Optimization Methods in Management Science 15.053

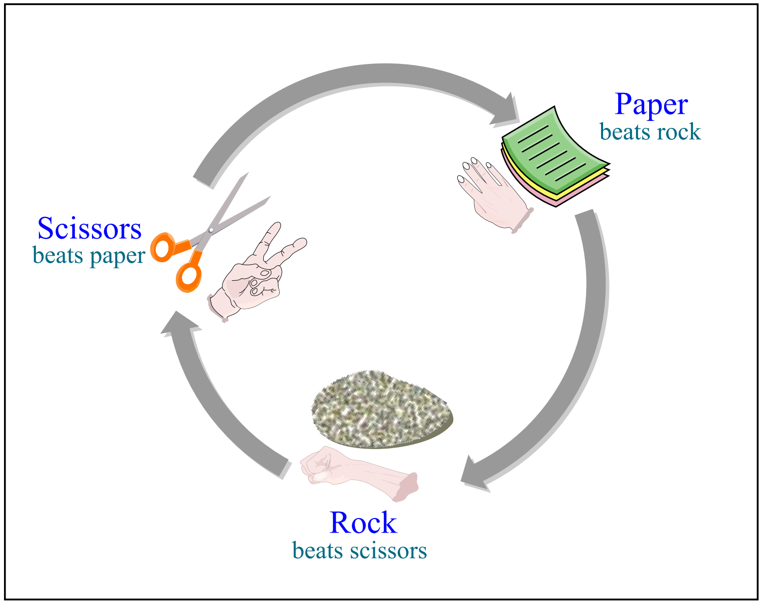

Image by MIT OpenCourseWare.

2-person 0-sum Games

written by Ebrahim Nasrabadi
with help from Jim Orlin

2-person 0-sum Games

Game theory is a branch of Mathematics with a wide variety of applications in economics, management science, political science, and engineering. It aims to models situations in which multiple participants interact or affect each other's outcome.

The 2-person 0-sum game is a basic model in game theory. There are two players, each with an associated set of strategies. While one player aims to maximize her payoff, the other player attempts to take an action to minimize this payoff. In fact, the gain of a player is the loss of another.

In this tutorial, we introduce 2-person 0-sum game theory, present some useful concepts, and discuss how each player can determine her optimal strategy.

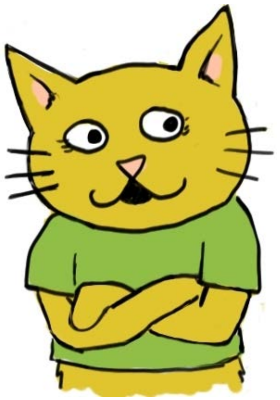

Key elements of a 2-person game:

Each 2-person game consists of

- 2 players;
- Strategies available to each player;
- Payoffs for each player;
  - the payoff is the amount of benefit or loss that a player derives if a particular outcome happens.
  - the payoff of each player depends on her choice, **and also** depends on the choice of the other player.

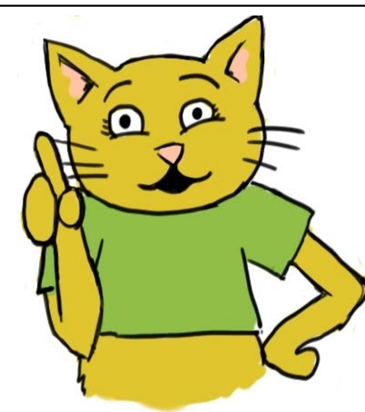

2-person 0-sum Games:

In 2-person 0-sum games the payoff function f can be represented as follows.

$$f : S_1 \times S_2 : \to \mathbb{R}$$
$$(s_1, s_2) \to f(s_1, s_2)$$

If Player 1 chooses strategy s_1 and Player 2 selects strategy s_2, then Player 1 will get $f(s_1,s_2)$ and Player 2 will get $-f(s_1,s_2)$. We will refer to $f(s_1,s_2)$ as the ***value*** of the game. Player 1 aims to maximize $f(s_1,s_2)$, while Player 2 attempts to minimize this value.

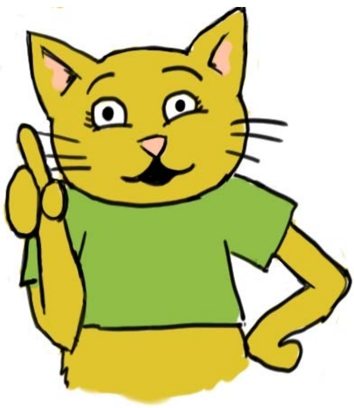

Even-or-Odd Game:

Let's play a game!

That sounds great! Where does one get these dollars?

OK. This game is called "evens and odds" and it is also called "coin matching." I'll describe it on the next slide.

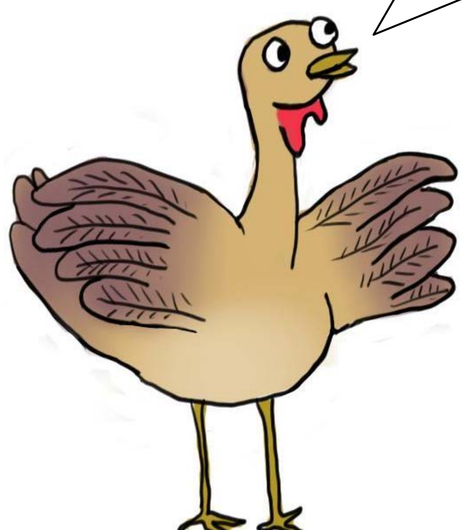

At the end of the tutorial, I'll let you know how to play the game optimally. If you can guess the optimum strategy before then, that would be very cool.

2-person 0-sum games

Here is another example of a game.

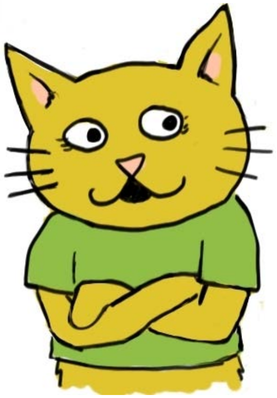

Coin Matching Problem (Even-Odd Game)

- There are two players: Player 1 (the row player) and Player 2 (the column player);
- Each player simultaneously shows a coin.
- If both coins are showing are heads, then Player 1 wins $2 (paid by Player 2). If both coins are tails, then Player 1 wins $4. If the coins do not match, then Player 1 loses $3.

Describing a game

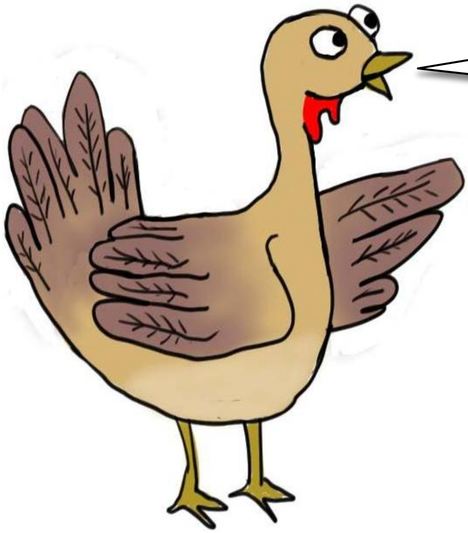

Is each game completely described by the payoff matrix?

That's right! Here is a more formal description.

Normal-form is a simple way to describe a 2-person 0-sum game by using a so-called ***payoff matrix***:

- Each row (column) of the matrix corresponds to a strategy available to Player 1 (Player 2). In this case, we refer to Player 1 as the row player (or simply R) and Player 2 as the column player (or simply C).
- The i-j element of the matrix gives the payoff to the row player if she chooses i-th row and the column player selects j-th column. The matrix is called a payoff matrix.

A payoff matrix

		Tom 1	Tom 2
Chaty	1	2	-3
	2	-3	4

Here is a payoff matrix for the Even-Odd game.

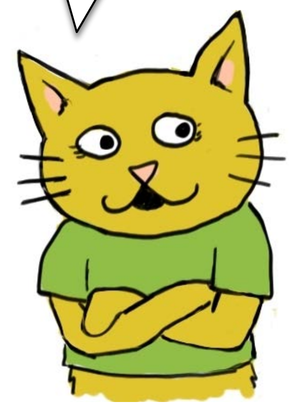

It is a good way to view all the basic elements of a game.

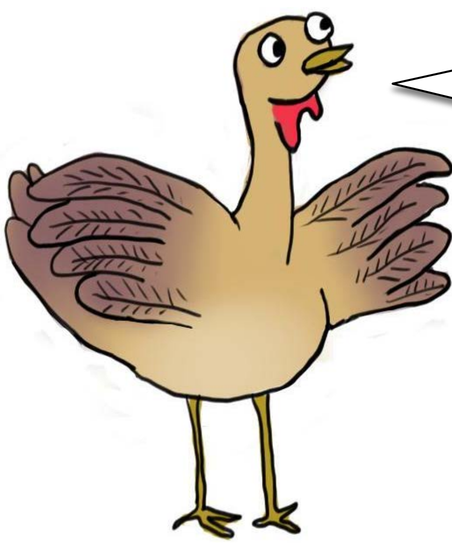

Yes, it is! In the rest of the tutorial, we will use this way to represent a game.

A payoff matrix

Player C

Player R

-2	1	2
2	-1	0
1	0	-2

Here is one more example. Each player has three strategies:

- Player R chooses a row: either row 1, or row 2, or row 3;
- Player C chooses a column: either column 1, or column 2, or column 3.

We will later refer to these as ***pure strategies***, for reasons that will become apparent when we describe mixed strategies

This matrix is the payoff matrix for Player R, and Player C gets the negative).

How much do R and C get if R chooses 1 and C selects 2?

How much do R and C get if R chooses 3 and C selects 3?

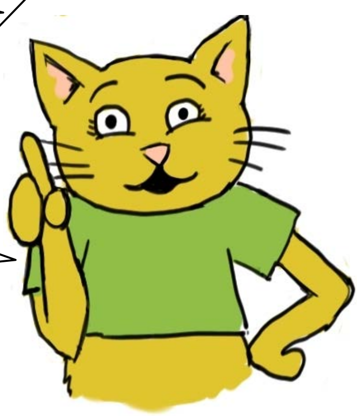

A payoff matrix

Player C

Player R

-2	1	2
2	-1	0
1	0	-2

Here is one more example. Each player has three strategies:

- Player R chooses a row: either row 1, or row 2, or row 3;
- Player C chooses a column: either column 1, or column 2, or column 3.

We will later refer to these as ***pure strategies***, for reasons that will become apparent when we describe mixed strategies

This matrix is the payoff matrix for Player R, and Player C gets the negative). For example,

- If R chooses 1 and C selects 2, then R gets +1 and C get -1
- If R chooses 3 and C selects 3, then R gets -2 and C get +2.

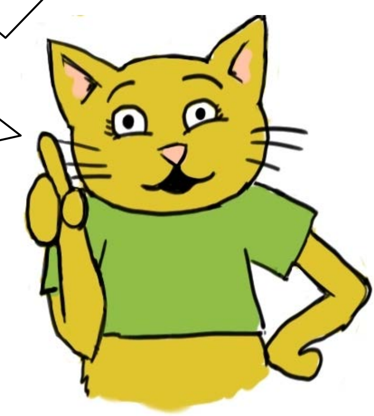

A guaranteed payoff for the Row Player

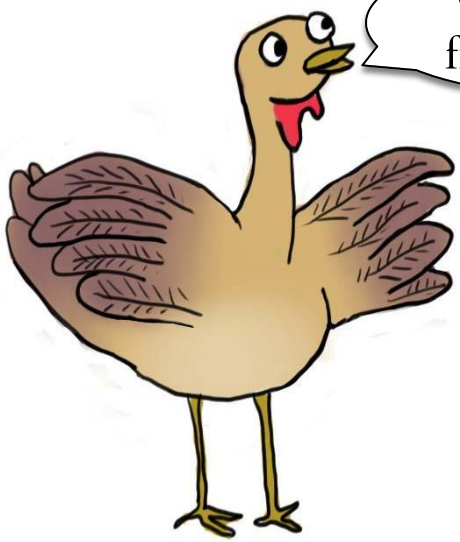

Who goes first?

In games, two players make decision **at the same time.**

Player C

Player R			
	-2	1	2
	2	-1	0
	1	0	-2

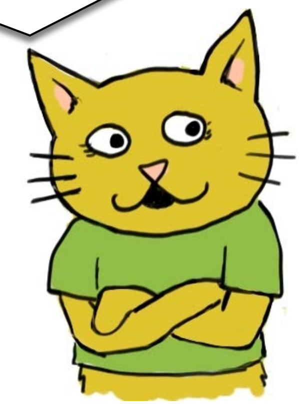

But for the purpose of this example, suppose that Player R were forced to announce a row before Player C makes her decision. If Player R announces

- Row 1, then the player's C best response is Column 1 and R will get -2;
- Row 2, then the player's C best response is Column 2 and R will get -1;
- Row 3, then the player's C best response is Column 3 and R will get -2;

Player R wishes to maximize her payoff and her best pure choice is to announce Row 2. In fact, she takes a ***maximin strategy*** to maximize her minimum payoff. This gurantees a payoff of at least -1 to Player R, regardless of the player's C strategy.

A guaranteed maximum payoff to the Row Player

If Player C announces

- Column 1, then R's best response is Row 2 and R gets 2.
- Column 2, then R's best response is Row 1 and R gets 1.
- Column 3, then R's best response is Row 1 and R gets 2.

Player C wants to minimize the payoff of Player R, and thus her best pure strategy (if she went first) is to announce Column 2. This strategy is called a ***minimax*** pure strategy. It minimizes the maximum payoff from Player C. This shows that the value of the game for the Row Player can always be limited to at most 1.

Player C

Player R

-2	1	2
2	-1	0
1	0	-2

If one player chooses prior to the other.

It seems that going first is always a disadvantage.

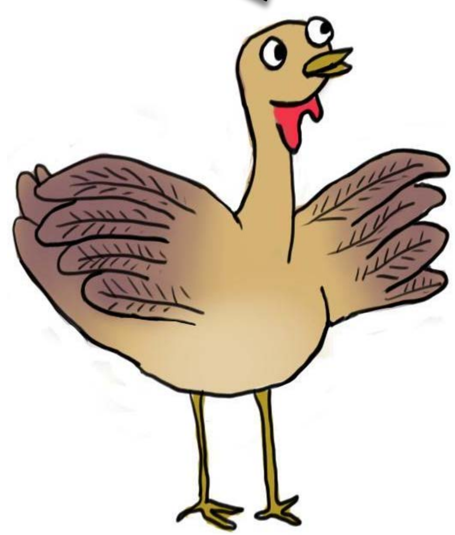

Good point. And it implies an important mathematical result. If Player R chooses a strategy ***before*** Player C, R can guarantee a payoff of at least $\max_{s_1 \in S_1} \min_{s_2 \in S_2} f(s_1, s_2)$.

If Player C chooses a stategy before Player R, then C can guarentee that R receives at most $\min_{s_2 \in S_2} \max_{s_1 \in S_1} f(s_1, s_2)$.

This implies that

$$\max_{s_1 \in S_1} \min_{s_2 \in S_2} f(s_1, s_2) \le \min_{s_2 \in S_2} \max_{s_1 \in S_1} f(s_1, s_2).$$

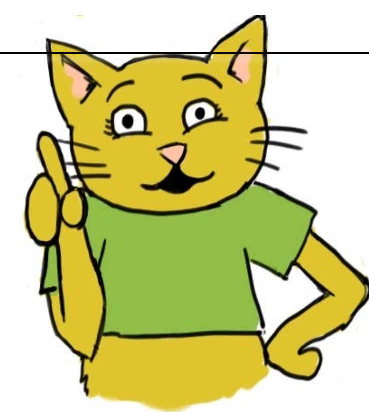

Saddle-point

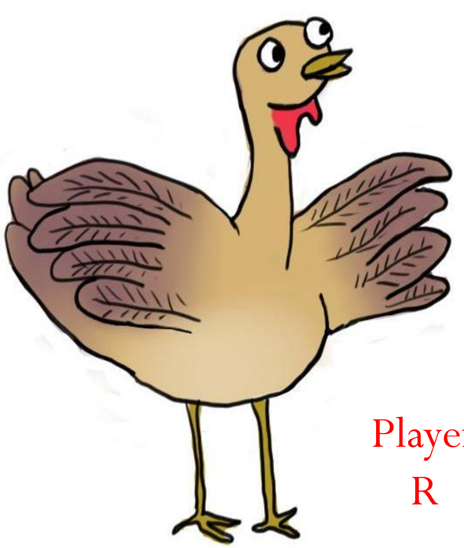

Can the equality hold in some games?

Yes! It may be the case that the lower bound and the upper bound on the value of the game coincide. In this case, there are strategies $s^*_1 \in S_1, s^*_2 \in S_2$ such that

$$f(s^*_1, s^*_2) = \max_{s_1 \in S_1} \min_{s_2 \in S_2} f(s_1, s_2) = \min_{s_2 \in S_2} \max_{s_1 \in S_1} f(s_1, s_2).$$

The pair $\boldsymbol{s^*_1, s^*_2}$ is called a ***saddle point*** of the game. It is also called a ***pure Nash equilibrium*** since no player has an intensive to change her strategy.

Player C

Player R

-1	-2	-1
2	2	1
-1	-1	0

In this example, the pure Nash equilibrium occurs when Player R chooses 2 and C selects 3.

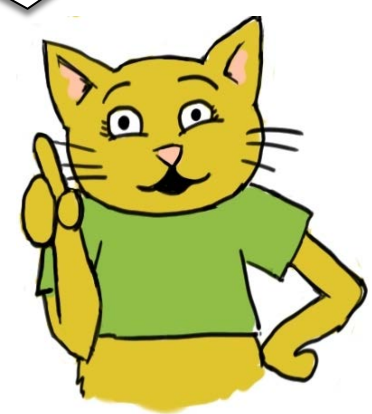

Mixed strategies

Since R and C go at the same time, it would seem pretty dumb for R to announce in advance what she will choose. And she needs to mix it up. What would happen if she used a coin to decide whether to select Row 1 or 2?

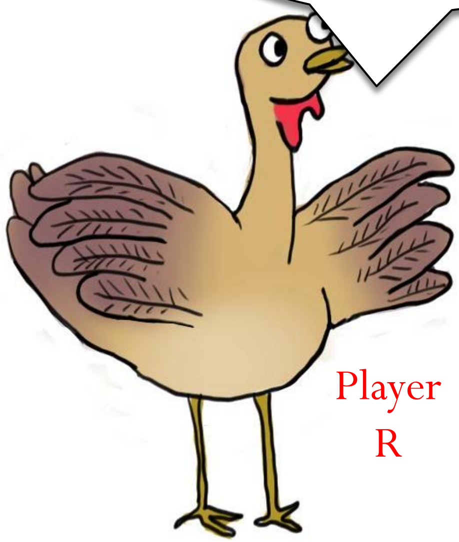

What you said is a ***mixed*** (also called ***randomized***) ***strategy*** since R assigns a probability to each row and selects a row accordingly. This plays a key role in game theory. In this case, R might do pretty well even if she (stupidly) announced her strategy in advance, so long as C couldn't see the result of the coin flip. Let's see what is she Player's C best response. If C chooses Column

- 1, then R gets -2 with probability 0.5 or 2 with probability 0.5; So R's expected payoff is 0.
- 2, then R gets 1 with probability 0.5 or 1 with probability 0.5; So R's expected payoff is 0.
- 3, then R gets 0 with probability 0.5 or 2 with probability 0.5; So R's expected payoff is 1.

Since C aims to minimize the player's R payoff, she will choose Column1 or Column 2. Thus, R can guarantee an expected payoff of at least 0, which is much better than she could guarantee before.

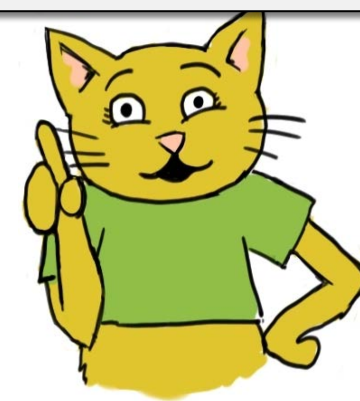

Player

Player R	Prob			
	0.5	-2	1	2
	0.5	2	-1	0
	0	1	0	-2

Optimal mixed strategies

Great! Can she do even better?

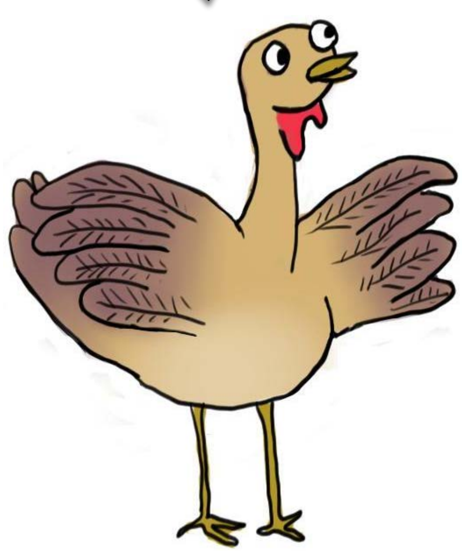

Yes. Here is the way to model the problem of finding her best strategy. For each row i, let the decision variable x_i be the probability of selecting row i. If C chooses

- Column 1, then R's expected value is $P_1 := (-2)x_1 + 2x_2 + 1x_3 = -2x_1 + 2x_2 + x_3$;
- Column 2, then R's expected value is $P_2 := 1x_1 + (-1)x_2 + 2(0)x_3 = x_1 - x_2$;
- Column 3, then R's expected value is $P_3 := 2x_1 + (0)x_2 + (-2)x_3 = 2x_1 - 2x_3$.

Thus Player R's expected value is at least

$$\min\{P_1, P_2, P_3\}.$$

Player R will assign the probabilities x_1, x_2 and x_3 in such a way to maximize $\min\{P_1, P_2, P_3\}$ in order to determine her best mixed strategy.

	Prob	Player C		
Player R	$x1$	-2	1	2
	$x2$	2	-1	0
	$x3$	1	0	-2

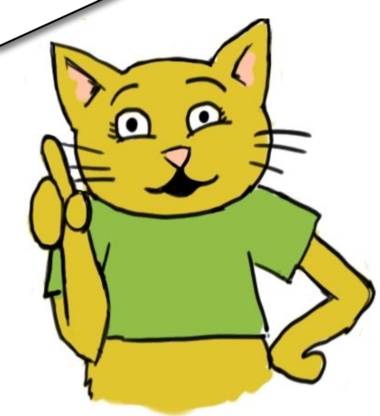

I see that computing the best mixed strategy for Player R is an optimization problem. Is it a linear program?

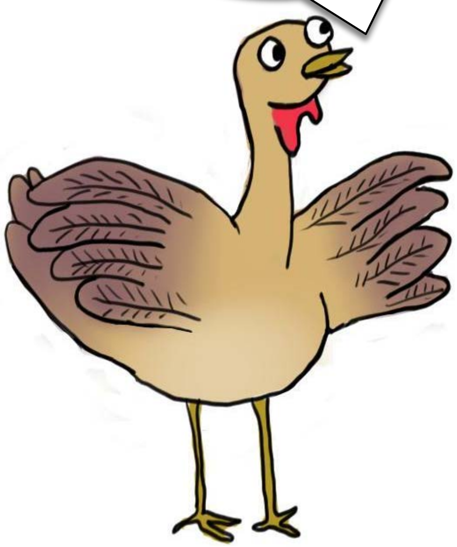

Yes it is. Here is the optimization problem:

$$\max \ \min\{P_1,P_2,P_3\}$$
$$P_1=-2x_1+2x_2+x_3,$$
$$P_2= x_1-x_2,$$
$$P_3=x_2+2x_3,$$
$$x_1+x_2+x_3=1,$$
$$x_1,x_2,x_3\geq 0.$$

Notice that probabilities must sum to 1, since Player R is obligated to choose a row. In addition, probabilities can never be negative! These are our constraints.

Introduce a new variable z to be $\min=\{P_1,P_2,P_3\}$.

Then, we can express the above problem as a linear program.

$$\max \ z$$
$$z \leq -2x_1+2x_2+x_3,$$
$$z \leq x_1-x_2,$$
$$z \leq x_2+2x_3,$$
$$x_1+x_2+x_3=1,$$
$$x_1,x_2,x_3 \geq 0.$$

The optimal solution is $x_1= 7/18$, $x_2=5/18$, $x_3=1/3$ with optimal value $Z=1/9$.

So, with a mixed strategy R guarantees obtaining at least 1/9.

How can we obtain an upper bound for R when she is permitted to choose a mixed strategy?

As we have already observed an upper bound for R (or equivalently a lower bound for C) can be computed if Player R is in a position to take her strategy after hearing the Player's C strategy.

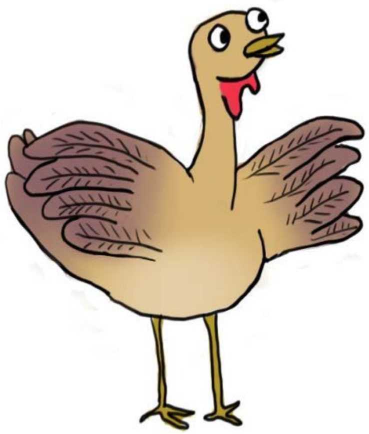

Player C

Prob	y_1	y_2	y_3
Player R	-2	1	2
	2	-1	0
	1	0	-2

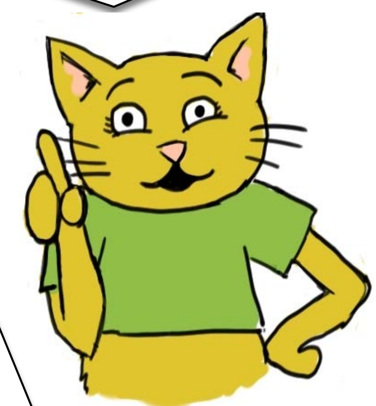

Let y_j be the probability of selecting column j, for j=1,2,3. If R chooses

- Row 1, then R's expected payoff is $P_1 := (-2)y_1 + y_2 + (2)y_3 = -2y_1 + y_2 + 2y_3$;
- Row 2, then R's expected payoff is $P_2 := 2y_1 + (-1)y_2 + (0)y_3 = 2y_1 - y_2$;
- Row 3, then R's expected payoff is $P_3 := (1)y_1 + (0)y_2 + (-2)y_3 = y_1 - 2y_3$.

Player R wants to maximize her expected payoff, so Player R max expected payoff is

$$\max\{P_1, P_2, P_3\}.$$

Therefore, Player C must assign the probabilities y_1, y_2 and y_3 in such a way to minimize $\max\{P_1, P_2, P_3\}$ in order to determine her best mixed strategy.

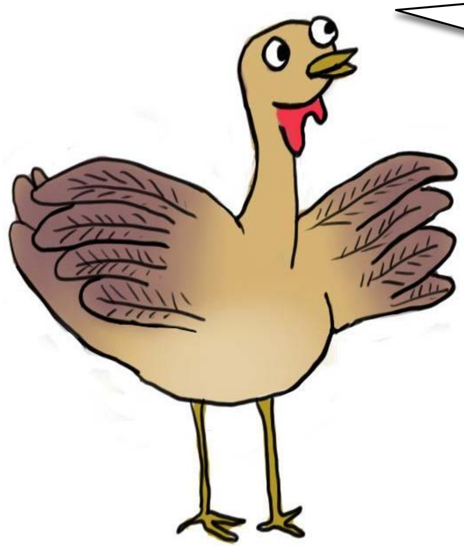

Thanks! I think we can write a linear problem to determine an optimal mixed strategy for Player C in the same manner.

Exactly! Here is the linear program:

$$
\begin{aligned}
&\min \ w \\
&\quad w \geq -2y_1+y_2+2y_3; \\
&\quad w \geq 2y_1-y_2\ ; \\
&\quad w \geq y_1-2y_3, \\
&\quad y_1+y_2+y_3 = 1, \\
&\quad y_1,y_2,y_3 \geq 0.
\end{aligned}
$$

The optimal solution is $y_1= 1/3$, $y_2=5/9$, $y_3=1/9$ with optimal value $w=1/9$.

So, with this random strategy R gets only 1/9.

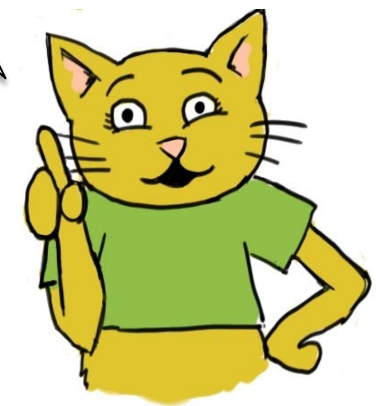

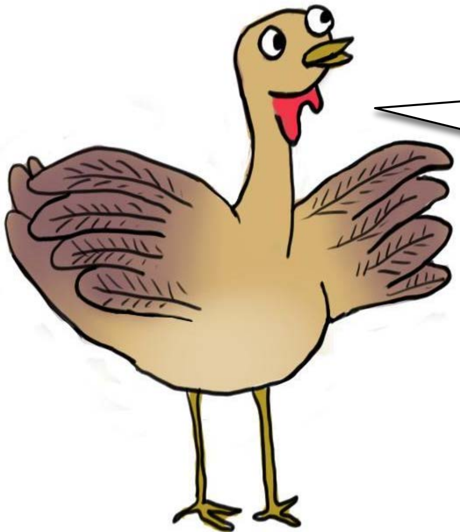

Wait a moment! I see that the payoff for player R is the same, whether Player R announces her strategy first or Player C. Is that a coincidence?

It's no coincidence that the optimal average payoff to the game is 1/9, assuming that both players play optimally, and it does not matter who goes first. This result holds for 2-person 0-sum games in general:

For 2-person 0-sum games, the maximum payoff that R can guarantee by choosing a random strategy is the minimum payoff to R that C can guarantee by choosing a random strategy.

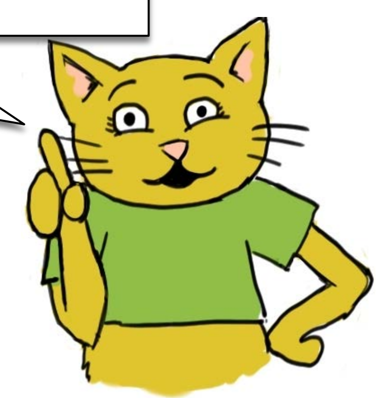

2-person 0-sum games in general

- Let x denote a random strategy for R, with value z(x) and let y denote a random strategy for C with value w(y). Then

$$z(x) \le w(y) \text{ for all } x, y.$$

- The optimum x* can be obtained by solving an LP. So can the optimum y*. In addition,

$$z(x^*) = w(y^*).$$

- In other words, the maximum payoff that R can guarantee by choosing a random strategy is the minimum payoff to R that C can guarantee by choosing a mixed strategy.

- Notice that z(x*) is an upper bound on the payoff of Player R and w(y*) is a lower bound on the payoff of Player C. Player R wishes to maximize her payoff, while Player C attempts to minimize her payoff. Since z(x*) = w(y*), neither player can benefit by a unilateral change in strategy, even when each player is aware of the other player's strategy. In this case, the pair (x*,y*) is called a ***mixed Nash equilibrium***.

A 2-dimensional view of game theory

In the class, I heard we can easily to solve 2-person 0-sum game graphically theory when there are two strategies per player. How does it work?

Notice that

- if R goes first and decides on a strategy of choosing row 1 with probability p and row 2 with probability 1-p, then the strategy for C is easily determined.
- so R can determine the payoff as a function of p, and then choose p to maximize the payoff.

Follow me in the next slides to illustrate this method with the Even-or-Odd game.

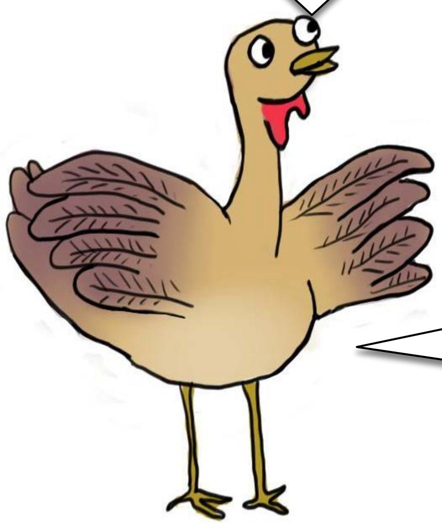

Thanks! Learning from examples is my favorite way of learning. It is way more fun than learning from mistakes.

I want to minimize the payoff. What would my optimal strategy be?

We graph the payoff as a function of the probability p given that I select Row 1.

I'll choose column 1 with probability p and column 2 with probability 1-p.

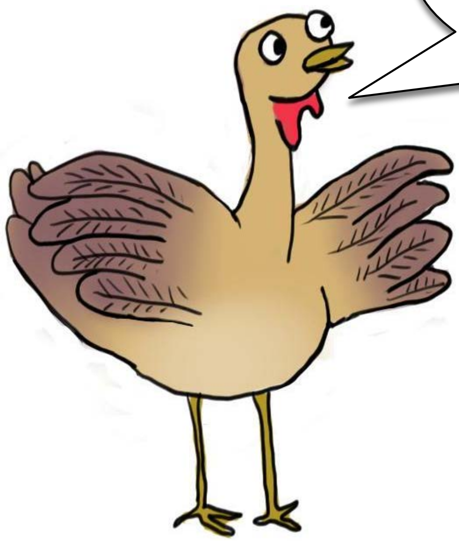

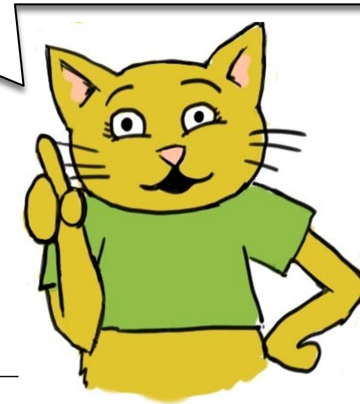

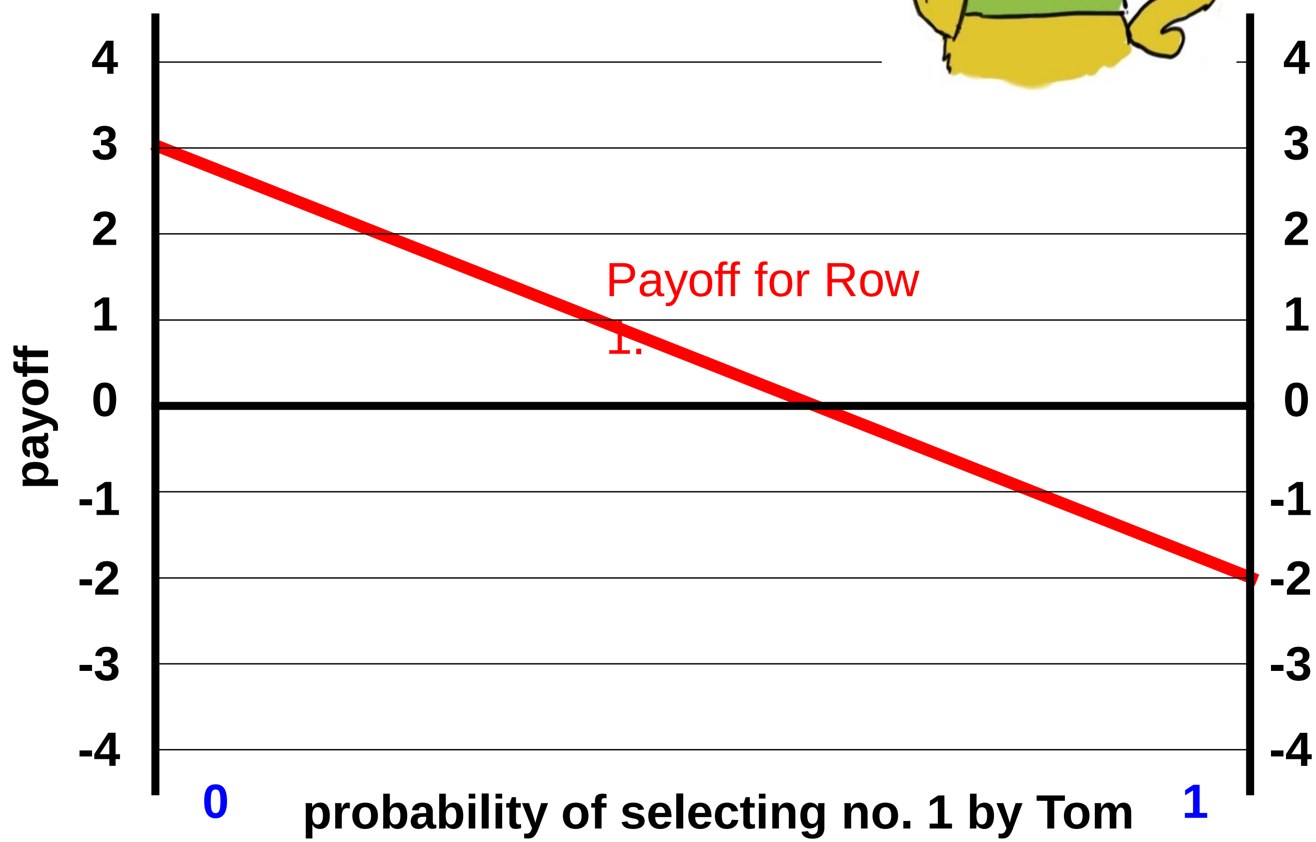

		Tom 1	Tom 2
Cathy	1	-2	3
	2	3	-4

We graph the payoff for Row 2 as p goes from 0 to 1.

Here is the payoff if you choose Row 2.

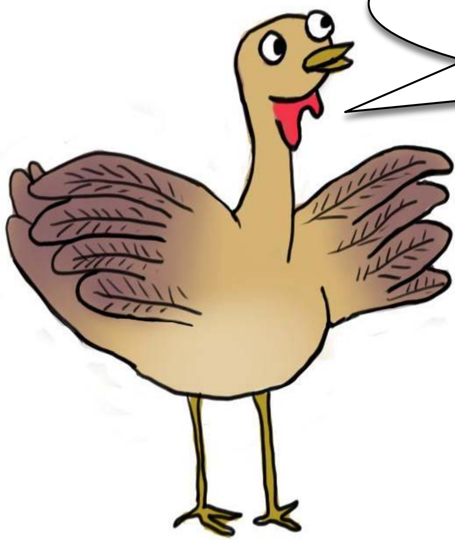

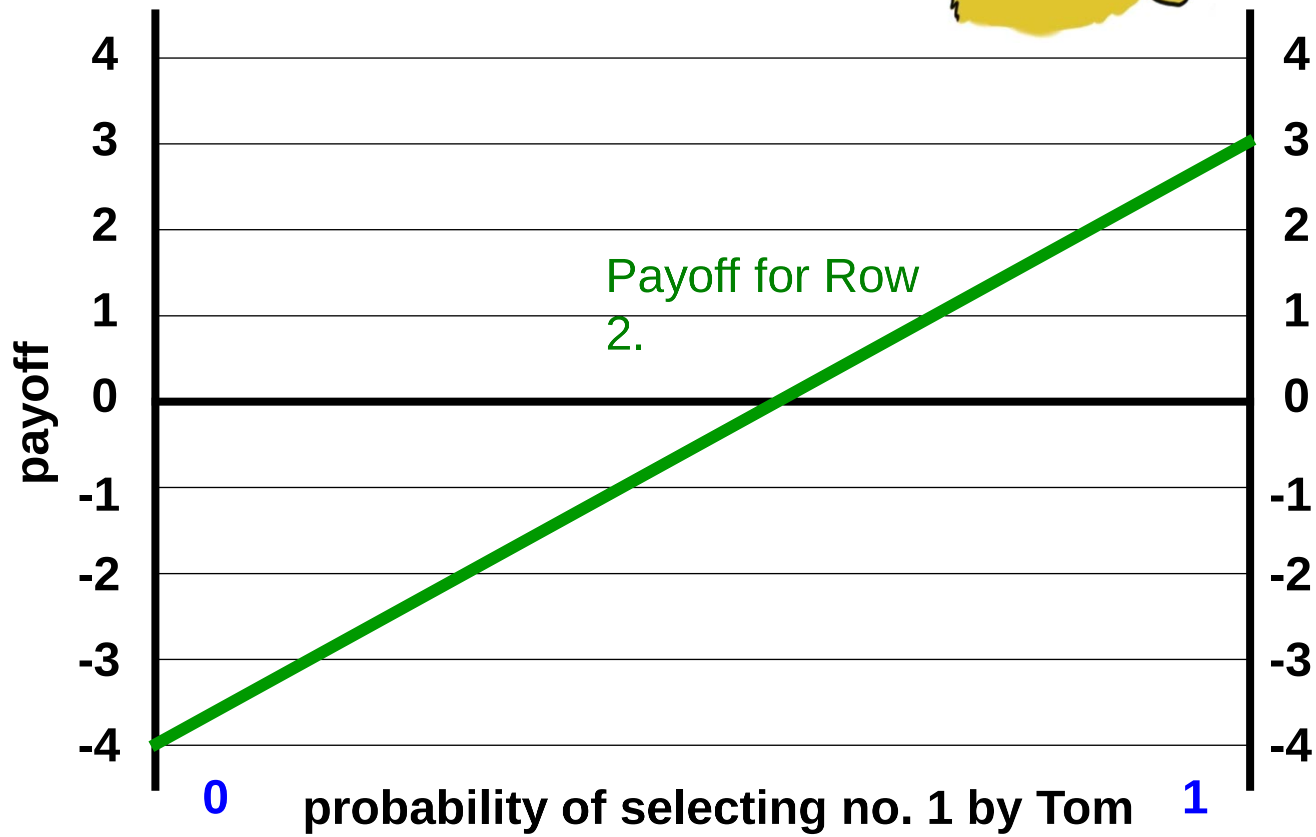

Tom

		1	2
Cathy	1	-2	3
	2	3	-4

So, what column should I pick?

I want to maximize the payoff to me. So here is my strategy as a function of p:

- if p<7/12, my best response is Row 1;
- If p>7/12, my best response will be Row 2;

Since you want to minimize my payoff, your best strategy is to set p=7/12. This guarantees that you can limit my expected winnings to at most 1/12.

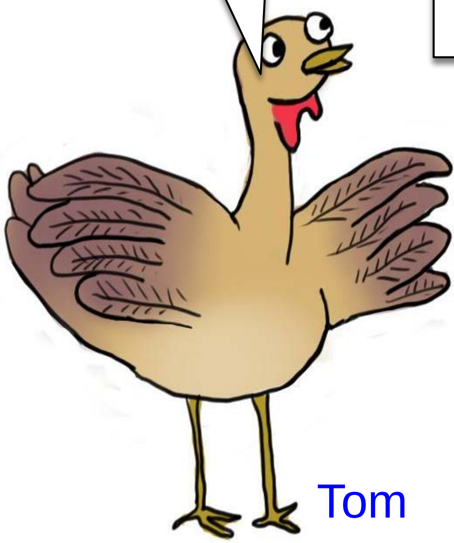

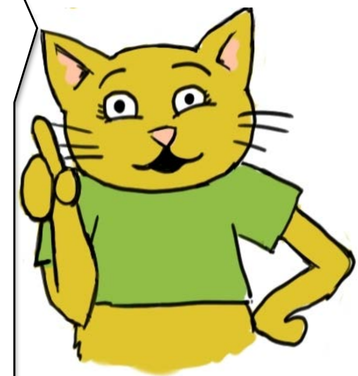

Tom

Cathy	1	2
1	-2	3
2	3	-4

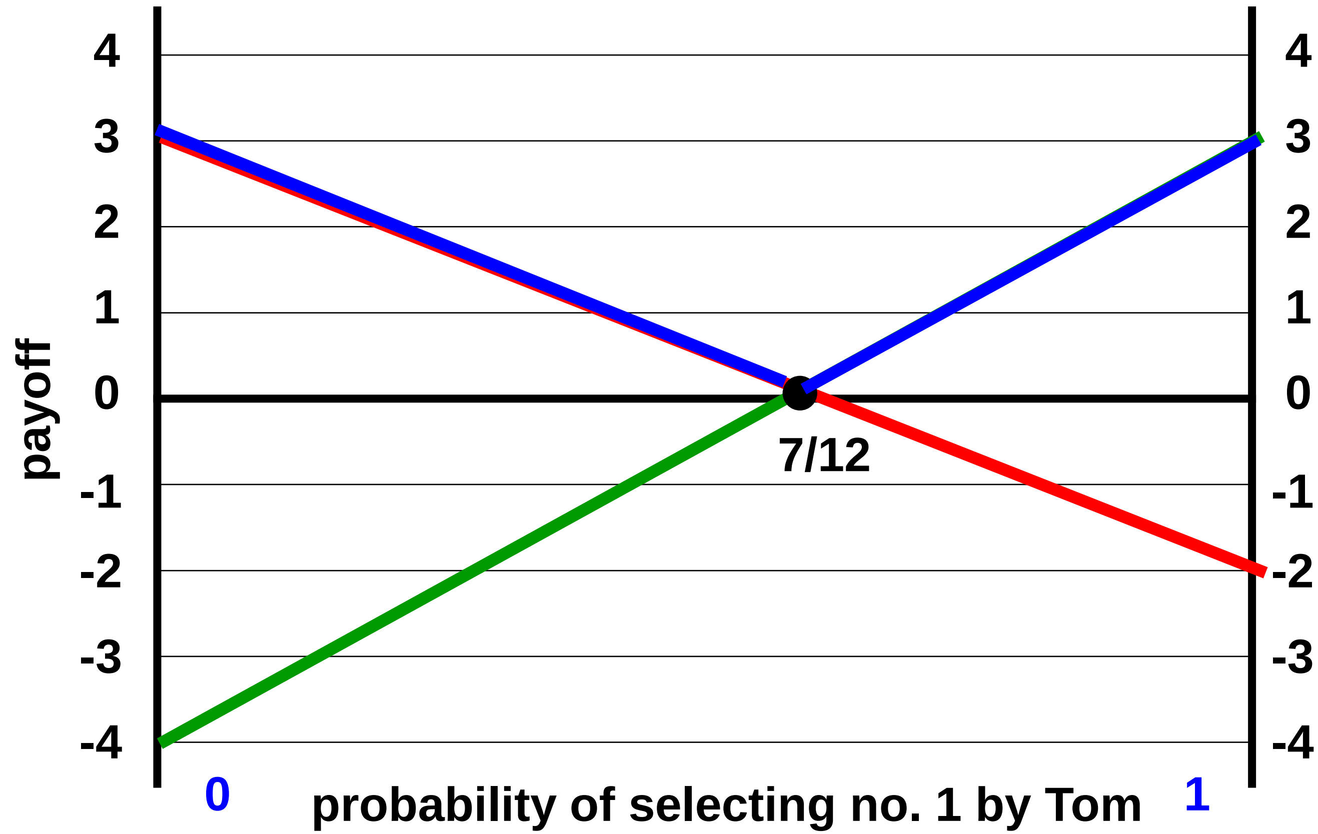

But what's the best that you can do?

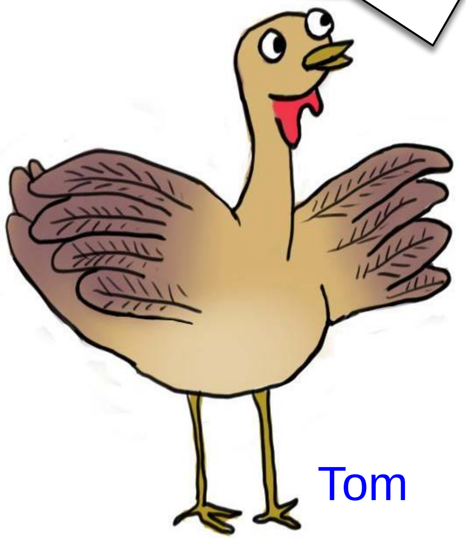

It now my turn to choose an optimal strategy. I'll let q denote the probability that I select Row 1. I need to figure out the best choice of q.

Here is my payoff if you select Column 1.

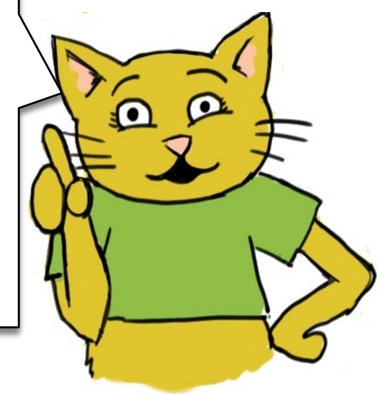

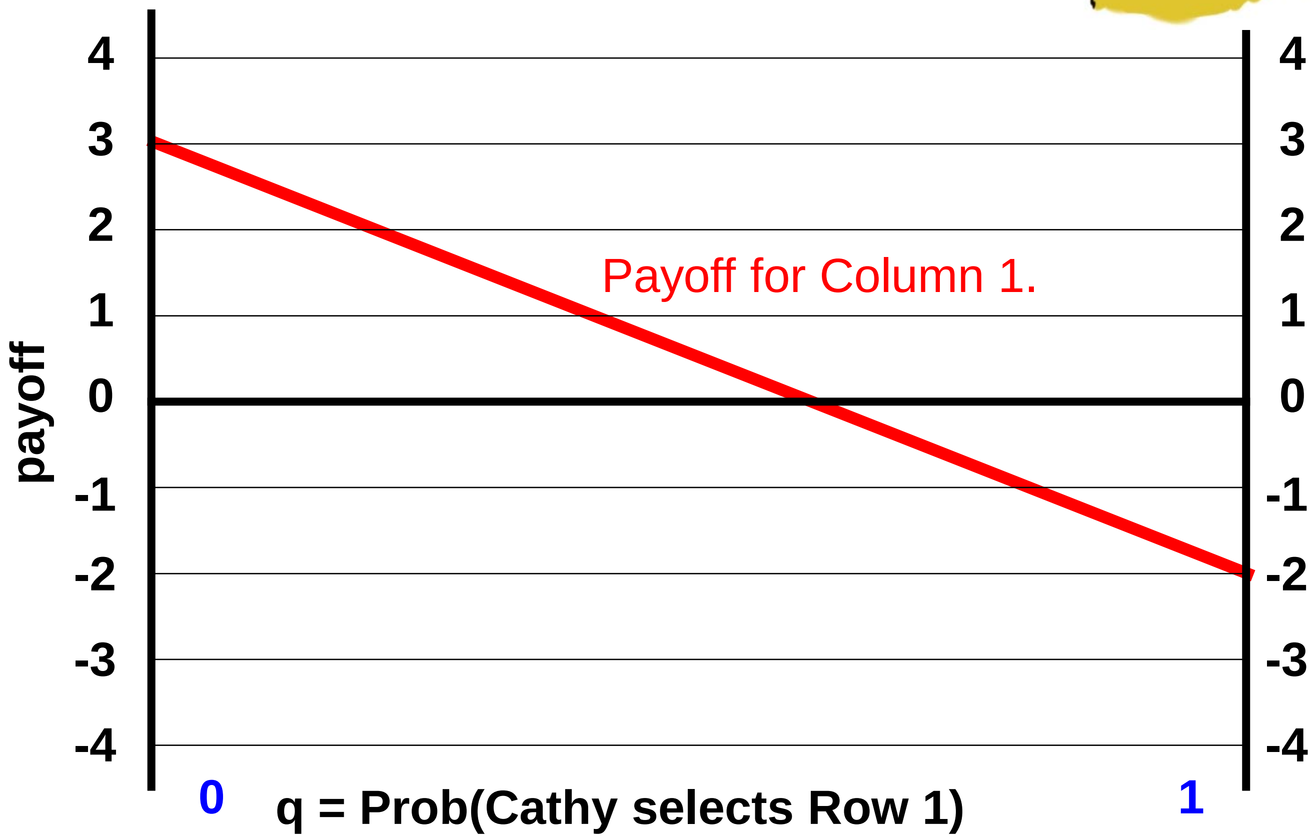

		Tom 1	Tom 2
Cathy	1	-2	3
	2	3	-4

Here is my payoff if you call out 2.

Remember that your strategy is ***minimax***, but mine is ***maximin***.

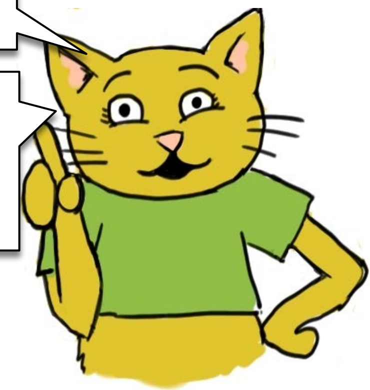

I see! You take a similar approach.

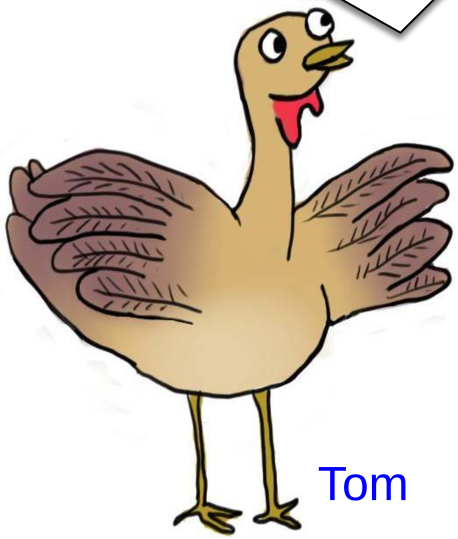

		Tom 1	Tom 2
Cathy	1	-2	3
	2	3	-4

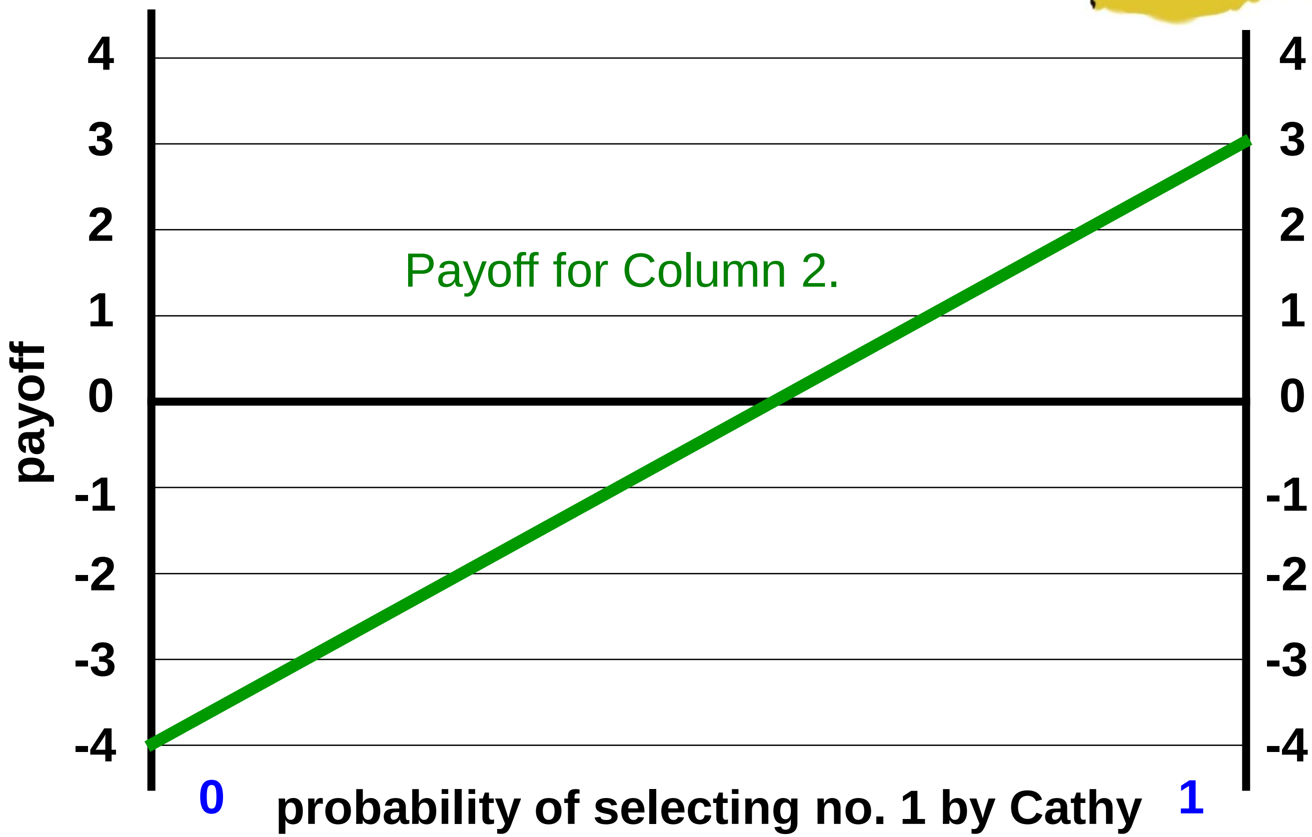

I can tell you now what is your best strategy:

I want to minimize the payoff. So here is my strategy as a function of q:

- if q<7/12, my best response is Column 1;
- If q>7/12, my best response will be Column 2;

Since you want to maximize the payoff, the best strategy is to set q=7/12. This guarantees you can get at least 1/12.

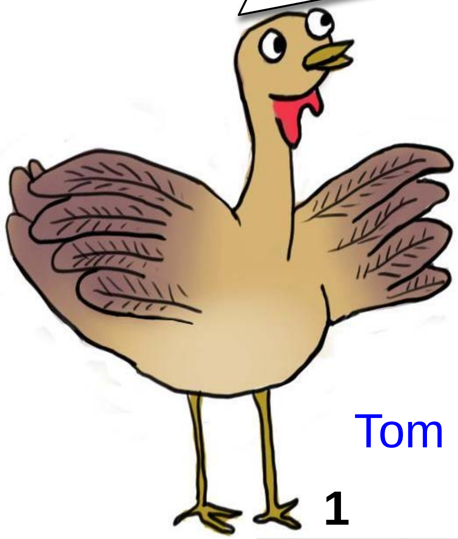

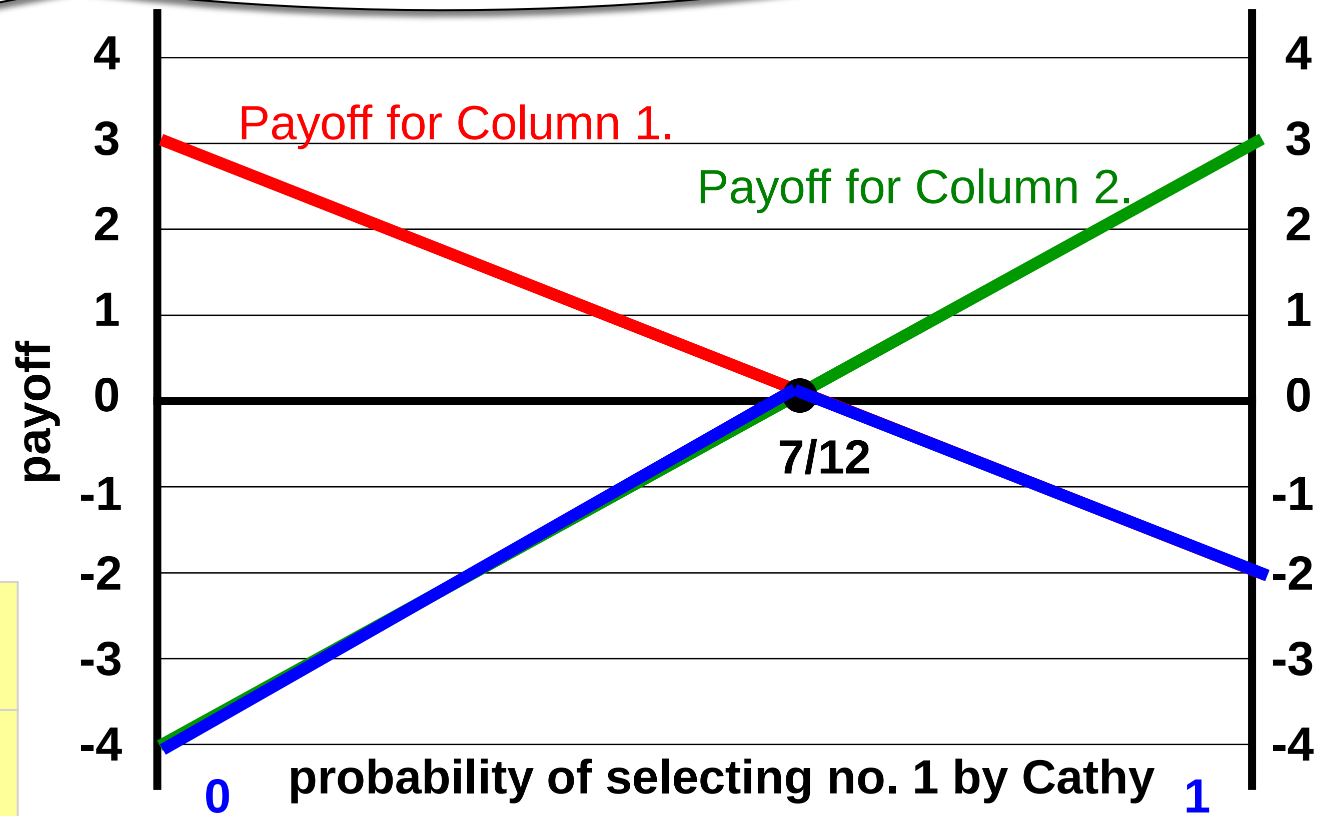

		Tom 1	Tom 2
Cathy	1	-2	3
	2	3	-4

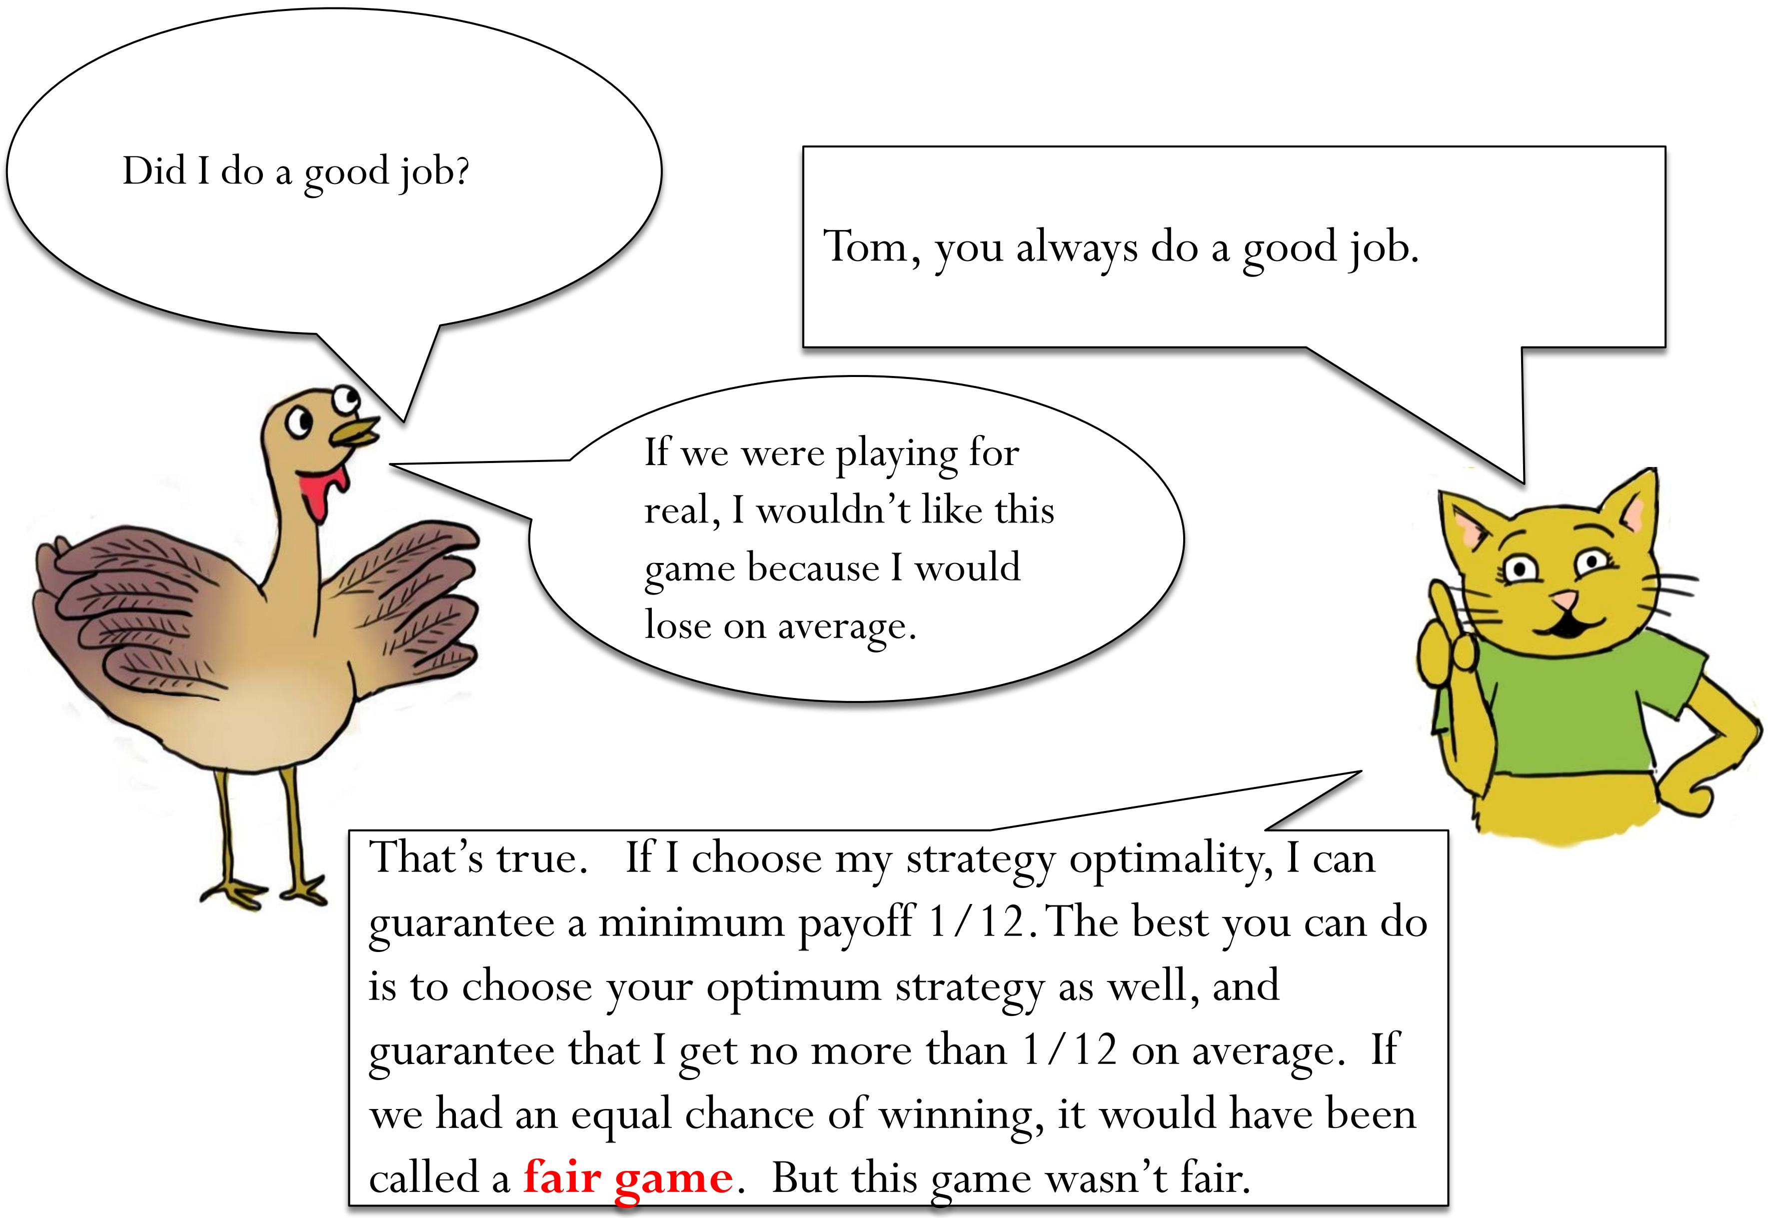
Did I do a good job?
Tom, you always do a good job.
If we were playing for real, I wouldn’t like this game because I would lose on average.
That’s true. If I choose my strategy optimality, I can guarantee a minimum payoff 1/12. The best you can do is to choose your optimum strategy as well, and guarantee that I get no more than 1/12 on average. If we had an equal chance of winning, it would have been called a **fair game**. But this game wasn’t fair.

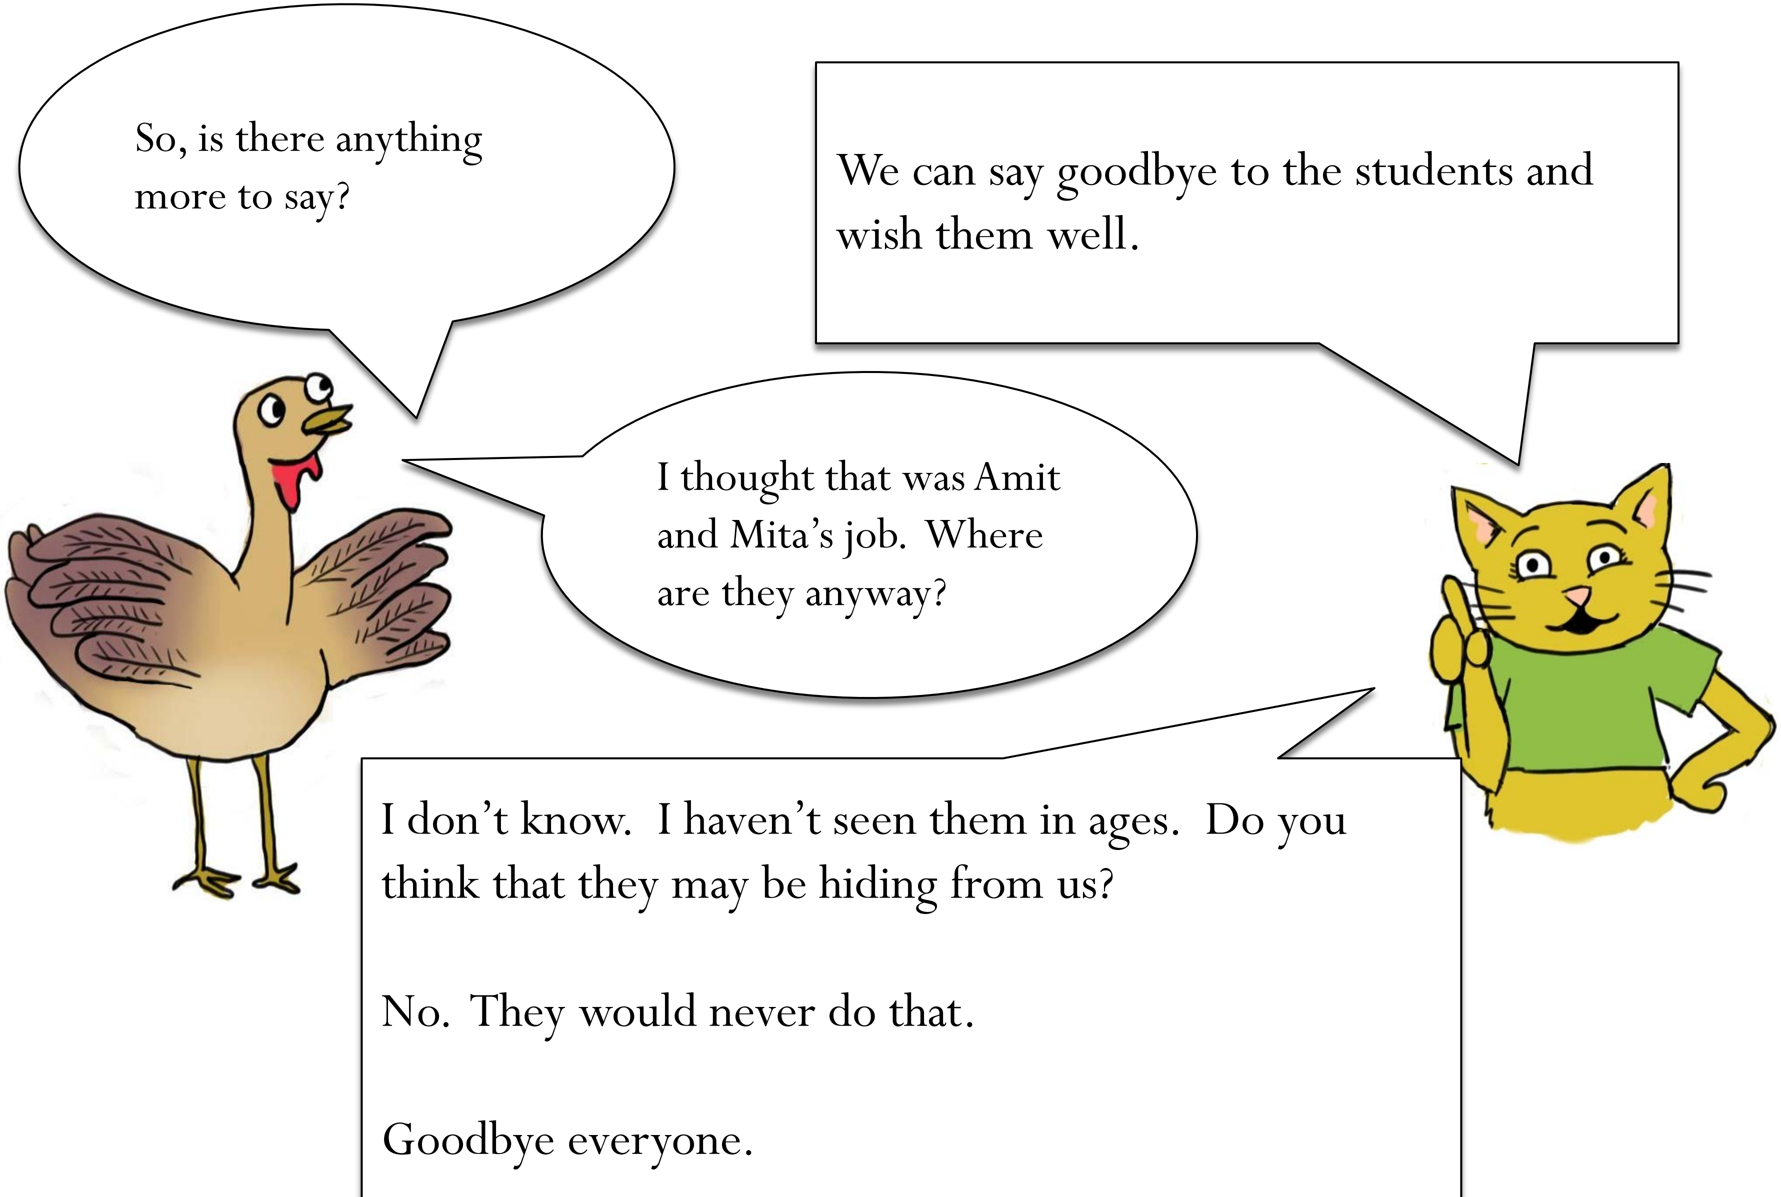
So, is there anything more to say?
We can say goodbye to the students and wish them well.
I thought that was Amit and Mita's job. Where are they anyway?
I don't know. I haven't seen them in ages. Do you think that they may be hiding from us?
No. They would never do that.
Goodbye everyone.

15.053 Optimization Methods in Management Science
Spring 2013