Game Theory

2.4 Cooperative Games & Bargaining

George Mason University, Spring 2018

In a *non-cooperative game*, each player acts alone and the outcome depends on their individual actions. They are unable to communicate and have no means of enforcing cooperative agreements.

In a *cooperative game*, the outcome depends on the joint action of a coalition (we'll be concerned with the case of coalitions of two in two-player games).

This action is coordinated via communication, enforcement of cooperative agreements, and bargaining.

Recall:
Ex. Bach or Stravinsky? (Osborne and Rubinstein)

Two people wish to go out together to a concert of music by either Bach or Stravinsky. Their main concern is to go out together, but one person prefers Bach and the other person prefers Stravinsky. Suppose that have expressed their wishes to one another, but have not been able to coordinate an outcome (perhaps someone's phone battery has died). What should they do?

This is a variant on the classic Battle of the Sexes (Luce and Raiffa), sometimes also called the Clash of Wills.

	Bach	Stravinsky
Bach	3,2	1,1
Stravinsky	0,0	2,3

This game has two Nash equilibria but no way to reach them if each acts in an individually rational way.

$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = 1$

$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = 1$

Playing maximinimizers will result in the non-equilibrium $\langle \text{Bach}, \text{Stravinsky} \rangle$. Individual rationality leads to an unstable solution, i.e., one that is not a Nash equilibrium.

Now, let's consider what happens if Row and Col can communicate.

Initially, it might seem they're no better off. They won't be able to agree on a strategy that reaches a Nash equilibrium because each will stand her ground in an attempt to maximize her own utility. We'll have a clash of wills.

But maybe there's a way to get around this. Let's start by considering what happens when the players play mixed strategies.

Recall that the EUs for Row of the minimax pure strategy is:
$EU_{Row} = 1$, $EU_{Col} = 1$

Consider what happens when the players play a *mixed strategy*, but don't yet coordinate. Each, knowing that playing their maximinimizer will lead to an unstable outcome, decides to flip a coin to decide what to do:

$$\langle\langle \text{Bach}[\tfrac{1}{2}], \text{Stravinsky}[\tfrac{1}{2}]\rangle, \langle \text{Bach}[\tfrac{1}{2}], \text{Stravinsky}[\tfrac{1}{2}]\rangle\rangle$$

$$EU_{Row} = \tfrac{1}{4}(3) + \tfrac{1}{4}(1) + \tfrac{1}{4}(0) + \tfrac{1}{4}(2) = \tfrac{3}{2}$$

$$EU_{Col} = \tfrac{1}{4}(2) + \tfrac{1}{4}(1) + \tfrac{1}{4}(0) + \tfrac{1}{4}(3) = \tfrac{3}{2}$$

They do better this way, but they can do even better if they are able to communicate.

They could, for example, decide to allow a single coin to determine what each of them does. Heads, they see Bach; tails, they see Stravinksy. This would be to adopt a *coordinated strategy*:

$\langle [\frac{1}{2}]\langle \text{Bach}, \text{Bach}\rangle, [\frac{1}{2}]\langle \text{Stravinsky}, \text{Stravinsky}\rangle\rangle$

$EU_{Row} = \frac{1}{2}(3) + \frac{1}{2}(2) = \frac{5}{2}$

$EU_{Col} = \frac{1}{2}(2) + \frac{1}{2}(3) = \frac{5}{2}$

Now, this is a better outcome.

More generally, if Row and Col play a coordinated strategy pair in this game:

$\langle [x]\langle R_1, C_1\rangle, [1-x]\langle R_2, C_2\rangle\rangle$

where $0 < x < 1$, their utilities are

$EU_{Row} = x(2) + (1-x)3 = 3 - x$

$EU_{Col} = x(3) + (1-x)2 = 2 + x$

We can visualize this on a plane where the x-axis is Row's EU and the y-axis is Col's. When x=0, their coordinated strategy gives a utility of (2,1), when x=1 it gives a utility of (1,2). If x is anything in between, it will give a set of utilities that falls somewhere on this line.

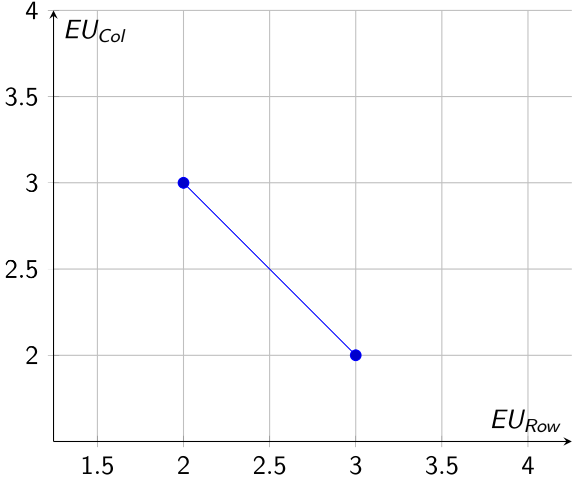

Row and Col can do no better than playing one of these coordinated strategies. Each point along this line is *Pareto optimal*, i.e., at each point, neither player can be made better off without making the other player worse off.

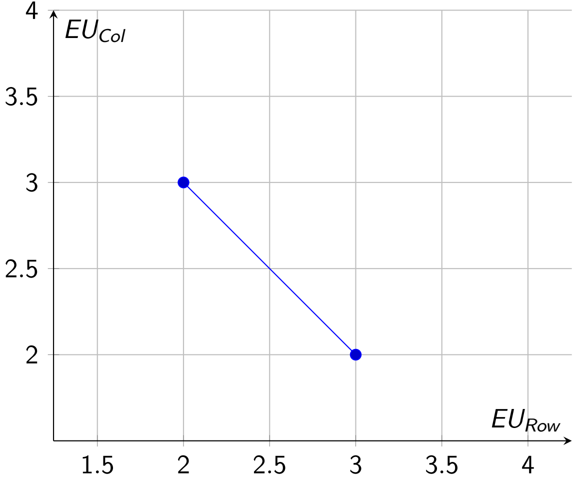

The question, then, is which one they should play. Which one is the solution to the game?

In a cooperative game, players can usually achieve outcomes with coordinated strategies that are better than the outcomes that they can achieve without coordination.

They can also *avoid* outcomes that are worse for both of them (like (0,0) and (1,1) in our game).

We call the full set of outcomes that can be reached by either coordinated or individual strategies the *achievable set*.

Achievable Set for Bach/Stravinsky Game

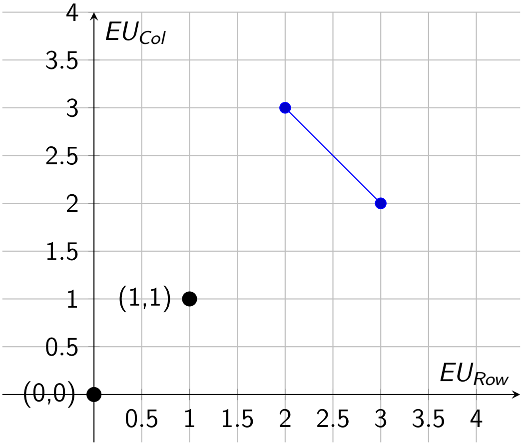

The Pareto optimal members of the achievable set are all those points along the blue line.

One possible response:

All the Pareto optimal members of the achievable set are solutions to the game.

Problem 1. This doesn't distinguish between what is intuitively a fair outcome of our game and those that are biased toward Row or Col. Similarly, in a cooperative version of Prisoner's Dilemma, the outcomes of one cooperating and one defecting as well as the outcome of both cooperating are all Pareto optimal, so this method lumps them together.

Problem 2. There are many Pareto optimal outcomes that yield many different values to Row and Col. There is no solution to the coordination problem this generates, no single best outcome for both players (as in a zero-sum game), so we'd find ourselves in a competitive game win which players vie for Pareto optimal outcomes favorable to themselves.

One solution is to allow players to *bargain*.

Bargaining Games.

What bargaining is not:

- Threats or violence
- Giving up the game
- Offering additional benefits

What bargaining is:

- Negotiating to one of the achievable Pareto optimal outcomes
- Since it's more rational to end disputes by negotiation than to leave them festering or to resolve them by force, bargaining games merit our attention.

Bargaining Games.

A bargaining game can be specified in terms of the *achievable set* of outcomes of a cooperative game.

The achievable set must contain at least *two* Pareto optimal outcomes, or there would be nothing to contest.

It also contains at least one non-Pareto optimal point. We call this the *failure point*. This is what the players end up with if they fail to cooperate.

The space defined by these points in a plane is the *bargaining region* of the game.

Finally, the *negotiation point* is the point within the bargaining region that would be selected by rational bargainers.

Bargaining Region for Bach/Stravinsky

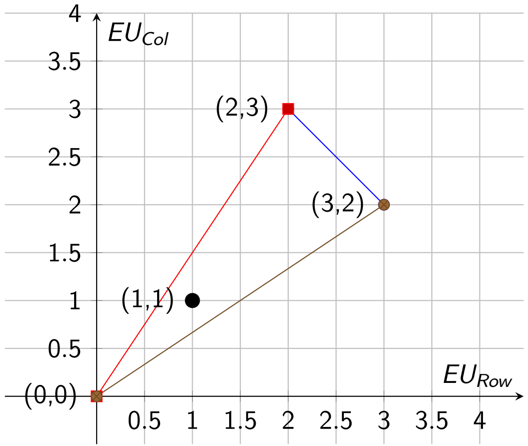

So, what negotiation point would a rational bargainer choose? Where would we be happy to find we ended up?

Nash Point.
Let the failure point be designated (f_1, f_2), i.e., the point where Row gets f_1 and Col gets f_2.

Let (u_1, u_2) designate one of the Pareto optimal points in the bargaining region R of a cooperative game.

Now consider the numbers $u_1 - f_1$ and $u_2 - f_2$. These are the gains that Row and Col will realize if they can manage cooperation and agree on (u_1, u_2). Rational players will want to maximize this, but which Pareto optimal point will they choose?

John Nash proposed taking the negotiation point to be the point within the bargaining region R at which the product

$(u_1 - f_1)(u_2 - f_2)$

is maximized.

Bargaining Region for Bach/Stravinsky

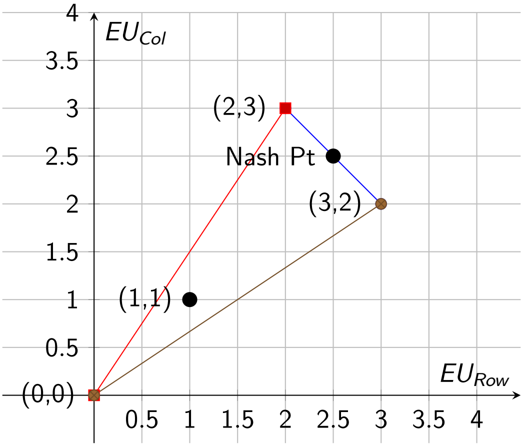

$(2.5-0)(2.5-0) = 6.25$ This is the max of the product for any Pareto optimal point. (Consider $(2.4)(2.6) = 6.24$.)

If we plug in $2\frac{1}{2}$ as the utility for Row and Col in our equations from earlier and solve for x:

$\frac{5}{2} = 3 - x$
$x = \frac{1}{2}$

$\frac{5}{2} = 2 + x$
$x = \frac{1}{2}$

So, $\langle[\frac{1}{2}]\langle \text{Bach}, \text{Bach}\rangle, [\frac{1}{2}]\langle \text{Stravinsky}, \text{Stravinsky}\rangle\rangle$ is the Nash point strategy.

It can be shown that the Nash point is a *unique* Pareto optimal point of the game, but this is insufficient reason for thinking it is a solution to the game.

Nash proposed four axioms that a solution to bargaining problems must satisfy and showed that his solution concept is the only one that can satisfy all four.

1. *Pareto optimality*: the negotiation point must be a unique Pareto optimal point.
2. *Invariance under utility transformations*: If two bargaining games can be obtained from each other by positive linear transformations of the players' utility functions, their negotiation points can be obtained from each other by the same transformations.
3. *Symmetry*: The players' strategies will be identical if and only if their utility functions are identical, i.e., if the game is symmetrical.
4. *Irrelevant expansion/contraction*: If (a) the bargain region of one game is contained within that of another and (b) the two games have the same failure points, then if the negotiation point for the game with the larger region happens to fall within the smaller region, the two games must have the same negotiation point.

We'll concern ourselves only with irrelevant expansion/contraction. This is the most controversial axiom. The reason is that it seems that expanding or contracting the bargaining space (maintaining the same failure point and keeping the original negotiation point in the space) can affect what we would intuitively think the negotiation point should be.

Trapezoid vs. Triangle Bargaining Space

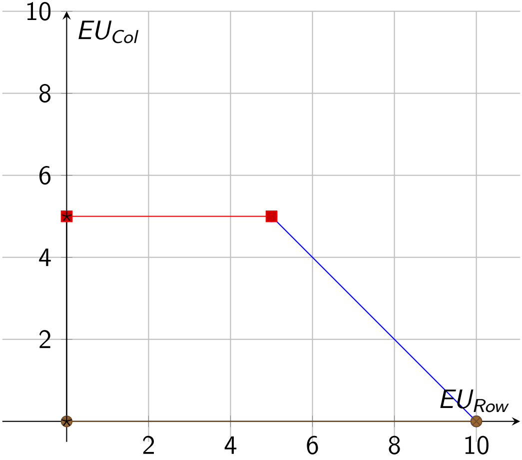

Nash point is (5,5). Verify this.

Trapezoid vs. Triangle Bargaining Space

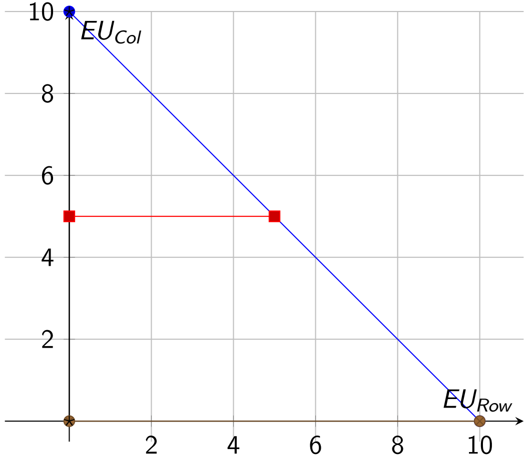

Nash point remains (5,5). But should it? Or, better, should it have been (5,5) in the trapezoid game?

One reasonable thought is that we should think about potential gains or concessions, as they matter to bargaining. In the trapezoid game, Row is conceding 5 utiles, while Col is receiving the max of what he could receive. Why would Row agree to this?

This has led theorists to propose solutions to bargaining games in terms of proportional gains and proportional concessions. The *equitable distribution point* is where the two players proportional gains (concessions) are equal.

Proportional Gains and Concessions.

To calculate a proportional gain, we need to identify the *ideal point* (i_1, i_2) of the game. This is likely a point outside of the bargaining region at which both players would max out their utility. In the trapezoid game, it is (5,10), in the triangle game it is (10,10).

The proportional gain for Row at a point is:

$\frac{u_1 - f_1}{i_1 - f_1}$.

The proportional concession for Row at a point is:

$\frac{i_1 - u_1}{i_1 - f_1}$.

What we need to find is the Pareto optimal point at which:

$\frac{u_1 - f_1}{i_1 - f_1} = \frac{u_2 - f_2}{i_2 - f_2}$

or

$\frac{i_1 - u_1}{i_1 - f_1} = \frac{i_2 - u_2}{i_2 - f_2}$.

Let's consider the Triangle and Trapezoid games from above. In both games the failure point is the same:

$(f_1, f_2) = (0, 0)$

In the **Triangle game** the ideal point is:

$(i_1, i_2) = (10, 10)$

The *equitable distribution point* of Triangle is (5,5) (the same as the Nash point) because at this point both players have a proportional gain of $\frac{1}{2}$ and as large as possible.

In the **Trapezoid game** the ideal point is:

$(i_1, i_2) = (5, 10)$

At (5,5) (the Nash point) the proportional gain for Col is 1 (the best she can do), but for Row it is only $1/2$. So this isn't the equitable distribution point of the game.

It turns out that at $(\frac{20}{3}, \frac{10}{3})$ the proportional gain for both players is $\frac{2}{3}$ and their gains are the largest available, so this is the *equitable distribution point* of the game.