


Game theory (Ch. 17.5)



Find best strategy

How do we formally find a Nash equilibrium?

If it is zero-sum game, can use minimax as neither player wants to switch for Nash (our PD example was not zero sum)

Let's play a simple number game: two players write down either 1 or 0 then show each other. If the sum is odd, player one wins. Otherwise, player 2 wins (on even sum)

Find best strategy

This gives the following payoffs:

		Player 2	
		Pick 0	Pick 1
Player 1	Pick 0	-1, 1	1, -1
	Pick 1	1, -1	-1, 1


(player 1's value first, then player 2's value)

We will run minimax on this tree twice:

1. Once with player 1 knowing player 2's move (i.e. choosing after them)
2. Once with player 2 knowing player 1's move

Find best strategy


Player 1 to go first (max):



If player 1 goes first, it will always lose

Find best strategy

Player 2 to go first (min):



If player 2 goes first, it will always lose

Find best strategy

This is not useful, and only really tells us that the best strategy is between -1 and 1 (which is fairly obvious)

This minimax strategy can only find pure strategies (i.e. you should play a single move 100% of the time)

To find a mixed strategy, we need to turn to linear programming

Find best strategy

A pure strategy is one where a player always picks the same strategy (deterministic)

A mixed strategy is when a player chooses actions probabilistically from a fixed probability distribution (i.e. the percent of time they pick an action is fixed)

If one strategy is better or equal to all others across all responses, it is a dominant strategy

Find best strategy

The definition of a Nash equilibrium is when no one has an incentive to change the combined strategy between all players

So we will only consider our opponent's rewards (and not consider our own)

This is a bit weird since we are not considering our own rewards at all, which is why the Nash equilibrium is sometimes criticized

Find best strategy

First we parameterize this and make the tree stochastic:

Player 1 will choose action “0” with probability p , and action “1” with $(1-p)$

If player 2 always picks 0, so the payoff for p2:

$$(1)p + (-1)(1-p)$$

If player 2 always picks 1, so the payoff for p2:

$$(-1)p + (1)(1-p)$$

Find best strategy


Plot these two lines:

$$U = (1)p + (-1)(1-p)$$

$$U = (-1)p + (1)(1-p)$$

As we maximize, the opponent gets to pick which line to play

Thus we choose the intersection



Find best strategy

Thus we find that our best strategy is to play 0 half the time and 1 the other half

The result is we win as much as we lose on average, and the overall game result is 0

Player 2 can find their strategy in this method as well, and will get the same 50/50 strategy (this is not always the case that both players play the same for Nash)

Find best strategy

We have two actions, so one parameter (p) and thus we look for the intersections of lines

If we had 3 actions (rock-paper-scissors), we would have 2 parameters and look for the intersection of 3 planes (2D)

This can generalize to any number of actions (but not a lot of fun)

		Player 2		
		Stone	Paper	Scissors
Player 1	Stone	(0, 0)	(-1, 1)	(1, -1)
	Paper	(1, -1)	(0, 0)	(-1, 1)
	Scissors	(-1, 1)	(1, -1)	(0, 0)

Find best strategy

How does this compare on PD?


	Confess	Lie
Confess	-8, -8	0, -10
Lie	-10, 0	-1, -1

Player 1: p = prob confess...

P2 Confesses: $-8 * p + 0 * (1 - p)$

P2 Lies: $-10 * p + (-1) * (1 - p)$

Cross at negative p , but red line is better (confess)



Find best strategy

In cases like this where you get wonky probabilities, it means there is a dominant strategy


You should then remove the dominated action and recalc.

	Confess	Lie
Confess	-8, -8	0, -10
Lie	-10, 0	-1, -1

never play

bottom row

	Confess	Lie
Confess	-8, -8	0, -10




Chicken

What is Nash for this game?
What is Pareto optimum?

	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0

Game of Chicken



Chicken

To find Nash, assume we (blue) play S probability p , C prob $1-p$

	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0

Column 1 (red=S): $p*(-10) + (1-p)*(1)$

Column 2 (red=C): $p*(-1) + (1-p)*(0)$

Intersection: $-11*p + 1 = -p, p = 1/10$

Conclusion: should always go straight $1/10$ and chicken $9/10$ the time

Chicken

We can see that 10% straight makes the opponent not care what strategy they use:

	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0

(Red numbers)

$$100\% \text{ straight: } (1/10)*(-10) + (9/10)*(1) = -0.1$$

$$100\% \text{ chicken: } (1/10)*(-1) + (9/10)*(0) = -0.1$$

$$\begin{aligned} 50\% \text{ straight: } & (0.5)*[(1/10)*(-10) + (9/10)*(1)] \\ & + (0.5)*[(1/10)*(-1) + (9/10)*(0)] \\ & = (0.5)*[-0.1] + (0.5)*[-0.1] = -0.1 \end{aligned}$$

Chicken

The opponent does not care about action, but you still do (never considered our values)



	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0

Your rewards, opponent 100% straight:

$$(0.1)*(-10) + (0.9)*(-1) = -1.9$$

Your rewards, opponent 100% curve:

$$(0.1)*(1) + (0.9)*(0) = 0.1$$

The opponent also needs to play at your value intersection to achieve Nash

Chicken

Pareto optimum?


All points except $(-10, 10)$

Going off the definition,
P1 loses point if move
off $(1, -1)$

... similar P2 on $(-1, 1)$

At $(0, 0)$ there is no point
with both vals positive

	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0




Chicken

We can define a mixed strategy
Pareto optimal points

	S	C
S	-10, -10	1, -1
C	-1, 1	0, 0

Can think about this
as taking a string from the
top right and bringing the
it down & left

Stop when string going
straight left and down



Find best strategy

We have two actions, so one parameter (p) and thus we look for the intersections of lines

If we had 3 actions (rock-paper-scissors), we would have 2 parameters and look for the intersection of 3 planes (2D)

This can generalize to any number of actions (but not a lot of fun)

		Player 2		
		Stone	Paper	Scissors
Player 1	Stone	(0, 0)	(-1, 1)	(1, -1)
	Paper	(1, -1)	(0, 0)	(-1, 1)
	Scissors	(-1, 1)	(1, -1)	(0, 0)

Repeated games

In repeated games, things are complicated

For example, in the basic PD, there is no benefit to “lying”

		PRISONER 2	
		Confess	Lie
PRISONER 1	Confess	-8, -8	0, -10
	Lie	-10, 0	-1, -1

However, if you play this game multiple times, it would be beneficial to try and cooperate and stay in the [lie, lie] strategy

Repeated games

One way to do this is the tit-for-tat strategy:

1. Play a cooperative move first turn
2. Play the type of move the opponent last played every turn after (i.e. answer competitive moves with a competitive one)

This ensure that no strategy can “take advantage” of this and it is able to reach cooperative outcomes

Repeated games

Two “hard” topics (if you are interested) are:

1. We have been talking about how to find best responses, but it is very hard to take advantage if an opponent is playing a sub-optimal strategy
2. How to “learn” or “convince” the opponent to play cooperatively if there is an option that benefits both (yet dominated)

Repeated game


In the example from earlier... the Nash would be to play (1,1)

3,3	0,4
3,0	1,1

But, if the player cooperate, they could both achieve better results

Specifically, if player 1 flips a coin between top and bottom and player 2 chooses left ... this will average to (3, 1.5) value for them

Repeated games



<http://ncase.me/trust/>