

Best counterstrategy for C

In the previous lecture we saw that if R plays a particular mixed strategy and shows no intention of changing it, the expected payoff for R (and hence C) varies as C varies her strategy. Obviously C wants to choose the strategy that keeps R 's expected payoff at a minimum.

In the next few slides we'll look at this question: for a particular mixed strategy played by R , what is the best possible response by C (C 's *counterstrategy*), i.e., the response that makes the expected payoff to R as small as possible?

As with the minimax method for strictly determined games, this will ultimately lead to a determination of the best strategy for R .

Best counterstrategy for C

We'll start with the earlier example of a zero-sum game with pay-off matrix for R given by

$$\begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix}$$

and let's assume that R plays $[\ .8 \ .2]$. Let's also assume that R is showing no signs of changing his strategy and C is exploring her options. C 's goal is to minimize R 's expected payoff, thus maximizing her own.

Best counterstrategy for C

First let's consider what happens when C plays a pure strategy.

We will start with pure strategy $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ where C **always plays column 1**. The expected payoff for R if C plays column 1 only is

$$[.8 \ .2] \begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = [-.4]$$

On the other hand, if C plays the pure strategy $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (where C **always plays column 2**), the expected payoff for R is

$$[.8 \ .2] \begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = [2].$$

Best counterstrategy for C

Thus, of these two pure strategies, **assuming that R continues to play the strategy $\begin{bmatrix} .8 & .2 \end{bmatrix}$** , the better one for C is the strategy $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$, giving an expected payoff of $-.4$ for R (so a payoff of $.4$ for C).

Suppose C chooses some mixed strategy, $\begin{bmatrix} q \\ 1-q \end{bmatrix}$, $0 \leq q \leq 1$.

Common sense might suggest that the expected payoff for R now falls somewhere *between* $-.4$ and 2 . Indeed, it's

$$\begin{bmatrix} .8 & .2 \end{bmatrix} \begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} q \\ 1-q \end{bmatrix} = [-.4(q) + 2(1-q)]$$

which varies between 2 (when $q = 0$) and $-.4$ (when $q = 1$).

Conclusion: If R continues to play the strategy $\begin{bmatrix} .8 & .2 \end{bmatrix}$, the best counterstrategy for C is a pure strategy; always playing the column that minimizes R 's payoff.

Best counterstrategy for C

In general suppose R plays a mixed strategy $[p, 1 - p]$, with a payoff matrix given by

$$\begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}$$

If C plays a pure strategy $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ or $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ the payoff for R will be

$$\begin{bmatrix} p & (1 - p) \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = [pa_{11} + (1 - p)a_{21}] \quad \text{or}$$

$$\begin{bmatrix} p & (1 - p) \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = [pa_{12} + (1 - p)a_{22}] \quad \text{respectively.}$$

The best counterstrategy for C to R 's mixed strategy $[p, 1 - p]$ is the pure strategy whose expected payoff is the minimum of $pa_{11} + (1 - p)a_{21}$ and $pa_{12} + (1 - p)a_{22}$

Summary

If R always plays the strategy $\begin{bmatrix} p & (1-p) \end{bmatrix}$, the best counterstrategy for C is a pure strategy; always playing the column that **minimizes** R 's expected payoff. Similarly, if C always plays the strategy $\begin{bmatrix} q \\ (1-q) \end{bmatrix}$, the best counterstrategy for R will be a pure strategy, where R always plays the row which *maximizes* R 's expected payoff.

Example

Recall the pay-off matrix for General Roadrunner:

		C. attacks	
		<i>B</i>	<i>S</i>
R. places bomb	<i>B</i>	80%	100%
	<i>S</i>	90%	50%

If General Roadrunner plays a strategy of $[\cdot 3, \cdot 7]$, what counterstrategy should general Coyote play in order to minimize the number of bombs which reach their target?

If Coyote responds with C1, expected payoff for Roadrunner is $(\cdot 3)80 + (\cdot 7)90 = 87$; if Coyote responds with C2, expected payoff for Roadrunner is $(\cdot 3)100 + (\cdot 7)50 = 65$; so Coyote's best counterexample to this mixed strategy of Roadrunner's is to always attack the fighter. In this way he can keep the percentage of bombs that reach their target down to 65%.

Example: Should I have the operation?

Surgeries for some conditions come with such a risk that it might be best not to undertake them if they can be avoided. It may not be easy to diagnose these conditions and often patients know only a probability that they have the condition.

Let us consider a simple example where someone is contemplating a risky surgery for a serious disease, given that their doctor says there is a 50% chance that they have the disease. Here the opponent is Nature and we will make Nature the Column player. The Patient will be the Row player. The payoff is given in years of life expectancy for each of the four situations, where D and ND denote having and not having the disease respectively and S and NS denote a decision to have or not have surgery respectively.

Example: Should I have the operation?

		Nature	
		<i>D</i>	<i>ND</i>
Patient	<i>S</i>	15	30
	<i>NS</i>	2	40

Nature probably is not influenced by the patient's choice to have the surgery or not, so Nature's strategy is $\begin{bmatrix} .5 \\ .5 \end{bmatrix}$ and the patient wants to find the best counterstrategy. In other words, is it better for the patient to have the surgery or not? Calculate the expected payoff (in life expectancy) for R for both decisions.

Strategy S has expected payoff 22.5, and strategy NS has expected payoff 21 — in this case surgery is the better choice.

Optimal mixed strategy for R

We can draw a picture representing the possible payoffs for R — we draw lines representing R 's payoff for each of C 's pure strategies (this payoff will vary as p varies in R 's strategy $[p, 1 - p]$). These lines are called *strategy lines*.

Example: Let's look at the example again where the payoff matrix is given by

$$\begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix}.$$

Let $[p, 1 - p]$ denote R 's strategy. We draw a co-ordinate system with the variable p on the horizontal axis and $y =$ the expected payoff for R on the vertical axis.

Optimal mixed strategy for R



If R plays $[p, 1 - p]$ and C plays $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$, the expected payoff for R is

$$[p \quad (1-p)] \begin{bmatrix} -1 & 3 \\ 2 & -2 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = [2 - 3p]$$

Thus the line showing the expected value for R when C always plays Col. 1 is $y = 2 - 3p$ (shown in blue on the next slide).

Similarly, we see that the line showing the expected value for R when C always plays Col. 2 is $y = 5p - 2$ (shown in red on the next slide).


Optimal mixed strategy for R



Now suppose that R chooses a value p . We assume that C will respond appropriately and choose her best possible counterstrategy. Recall that whatever strategy C chooses, R 's expected value will be in the shaded region on the right. The best counterstrategy for C is to choose the pure strategy that will give the minimum expected payoff for R .

Optimal mixed strategy for R

Hence, if C chooses the best counterstrategy for a particular choice of p made by R , R 's payoff will be minimized and will appear along the line highlighted in green below.



Optimal mixed strategy for R

Now since R wants to maximize his payoff, R chooses the strategy corresponding to the value of p which gives the maximum along the green line. This is the value of p at which the lines meet.

Find the value of p for which the above strategy lines meet and find R 's best mixed strategy.

The red line is $y = 5p - 2$ and the blue line is $y = 2 - 3p$. Hence $5p - 2 = 2 - 3p$ or $8p = 4$ or $p = \frac{1}{2}$. Hence R 's best mixed strategy is to play each option with probability of 50%, for an expected payoff of $1/2$.

Notice that when R makes this choice of p , it doesn't matter what C does ... whatever pure or mixed strategy C chooses, the expected payoff for R will *always* be $1/2$.

Example: Roadrunner & Coyote

Recall the pay-off matrix for General Roadrunner:

		C. attacks	
		<i>B</i>	<i>S</i>
R. places bomb	<i>B</i>	80%	100%
	<i>S</i>	90%	50%


(a) Draw the strategy lines for *R* for this game.

$$[p \quad 1-p] \begin{bmatrix} 80 & 100 \\ 90 & 50 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} = [p80 + (1-p)90] = [90 - 10p]$$

$$[p \quad 1-p] \begin{bmatrix} 80 & 100 \\ 90 & 50 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} = [p100 + (1-p)50] = [50 + 50p].$$

The lines are $y = 90 - 10p$ and $y = 50 + 50p$.

Example: Roadrunner & Coyote



Example: Roadrunner & Coyote

(b) Use the strategy lines above to determine the optimal mixed strategy for General Roadrunner.

The strategy lines intersect when $90 - 10p = 50 + 50p$ so $40 = 60p$ and $p = \frac{2}{3}$. The payoff to Roadrunner is

$$\frac{250}{3} = 90 - 10 \cdot \frac{2}{3} = 50 + 50 \cdot \frac{2}{3} = 83.33\dots$$

Notice that is is higher than the payoff the Roadrunner when he plays his optimum pure strategy. That payoff was 80. This example shows the advantage of moving to a mixed strategy.

If C has more than two options

If C has more than two options, say payoff matrix has dimensions $2 \times n$, the thinking is still the same:

For each possible mixed strategy $[p \ 1 - p]$ for R , C has n possible pure strategy options in response. Each of these has an expected payoff. Any mixed strategy that C might come up with is going to be a mix or (weighted) average of these pure responses, so the expected payoff will be a mix or (weighted) average of the expected payoffs. The average of a bunch of numbers can't be smaller than the smallest of the numbers, so the pure strategy with the smallest expected payoff will be the best counterstrategy for C .

R can again draw strategy lines (now n of them), identify the lower boundary of these lines, and identify the value of p at the point where the lower boundary is as high as possible. This will give R his optimum mixed strategy.

Example

In tennis the server can serve to the forehand or serve to the backhand. The opponent can make a guess as to the type of serve and prepare for that serve, or not guess at all. The payoff matrix for two players, Roger and Caroline, shown below shows the percentage of points ultimately won by the server in each situation.

		Caroline		
		Guess Forehand	Guess Backhand	No Guess
Roger	Serve To Forehand	40	70	45
	Serve To Backhand	80	60	65

(a) Does this payoff matrix have a saddle point?

The maximum of the row minima is 60. The minimum of the column maxima is 65. Since $60 \neq 65$, there is no saddle point.

Example


(b) Plot Roger's the three strategy lines and highlight the lowest path.

$$[p \quad 1-p] \begin{bmatrix} 40 & 70 & 45 \\ 80 & 60 & 65 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} = [80 - 40p].$$

$$[p \quad 1-p] \begin{bmatrix} 40 & 70 & 45 \\ 80 & 60 & 65 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = [60 + 10p].$$

$$[p \quad 1-p] \begin{bmatrix} 40 & 70 & 45 \\ 80 & 60 & 65 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = [65 - 20p].$$

Example



The lowest path follows the red line from the y axis, since $p \geq 0$, until it crosses the black line. Then it follows the black line until it crosses the green line and then it follows the green line until it crosses the gray line $p = 1$.

Example

(c) Determine the optimal mixed strategy for Roger for this game.

Roger's optimal strategy occurs at the value of p where the lowest path is as high as possible. This occurs where the red and black lines intersect. These occur where

$65 - 20p = 60 + 10p$ or $5 = 30p$ or $p = \frac{1}{6}$. Roger should

serve to Caroline's forehand $\frac{1}{6}$ of the time and to his

backhand $\frac{5}{6}$ of the time.

Optimal mixed strategy for C

In the above examples we have found R 's best mixed strategy keeping in mind that C will respond with the optimal counterstrategy. We saw that the analysis led to a solution where the payoff for R was the same no matter what strategy C chooses.

The dynamic however does not necessarily end there. If C is not playing optimally and R can increase his/her payoff by responding with their best counterstrategy, we can assume that they will do so. So in order to find an equilibrium, we must also find the optimal strategy for C , which is a similar problem to that of finding the optimal strategy for R .

Optimal mixed strategy for C

If C has only two options, and the payoff matrix does not have a saddle point, then we can determine C 's optimal mixed strategy in a manner similar to how we found R 's optimal strategy.

Namely, if C 's strategy is denoted by $\begin{bmatrix} q \\ 1-q \end{bmatrix}$, we plot strategy lines in a Cartesian plane with horizontal axis q and vertical axis y , where each line determines the expected payoff for R , if R plays a particular pure counterstrategy. We then highlight the *upper* boundary, which corresponds to the *maximum* payoff for R . C will then choose the value of q that makes this maximum payoff as small as possible; i.e., she we find the place where the upper boundary line is as *low* as possible.

Example: Roadrunner & Coyote


Let us determine General Coyote's optimal strategy. The payoff matrix for General Roadrunner is $\begin{bmatrix} 80 & 100 \\ 90 & 50 \end{bmatrix}$. If Coyote plays strategy $\begin{bmatrix} q \\ 1-q \end{bmatrix}$, we need to work out the payoff for Roadrunner's two pure strategies:

$$\begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} 80 & 100 \\ 90 & 50 \end{bmatrix} \begin{bmatrix} q \\ 1-q \end{bmatrix} = [100 - 20q]$$

$$\begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} 80 & 100 \\ 90 & 50 \end{bmatrix} \begin{bmatrix} q \\ 1-q \end{bmatrix} = [50 + 40q]$$

The next slide shows Coyote's strategy lines. The red line is Roadrunner's payoff if he plays strategy R1 and the blue line is his payoff if Roadrunner plays R2.

Example: Roadrunner & Coyote



The intersection point is at $100 - 20q = 50 + 40q$ or
 $50 = 60q$ or $q = \frac{5}{6}$.

Example: Roadrunner & Coyote

If Coyote plays a mixed strategy $\begin{bmatrix} q \\ 1-q \end{bmatrix}$ Roadrunner's payoff will be somewhere in the shaded region over that value of q . Hence Roadrunner can always play so as to force his payoff to be along the red line for $0 \leq q \leq \frac{5}{6}$ and along the blue line if Coyote plays a value of q between $\frac{5}{6}$ and 1.

Remember that Roadrunner wants big payoffs so the higher up the y -axis the point is, the better.

Coyote on the other hand wants the payoff to be as small as possible. Hence Coyote will play $\begin{bmatrix} \frac{5}{6} \\ \frac{1}{6} \end{bmatrix}$ since anything else will allow Roadrunner to increase his payoff.

Roadrunner's payoff if Coyote plays his optimal mixed strategy is

$$50 + 40 \cdot \frac{5}{6} = 100 - 20 \cdot \frac{5}{6} = 80.33\dots$$

Summary of Roadrunner versus Coyote

We've seen that Roadrunner's optimal mixed strategy is $\left[\frac{2}{3} \quad \frac{1}{3}\right]$, with expected payoff $250/3$.

- ▶ If Roadrunner plays any strategy other than his optimal one, Coyote can, by appropriate choice of counterstrategy, force him to accept a payoff of less than $250/3$.

Coyote's optimal mixed strategy is $\left[\frac{5}{6} \quad \frac{1}{6}\right]$, also with expected payoff $250/3$.

- ▶ If Coyote plays any strategy other than his optimal one, Roadrunner can, by appropriate choice of counterstrategy, force a payoff greater than $250/3$.

Neither player has any incentive to move away from their optimum mixed strategy; the game is stable, or in equilibrium, when both players play their optimum mixed strategy.

The *value* of the game is defined to be the common expected payoff, $250/3$.

Example — hard Irish names

Ruaidrí and Caoimhín play a game in which they try to pronounce each other's names. If they both do so correctly, Ruaidrí gets 6 euro from Caoimhín. If they both do so incorrectly, he gets 4 euro. If Ruaidrí messes up and Caoimhín doesn't, Caoimhín gets 3 euro from Ruaidrí, and if Caoimhín messes up but Ruaidrí doesn't, Caoimhín gets 8 euro.

Here's the payoff matrix for Ruaidrí, with G the strategy of making a good pronunciation and B the strategy of making a bad pronunciation:

		Caoimhín	
		G	B
Ruaidrí	G	6	-8
	B	-3	4

Example — hard Irish names

		Caoimhín	
		<i>G</i>	<i>B</i>
Ruaidrí	<i>G</i>	6	-8
	<i>B</i>	-3	4

The row minima are -8 and -3 with -3 the biggest, so Ruaidrí's best pure strategy is B with payoff -3 .

The column maxima are 6 and 4 with 4 the smallest, so Caoimhín's best pure strategy is B with payoff 4

The game has no saddle point.

If Ruaidrí plays the strategy $[p \ 1 - p]$ the strategy lines are $y = 6p - 3(1 - p) = 9p - 3$ and $y = -8p + 4(1 - p) = 4 - 12p$. These intersect when $9p - 3 = 4 - 12p$ or $p = 1/3$, and which point $y = 0$. So Ruaidrí's best mixed strategy is $[1/3 \ 2/3]$ with expected payoff 0.

Example — hard Irish names

		Caoimhín	
		G	B
Ruaidrí	G	6	-8
	B	-3	4

If Caoimhín plays the strategy $\begin{bmatrix} q \\ 1-q \end{bmatrix}$ the strategy lines are $y = 6q - 8(1 - q) = 14q - 8$ and $y = -3q + 4(1 - q) = 4 - 7q$. These intersect when $14q - 8 = 4 - 7q$ or $q = 12/21 = 4/7$, and which point $y = 0$. So Caoimhín's best mixed strategy is $\begin{bmatrix} 4/7 \\ 3/7 \end{bmatrix}$ with expected payoff 0.

As with Roadrunner and Coyote, the expected payoff to both players at their optimum mixed strategy is the same, and so the game is stable — neither player has an incentive to move away from their optimum mixed strategy.

The value of the game, the common expected payoff, is 0. Games with value 0 are called *fair* since neither player has a long-run advantage.

Games with more than two options

What's going on should “feel like” the linear optimization that we have done. To find his optimal mixed strategy, R has to maximize a linear function — y — inside a region bounded by straight lines — the strategy lines and the lines $p = 0$ and $p = 1$, — and C has to minimize y inside a similarly defined region.

If R has more than two options to choose from, things get more complicated. If he has three options then a typical mixed strategy is $[p \ q \ r]$ with $p, q, r \geq 0$ and $p + q + r = 1$. For each of C's strategies there is now a strategy *plane*, and to find his optimal mixed strategy R has to find the highest point inside some 3-dimensional region bounded by all the strategy planes. C has to do something similar.

The minimax theorem

In 1928 John von Neumann used linear programming to prove that for every two-person zero-sum game with finitely many options for each player, there is an optimal mixed strategy for R with expected payoff some number v (the *value* of the game), and there is an optimal mixed strategy for C with expected payoff *the same number* v

It follows that all games have a stable equilibrium — R has no incentive to move away from his optimal mixed strategy, since if he did, C could engineer it that R gets an expected payoff less than v , and C has no incentive to move away from her optimal mixed strategy, since if she did, R could engineer it that he (R) gets an expected payoff more than v .

Old exam questions

1: Rocky and Creed play a zero-sum game. The payoff matrix for Rocky is given by: $\begin{bmatrix} 3 & 1 \\ 2 & 4 \end{bmatrix}$. If Rocky plays the mixed strategy $(.6 \ .4)$, which of the following mixed strategies should Creed play to maximize his (Creed's) expected payoff in the game? (a) $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$ (b) $\begin{bmatrix} .6 \\ .4 \end{bmatrix}$ (c) $\begin{bmatrix} .4 \\ .6 \end{bmatrix}$ (d) $\begin{bmatrix} .3 \\ .7 \end{bmatrix}$ (e) $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$

Creed should always play a pure strategy, hence either $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ or $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. Since $[.6 \ .4] \begin{bmatrix} 3 & 1 \\ 2 & 4 \end{bmatrix} = [1.8 + .8 \ .6 + 1.6] = [2.6 \ 2.2]$, Creed's best pure is strategy $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$. Remember the payoff matrix is for R 's payoff so Creed wants this number to be as small as possible.

Old exam questions

2: Rose (R) and Colm (C) play a zero-sum game. The payoff matrix for Rose is given by:

$$\begin{bmatrix} 2 & -1 \\ 1 & 3 \end{bmatrix}.$$

Which of the following give the strategy lines corresponding to the fixed strategies of Colm where Rose's strategy is given by $[p \ (1 - p)]$ and Rose's expected payoff is denoted by y ?

(a) $y = p + 1$
 $y = 3 - 4p$ (b) $y = 3p - 1$
 $y = 3 - 2p$ (c) $y = 2p + 1$
 $y = 3p - 1$

(d) $y = 2 - p$
 $y = 4p - 1$ (e) $y = 2 - 3p$
 $y = 4p + 1$

Old exam questions

Since $[p \ 1 - p] \begin{bmatrix} 2 & -1 \\ 1 & 3 \end{bmatrix} = [(2p + 1 - p) \ (-p + 3 - 3p)] = [(p + 1) \ (3 - 4p)]$. Hence the strategy lines corresponding to the fixed strategies of Colm are

$$y = p + 1$$

$$y = 3 - 4p$$