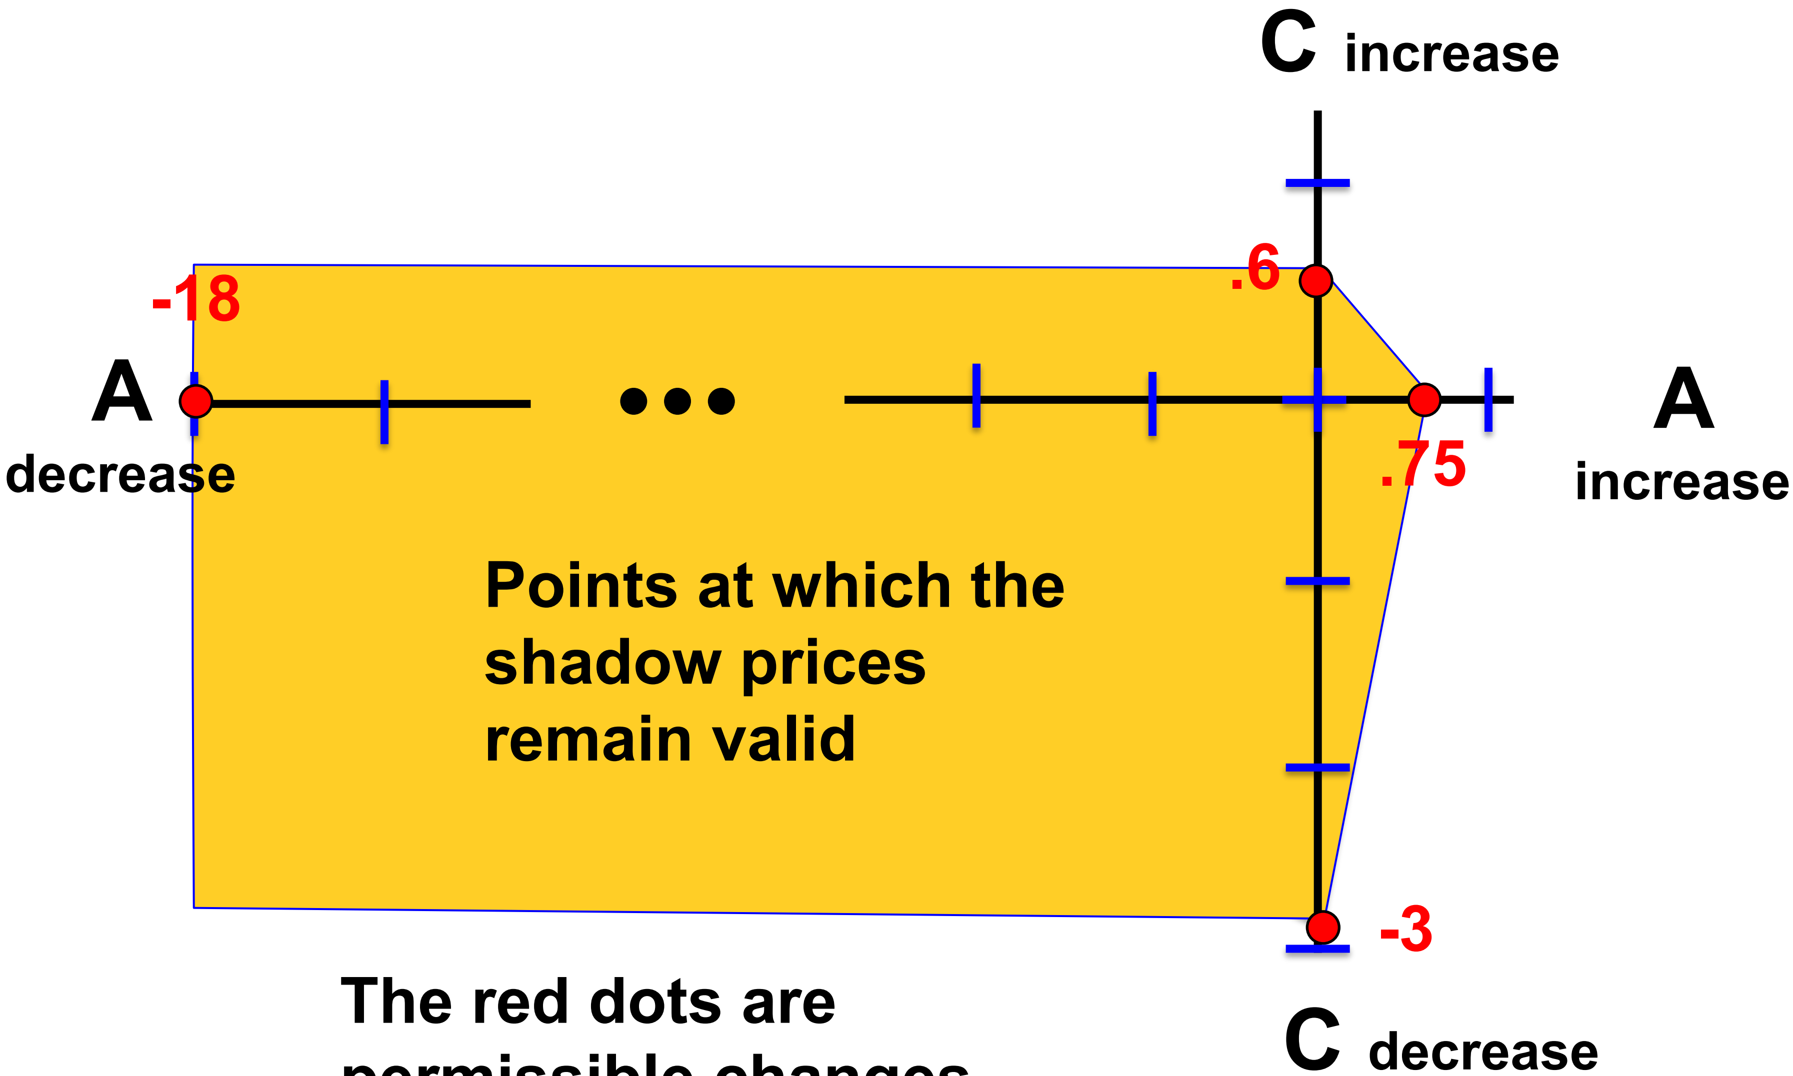

The red dots are permissible changes.

The 100% rule

For each RHS that changes, compute the amount of change divided by the total allowable change. Add up these fractions. If the total value is less than 1, then the shadow prices are valid.

Name	Shadow Price	R.H.S	Allowable Increase	Allowable Decrease	Proposed Decrease	proposed /allowable
Max for A	1.04	18	0.75	18	3	1/6
Max for C	0.6	3	0.6	3	1	1/3
					Total	1/2

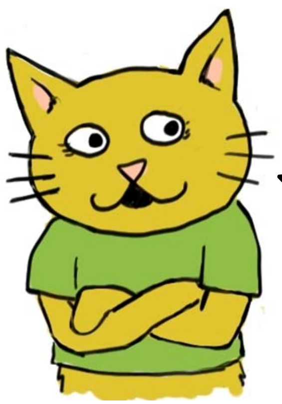

You can learn more about the 100% rule in Section 3.7 of AMP.

Objective "increases" by
$-3 \times 1.04 \;\; + -1 \times .6 = -3.72$

Revenue decreases by $3.72 million

Question. Suppose that the Maximum for A Decreased by 1000 (to 17,000), and the Maximum for C increased by 500 (to 3500). What would be the increase in the optimum revenue?

Name	Shadow Price	R.H.S	Allowable Increase	Allowable Decrease
Max for A	1.04	18	0.75	18
Max for C	0.6	3	0.6	3

1. It cannot be determined from the data
2. - $.74 million
3. $1.34 million
4. -$1.34 million

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
D	0	-0.1	0.6	0.1	1.00E+30
E	32	0	0.3	1.00E+30	0.1

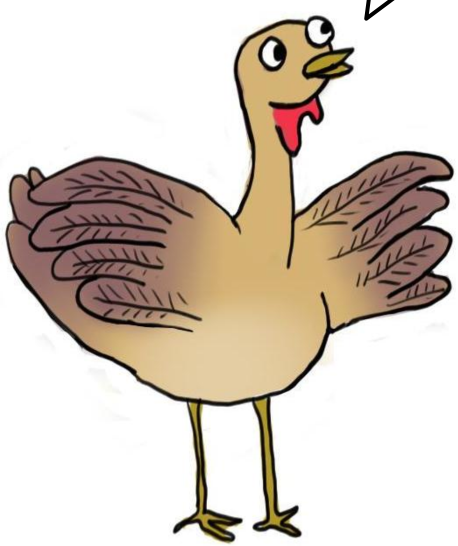

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
D	0	-0.1	0.6	0.1	1.00E+30
E	32	0	0.3	1.00E+30	0.1

Definition 1. The reduced cost for a variable is the shadow price for the non-negativity constraint.

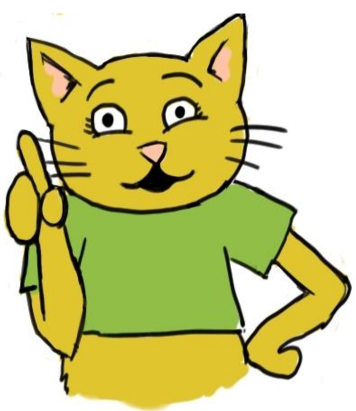

If we change the constraint "D ≥ 0" to D ≥ 1" the objective "increases" by -0.1

If we change the constraint "E ≥ 0" to E ≥ 1" the objective "increases" by 0.

Name	Final Value	Reduced Cost	Objective Coef	Allowable Increase	Allowable Decrease
D	0	-0.1	0.6	0.1	1.00E+30
E	32	0	0.3	1.00E+30	0.1

Although the reduced cost is a kind of shadow price, the allowable increase column refers to the change that can be made in the cost while keeping the same solution optimal.

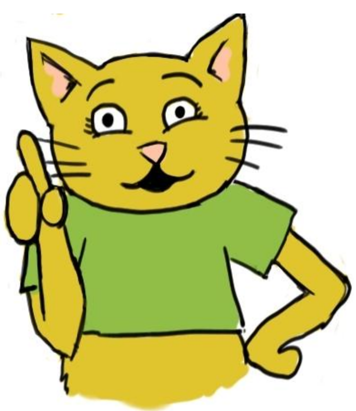

If we increase the profit of D by less than .1, the same solution stays optimal. If we increase it by more than .1, then the solution will change. In this case, D will be positive in the new optimal solution.

-z	A	B	C	D	E
1	0	0	0	-0.1	0

The objective row of the final tableau.

The second definition is that the reduced cost is the objective value coefficient of a variable in the final and optimal tableau.

Wow. I am dumbfounded by the interconnections.

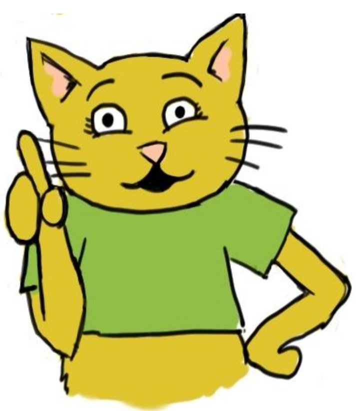

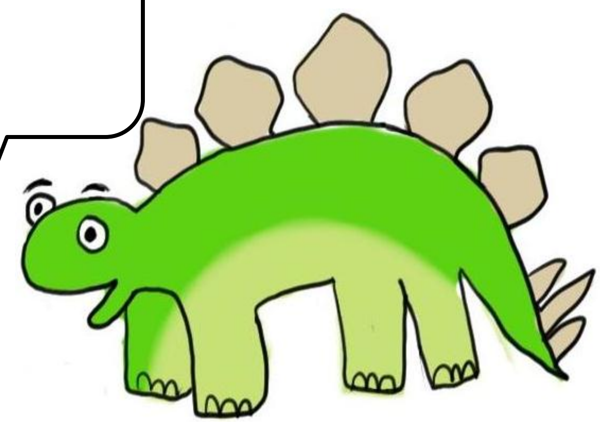

Reduced costs are the costs in the z-row.

Basic Var	-z	x_1	x_2	x_3	x_4	x_5		RHS
-z	1	0	-2	0	0	6	=	-11
x_3	0	0	2	1	0	2	=	4
x_4	0	0	-1	0	1	-2	=	1
x_1	0	1	6	0	0	3	=	9

The reduced costs in the S.A. refer to the optimal reduced costs (those of the optimal tableau).

Basic Var	-z	x_1	x_2	x_3	x_4	x_5		RHS
-z	1	0	-8	-3	0	0	=	-23
x_5	0	0	1	0.5	0	1	=	2
x_4	0	0	1	1	1	0	=	5
x_1	0	1	3	-1.5	0	0	=	3

-z	A	B	C	D	E
1	0	0	0	-0.1	0

The objective row of the final tableau.

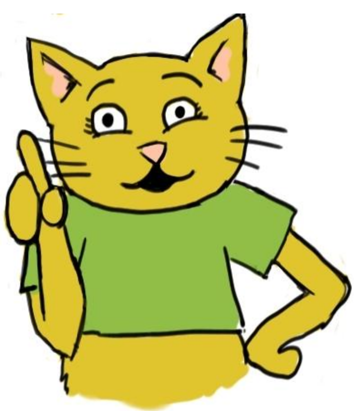

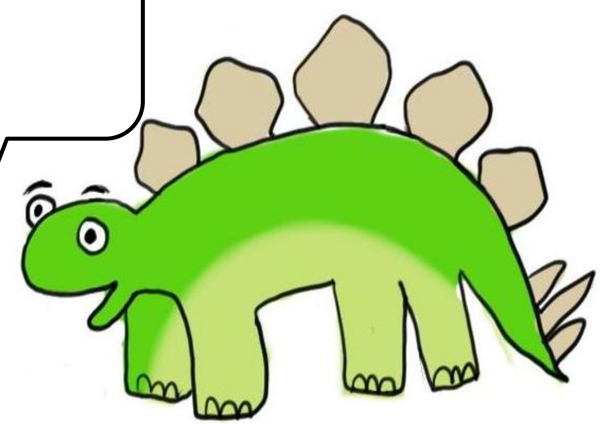

Pricing Out

Reduced cost of D =

original objective value minus the sumproduct of the RHS and the constraint coefficients for the variable.

	D	E	Shadow Price		Pricing out D
obj.	0.6	0.3			.6
	0	0	0		- 0 × 0
	2	1	0		- 2 × 0
	0.5	0	0.8		- .5 × .8
	0	0	1.04		- 0 × 1.04
	0	0	0.6		- 0 × .6
	0	0	0		- 0 × 0
	1	1	0.3		- 1 × .3
red. cost	-.1	0			= -.1

Pricing Out the Constraints

	A	B	Shadow Price	
obj.	1.2	0.8		
	2	0	0	
	0	2	0	
	0.2	1	0.8	
	1	0	1.04	
	0	0	0.6	
	1	1	0	
	0	0	0.3	
red. cost				

Pricing Out a new Variable

F	Shadow Price	
.65		
0	0	
2	0	
0.4	0.8	
0	1.04	
0	0.6	
0	0	
1	0.3	

Suppose that a new e-reader is introduced, codenamed the fox.

F uses 2 low capacity memory chips and has .4 new hard drives on average.

It lists for $650. Is it profitable to produce F?

Summary

- Shadow Prices
- Ranges
- Reduced costs
- Pricing out