CS 440/ECE448 Lecture 31: Game Theory

Mark Hasegawa-Johnson

Prisoner B / Prisoner A	Prisoner B stays silent (*cooperates*)	Prisoner B betrays (*defects*)
Prisoner A stays silent (*cooperates*)	Each serves 1 year	Prisoner A: 3 years Prisoner B: goes free
Prisoner A betrays (*defects*)	Prisoner A: goes free Prisoner B: 3 years	Each serves 2 years

https://en.wikipedia.org/wiki/Prisoner’s_dilemma

Today: Games with Simultaneous Moves

Assume:

- two-player game, deterministic environment (not necessary, but simplifies the problem),
- rational players (each player tries to maximize their own reward),
- not zero-sum (game can have 0, 1, or 2 winners),
- simultaneous moves.

Some surprising results:

1. The rational course of action changes may depend on your belief about what the other player will do (Nash equilibrium).
2. There are different ways to define “optimum” (Pareto optimal outcomes).
3. There may be a Pareto optimal outcome that a rational player is forced to reject (Dominant strategy).
4. In some cases, the rational thing to do is to play randomly (Mixed-strategy equilibrium).

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
  - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
  - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Chicken (Anti-Coordination Games)
  - **Mixed Nash Equilibrium**: Randomness can be the rational thing to do

Notation: sequential games

- Players take turns acting (e.g., dog moves first, then cat)
- Each node represents the action of one player (e.g., each animal can go either L or R)
- Terminal node is marked with the value for each player

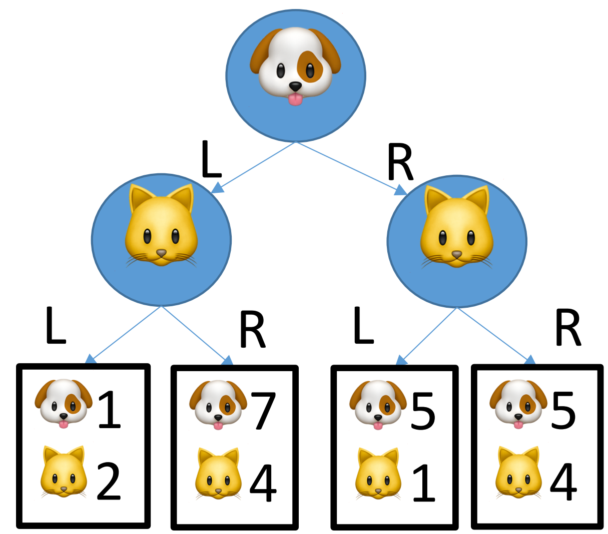

Notation: simultaneous games

The payoff matrix shows:

- Each column is a different move for player 1.
- Each row is a different move for player 2.
- Each square is labeled with the rewards earned by each player in that square.

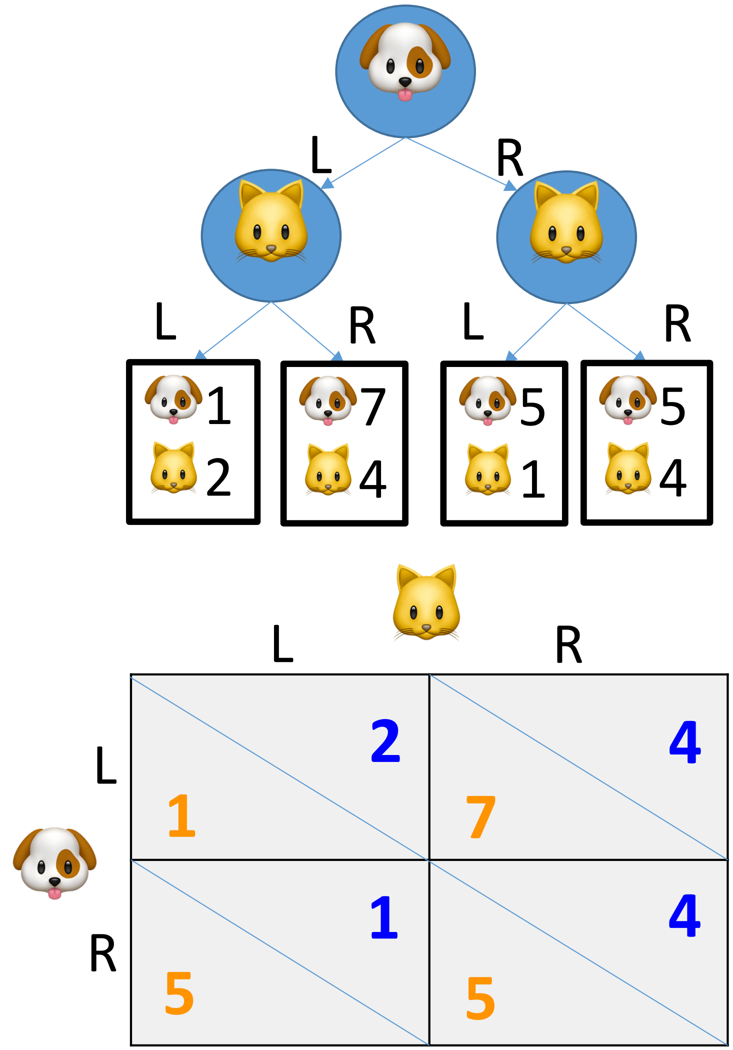

Payoff matrix

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
  - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
  - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Chicken (Anti-Coordination Games)
  - **Mixed Nash Equilibrium**: Randomness can be the rational thing to do

Stag hunt

Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

Bob \ Alice	Defect	Cooperate
Defect	Bob: 10, Alice: 10	Bob: 10, Alice: 0
Cooperate	Bob: 0, Alice: 10	Bob: 100, Alice: 100

By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

Apparently first described by Jean-Jacques Rousseau:

- If both hunters (Bob and Alice) cooperate in hunting for the stag → each gets to take home half a stag (100lbs)
- If one hunts for the stag, while the other wanders off and bags a hare → the defector gets a hare (10lbs), the cooperator gets nothing.
- If both hunters defect → each gets to take home a hare.

Nash Equilibrium

Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

		Alice	
		Defect	Cooperate
Bob	Defect	10 (Bob), 10 (Alice)	10 (Bob), 0 (Alice)
	Cooperate	0 (Bob), 10 (Alice)	100 (Bob), 100 (Alice)

By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

A Nash Equilibrium is a game outcome such that each player, **knowing the other player's move in advance**, responds rationally.

Respond rationally = the player behaves in a manner that maximizes their reward. If all rewards are the same, then all actions are rational.

Nash Equilibrium

Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

Bob \ Alice	Defect	Cooperate
Defect	10 (Bob), 10 (Alice)	10 (Bob), 0 (Alice)
Cooperate	0 (Bob), 10 (Alice)	100 (Bob), 100 (Alice)

By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

Example: (Defect,Defect) is a Nash equilibrium.

- Alice knows that Bob will defect, so she defects.
- Bob knows that Alice will defect, so he defects.
- Neither player can ***rationally*** change his or her move, unless the other player also changes.

Nash Equilibrium

Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

		Alice: Defect	Alice: Cooperate
Bob	Defect	Bob 10, Alice 10	Bob 10, Alice 0
	Cooperate	Bob 0, Alice 10	Bob 100, Alice 100

By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

(Cooperate,Cooperate) is also a Nash equilibrium!

- Alice knows that Bob will cooperate, so she cooperates!
- Bob knows that Alice will cooperate, so she cooperates!
- Neither player can ***rationally*** change his or her move, unless the other player also changes.

Surprising result #1: Nash equilibrium depends on belief


Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

		Alice	
		Defect	**Cooperate**
Bob	**Defect**	Bob: 10, Alice: 10	Bob: 10, Alice: 0
	Cooperate	Bob: 0, Alice: 10	Bob: 100, Alice: 100


By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

Surprising result:

The rational course of action depends on what you believe the other player will do.

How is "belief" formed? Answer: usually, by watching them play the game against other players, and learning their personality.

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Games)
  - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
  - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Chicken (Anti-Coordination Games)
  - **Mixed Nash Equilibrium**: Randomness can be the rational thing to do

Mending Fences

I let my neighbor know beyond the hill;
And on a day we meet to walk the line
And set the wall between us once again.
We keep the wall between us as we go.
…
There where it is we do not need the wall:
He is all pine and I am apple orchard.
My apple trees will never get across
And eat the cones under his pines, I tell him.
- from *Mending Wall*, Robert Frost

https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

Asymmetric Games

https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

Bob \ Alice	Wheat	Cattle
Wheat	Bob 20, Alice 10	Bob 10, Alice 10
Cattle	Bob 10, Alice 10	Bob 10, Alice 20

https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

Alice has lots of cattle pasture, and little wheat field.

Bob has lots of wheat field, and little cattle pasture.

- Mend the wheat fields: Bob earns $20, Alice earns $10
- Mend the cattle pastures: Alice earns $20, Bob earns $10
- If they don’t cooperate, they each earn $10

Asymmetric Games

https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

Bob \ Alice	Wheat	Cattle
Wheat	20 / 10	10 / 10
Cattle	10 / 10	10 / 20

https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

These 3 outcomes are Nash equilibria.

- (W,W): Bob has no reason to change, Alice has no reason to change
- (W,C): Bob has no reason to change, Alice has no reason to change
- (C,C): Bob has no reason to change, Alice has no reason to change

What happens if they trust one another?

https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

Bob \ Alice	Wheat	Cattle
Wheat	Bob: 20, Alice: 10	Bob: 10, Alice: 10
Cattle	Bob: 10, Alice: 10	Bob: 10, Alice: 20

https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

What happens if they discuss their actions, and make promises, and trust one another? It depends on whose needs are more urgent.

- If Bob's needs are more urgent, they will work together to mend the wheat fields
- If Alice's needs are more urgent, they will work together to mend the cattle pasture

Pareto optimal outcomes

https://commons.wikimedia.org/wiki/File:Bolivar_Heights_Battlefield,_fence_on_Bakerton_Road.jpg

Bob \ Alice	Wheat	Cattle
Wheat	Bob: 20, Alice: 10	Bob: 10, Alice: 10
Cattle	Bob: 10, Alice: 10	Bob: 10, Alice: 20

https://commons.wikimedia.org/wiki/File:Split_rail_fencing.jpg

An outcome is Pareto optimal if there is no way to increase value for one player except by decreasing value for the other.

- (W,C) and (C,W) are **not** Pareto-optimal: either player, by changing their action, can increase reward for the other player ***without decreasing reward for themselves***
- (C,C) and (W,W) are **both** Pareto-optimal: the only way to increase reward for Bob is by decreasing reward for Alice, or vice versa

Stag hunt

Photo by Scott Bauer, Public Domain, https://commons.wikimedia.org/w/index.php?curid=245466

Bob \ Alice	Defect	Cooperate
Defect	Bob: 10, Alice: 10	Bob: 10, Alice: 0
Cooperate	Bob: 0, Alice: 10	Bob: 100, Alice: 100

By Ancheta Wis, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=68432449

An outcome is Pareto optimal if there is no way to increase value for one player except by decreasing value for the other.

For the stag hunt, only (C,C) is Pareto-optimal.

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
  - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
  - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Chicken (Anti-Coordination Games)
  - **Mixed Nash Equilibrium**: Randomness can be the rational thing to do

Prisoner's dilemma

- Two criminals have been arrested and the police visit them separately
- If one player testifies against the other and the other refuses, the one who testified goes free and the one who refused gets a 10-year sentence
- If both players testify against each other, they each get a 5-year sentence
- If both refuse to testify, they each get a 1-year sentence

By Monogram Pictures, Public Domain, https://commons.wikimedia.org/w/index.php?curid=50338507

	Alice: Testify	Alice: Refuse
Bob: Testify		
Bob: Refuse		

Prisoner’s dilemma

- Two criminals have been arrested and the police visit them separately
- If one player testifies against the other and the other refuses, the one who testified goes free and the one who refused gets a 10-year sentence
- If both players testify against each other, they each get a 5-year sentence
- If both refuse to testify, they each get a 1-year sentence

By Monogram Pictures, Public Domain, https://commons.wikimedia.org/w/index.php?curid=50338507

	Alice: Testify	Alice: Refuse
Bob: Testify	-5 / -5	-10 / 0
Bob: Refuse	0 / -10	-1 / -1

Pareto optimality

	Alice: Testify	Alice: Refuse
Bob: Testify	-5 -5	-10 0
Bob: Refuse	0 -10	-1 -1

If you were permitted to discuss options with the other player, what are the different possible outcomes that might result from that discussion?

- If Bob's needs are considered most important, the (-10,0) outcome might result.
- If Alice's needs are considered more important, the (0,-10) outcome might result.
- If their needs are equally important, the (-1,-1) outcome might result.

A ***Pareto optimal*** outcome is an outcome whose cost to player A can only be reduced by increasing the cost to player B.

By Monogram Pictures, Public Domain, https://commons.wikimedia.org/w/index.php?curid=50338507

Nash equilibrium

If you knew in advance what your opponent was going to do, what would you do?

- If Bob knew that Alice was going to refuse, then it be rational for Bob to testify (he'd get 0 years, instead of 1).
- If Alice knew that Bob was going to testify, then it would be rational for her to testify (she'd get 5 years, instead of 10).
- If Bob knew that Alice was going to testify, then it would be rational for him to testify (he'd get 5 years, instead of 10).

A ***Nash equilibrium*** is an outcome such that foreknowledge of the other player's action does not cause either player to change their action.

	Alice: Testify	Alice: Refuse
Bob: Testify	Alice: -5, Bob: -5	Alice: -10, Bob: 0
Bob: Refuse	Alice: 0, Bob: -10	Alice: -1, Bob: -1

By Monogram Pictures, Public Domain, https://commons.wikimedia.org/w/index.php?curid=50338507

Dominant strategy

If you didn't know in advance what your opponent was going to do, what would you do?

- If Bob knew that Alice was going to refuse, then it be rational for Bob to testify (he'd get 0 years, instead of 1).
- If Bob knew that Alice was going to testify, then it would still be rational for him to testify (he'd get 5 years, instead of 10).

A ***dominant strategy*** is an action that maximizes reward, for one player, regardless of what the other player does.

	Alice: Testify	**Alice: Refuse**
Bob: Testify	Alice: -5, Bob: -5	Alice: -10, Bob: 0
Bob: Refuse	Alice: 0, Bob: -10	Alice: -1, Bob: -1

By Monogram Pictures, Public Domain, https://commons.wikimedia.org/w/index.php?curid=50338507

What makes it a Prisoner’s Dilemma?

We use that term to mean a game in which

- Defecting is the ***dominant strategy*** for each player, therefore
- (Defect,Defect) is the only ***Nash equilibrium***, even though
- (Defect,Defect) is not a ***Pareto-optimal solution***.

	Defect	**Cooperate**
Defect	Lose / Lose	Lose Big / Win Big
Cooperate	Win Big / Lose Big	Win / Win

http://en.wikipedia.org/wiki/Prisoner’s_dilemma

Prisoner's Dilemma vs. Stag Hunt

Prisoner's Dilemma

	Defect	Cooperate
Defect	Lose Lose	Lose Big Win Big
Cooperate	Win Big Lose Big	Win Win

Players ***improve*** their winnings by defecting unilaterally

Stag Hunt

	Defect	Cooperate
Defect	Win Win	Lose Win
Cooperate	Win Lose	Win Big Win Big

Players ***reduce*** their winnings by defecting unilaterally

Outline of today's lecture

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - **Nash Equilibrium**: Each player knows what the other will do, and responds rationally
- Example: Mending Fences (Asymmetric Coordination)
  - **Pareto Optimal outcome**: No player can win more w/o some other player winning less
- Example: Prisoners' Dilemma (Betrayal Games)
  - **Dominant Strategy**: an action that is rational regardless of what the other player does
- Example: Chicken (Anti-Coordination Games)
  - **Mixed Nash Equilibrium**: Randomness can be the rational thing to do

Game of Chicken

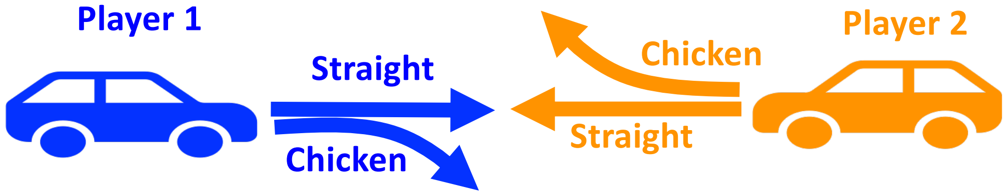

		Alice: Straight	Alice: Chicken
Bob	**Straight**		Bob: 2, Alice: -1
	Chicken	Bob: -1, Alice: 2	Bob: 1, Alice: 1

- Two players each bet $1000 that the other player will chicken out
- Outcomes:
  - If one player chickens out, he loses $1000, and the other wins $2000
  - If both players chicken out, they each keep their original $1000
  - If neither player chickens out, they both lose $10,000 (the cost of the car)

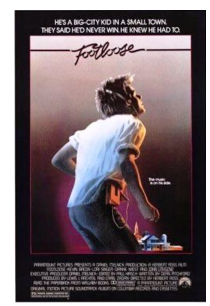

http://en.wikipedia.org/wiki/Game_of_chicken

Prisoner's Dilemma vs. Game of Chicken

Prisoner's Dilemma

	Defect	Cooperate
Defect	Lose Lose	Lose Big Win Big
Cooperate	Win Big Lose Big	Win Win

Players cut their losses by defecting if the other player defects

Game of Chicken

	Straight	Chicken
Straight	Lose Big Lose Big	Lose Win Big
Chicken	Win Big Lose	Win Win

Defecting, if the other player defects, is the worst thing you can do

Game of Chicken

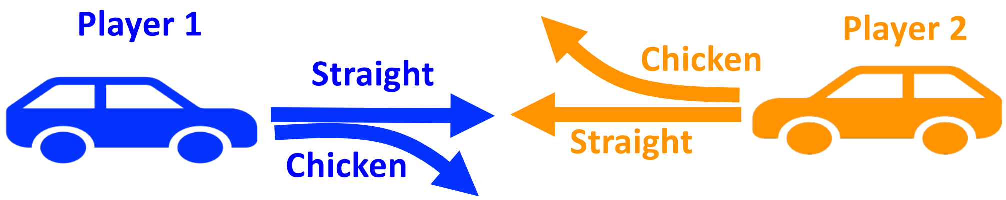

Bob \ Alice	Straight	Chicken
Straight	-10 (Alice), -10 (Bob)	-1 (Alice), 2 (Bob)
Chicken	2 (Alice), -1 (Bob)	1 (Alice), 1 (Bob)

- Is there a dominant strategy for either player?
- Is there a Nash equilibrium?
  (straight, chicken) or (chicken, straight)
- *Anti-coordination* game: it is mutually beneficial for the two players to choose different strategies
  - Model of escalated conflict in humans and animals (hawk-dove game)
- How are the players to decide what to do?
  - Bluff! You have to somehow convince your opponent that you will drive straight, no matter what happens, even if it's irrational for you to do so.
  - In that case, the rational thing for your opponent to do is to chicken out.

http://en.wikipedia.org/wiki/Game_of_chicken

Game of Chicken

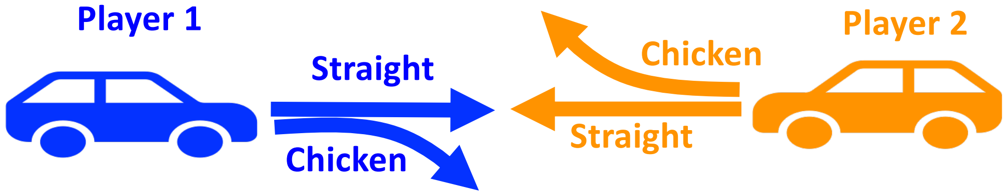

		Alice	
		Straight	Chicken
Bob	Straight	-10 / -10	2 / -1
	Chicken	-1 / 2	1 / 1

- Is there a dominant strategy for either player?
- Is there a Nash equilibrium?

  (straight, chicken) or (chicken, straight)
- *Anti-coordination* game: it is mutually beneficial for the two players to choose different strategies
  - Model of escalated conflict in humans and animals (hawk-dove game)
- How are the players to decide what to do?
  - Bluff! You have to somehow convince your opponent that you will drive straight, no matter what happens, even if it's irrational for you to do so.
  - In that case, the rational thing for your opponent to do is to chicken out.

Seriously??!! Is there no way to win this game without convincing the other player that you are irrational??!!

http://en.wikipedia.org/wiki/Game_of_chicken

Irrational versus Random

The game of chicken has two different types of Nash equilibria:

- **Be irrational**: Bluff. One player convinces the other that he or she will behave irrationally. The other player concedes the game. Result: (straight,chicken) or (chicken,straight).
- **Be random:** Mixed Nash Equilibrium.
  - Alice chooses a move at random, according to some probability distribution. She tells Bob, in advance, what probability distribution she will use.
  - Bob responds rationally.
  - One of Bob's rational options is to choose his move, also, at random.

Mixed strategy = Random

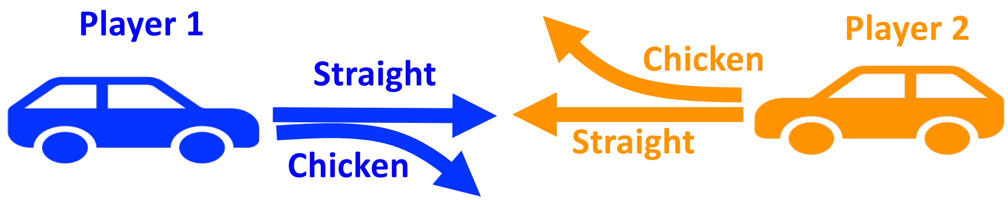

Bob \ Alice	Straight	Chicken
Straight	Bob: -10, Alice: -10	Bob: 2, Alice: -1
Chicken	Bob: -1, Alice: 2	Bob: 1, Alice: 1

- **Mixed strategy:** a player chooses between the different possible actions according to a probability distribution.
- **Mixed Nash equilibrium**:
  - One or both players choose their actions at random
  - For both players, this is a rational thing to do, i.e., they can’t increase their expected reward by using a non-random strategy

Mixed strategy = Random

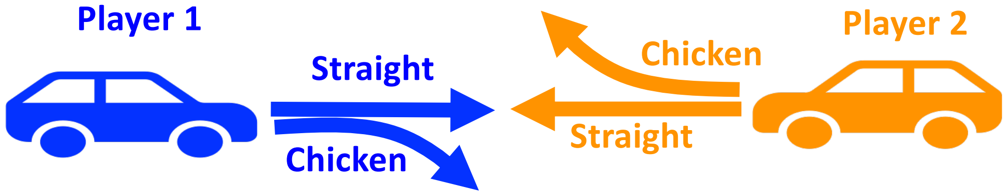

		Alice: Straight	Alice: Chicken
Bob	Straight	-10 (Alice), -10 (Bob)	-1 (Alice), 2 (Bob)
	Chicken	2 (Alice), -1 (Bob)	1 (Alice), 1 (Bob)

For example, suppose that both players, independently, decide to go straight with probability 1/10.

- Random variable A: Alice's action

$$P(A = s) = \frac{1}{10}, \qquad P(A = c) = \frac{9}{10}$$

- Random variable B: Bob's action

$$P(B = s) = \frac{1}{10}, \qquad P(B = c) = \frac{9}{10}$$

Expected reward

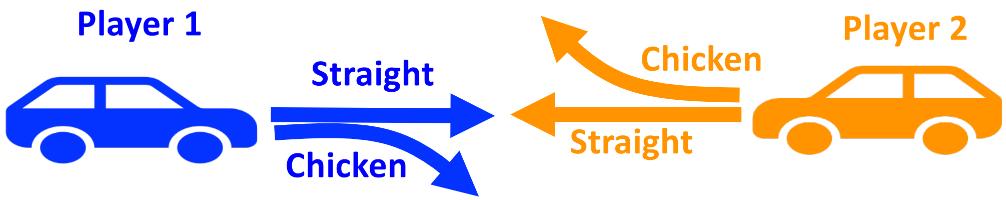

Bob \ Alice	Alice: Straight	Alice: Chicken
Straight	Bob: -10, Alice: -10	Bob: 2, Alice: -1
Chicken	Bob: -1, Alice: 2	Bob: 1, Alice: 1

The expected reward for Alice is independent of her action:

$$E[r_A|\mathrm{A}=\mathrm{s}] = \mathrm{P}(\mathrm{B}=s)r_A(s,s) + \mathrm{P}(\mathrm{B}=c)\, r_A(s,c) = \left(\frac{1}{10}\right)(-10) + \left(\frac{9}{10}\right)(2) = \frac{8}{10}$$

$$E[r_A|\mathrm{A}=\mathrm{c}] = \mathrm{P}(\mathrm{B}=s)r_A(c,s) + \mathrm{P}(\mathrm{B}=c)r_A(c,c) = \left(\frac{1}{10}\right)(-1) + \left(\frac{9}{10}\right)(1) = \frac{8}{10}$$

The expected reward for Bob is independent of his action:

$$E[r_B|\mathrm{B}=\mathrm{s}] = \mathrm{P}(\mathrm{A}=s)r_B(s,s) + \mathrm{P}(\mathrm{A}=c)r_B(s,c) = \left(\frac{1}{10}\right)(-10) + \left(\frac{9}{10}\right)(2) = \frac{8}{10}$$

$$E[r_B|\mathrm{B}=\mathrm{c}] = \mathrm{P}(\mathrm{A}=s)r_B(c,s) + \mathrm{P}(\mathrm{A}=c)r_B(c,c) = \left(\frac{1}{10}\right)(-1) + \left(\frac{9}{10}\right)(1) = \frac{8}{10}$$

Mixed Nash equilibrium

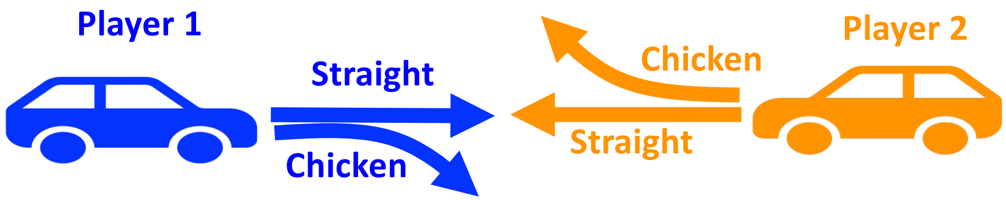

		Alice: Straight	Alice: Chicken
Bob	Straight	-10 (Bob), -10 (Alice)	2 (Bob), -1 (Alice)
	Chicken	-1 (Bob), 2 (Alice)	1 (Bob), 1 (Alice)

- Neither player can increase their reward by choosing a deterministic action:

$$E[r_A|\mathrm{A}=\mathrm{s}] = E[r_A|\mathrm{A}=\mathrm{c}] = \frac{8}{10}$$

$$E[r_B|\mathrm{B}=\mathrm{s}] = E[r_B|\mathrm{B}=\mathrm{c}] = \frac{8}{10}$$

- Going straight, chickening out, or choosing at random between the two are all rational actions. Random choice with probability $\mathrm{P}(c) = \frac{9}{10}$ is a rational action, as would be any other strategy.

Unstable equilibrium

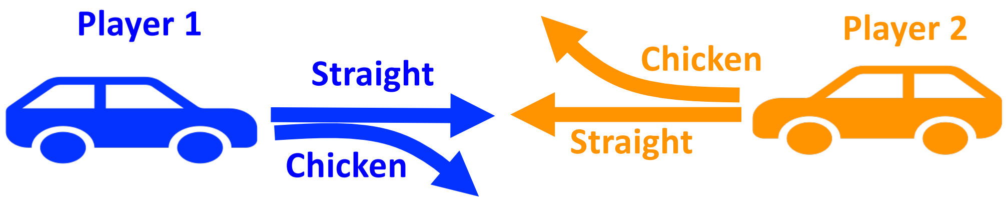

Bob \ Alice	Straight	Chicken
Straight	Alice: -10, Bob: -10	Alice: -1, Bob: 2
Chicken	Alice: 2, Bob: -1	Alice: 1, Bob: 1

- Random choice with probability $\mathrm{P}(A=c)=\frac{9}{10}$ is a rational action; but any other strategy would also be rational. So why should Alice choose this strategy in particular?
- Notice: if Alice increases her probability of going straight ***even a little bit***, then it suddenly becomes rational for Bob to always chicken out.
- The mixed Nash equilibrium in this case is an ***unstable equilibrium***: as long as $\mathrm{P}(A=c)=\mathrm{P}(B=c)=\frac{9}{10}$, then there is no reason for either of them to change, but if one of them changes by even a little bit, then the other is forced to change in response (if behaving rationally).

Calculating the mixed equilibrium

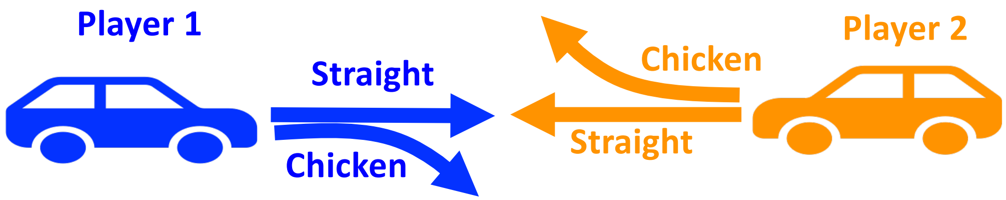

Bob \ Alice	Defect	Cooperate
Defect	Alice: -10; Bob: -10	Alice: -1; Bob: 2
Cooperate	Alice: 2; Bob: -1	Alice: 1; Bob: 1

There is a Mixed Nash equilibrium if and only if there are some probabilities $0 \leq P_A \leq 1$ and $0 \leq P_B \leq 1$ such that Alice's expected reward is independent of her action:

$$(1 - P_B)r_A(d,d) + P_B r_A(d,c) = (1 - P_B)r_A(c,d) + P_B r_A(c,c)$$

…and Bob's expected reward is independent of his action:

$$(1 - P_A)r_B(d,d) + P_A r_B(d,c) = (1 - P_A)r_B(c,d) + P_A r_B(c,c)$$

Try the quiz!

Try the quiz!

Does every game have a mixed-strategy equilibrium?

A mixed-strategy equilibrium exists only if there are some $0 \leq p \leq 1$ and $0 \leq q \leq 1$ that solve these equations:

$$(1-p)w + px = (1-p)y + pz$$
$$(1-q)a + qc = (1-q)b + qd$$

That's not necessarily possible for every game. For example, it's not true for Prisoner's Dilemma.

- Prisoner's Dilemma has only one fixed-strategy Nash equilibrium (both players defect).
- Stag Hunt has two fixed-strategy Nash equilibria (either both players cooperate, or both players defect), and one mixed-strategy equilibrium (each player cooperates with probability 1/10).
- The Game of Chicken has:
  - 2 fixed strategy Nash equilibria (Alice defects while Bob cooperates, or vice versa)
  - 1 mixed-strategy Nash equilibrium (both Alice and Bob each defect with probability 1/10).

	Defect w/ Prob. $1-p$	Coop. w/ Prob. p
Defect w/ Prob. $1-q$	$\boldsymbol{a}$ / $\boldsymbol{w}$	$\boldsymbol{b}$ / $\boldsymbol{x}$
Coop. w/ Prob. q	$\boldsymbol{c}$ / $\boldsymbol{y}$	$\boldsymbol{d}$ / $\boldsymbol{z}$

Existence of Nash equilibria

- Any game with a finite set of actions has at least one Nash equilibrium (which may be a mixed-strategy equilibrium).
- If a player has a dominant strategy, there exists a Nash equilibrium in which the player plays that strategy and the other player plays the *best response* to that strategy.
- If both players have dominant strategies, there exists a Nash equilibrium in which they play those strategies.

Summary

- Dominant strategy
  - a strategy that's optimal for one player, regardless of what the other player does
  - Not all games have dominant strategies
- Nash equilibrium
  - an outcome (one action by each player) such that, knowing the other player's action, each player has no reason to change their own action
  - Every game with a finite set of actions has at least one Nash equilibrium, though it might be a mixed-strategy equilibrium.
- Pareto optimal
  - an outcome such that neither player would be able to win more without simultaneously forcing the other player to lose more
  - Every game has at least one Pareto optimal outcome. Usually there are many, representing different tradeoffs between the two players.
- Mixed strategies
  - A mixed strategy is optimal only if there's no reason to prefer one action over the other, i.e., if $0 \leq P_A \leq 1$ and $0 \leq P_B \leq 1$ such that:

$$(1 - P_B) r_A(d, d) + P_B r_A(d, c) = (1 - P_B) r_A(c, d) + P_B r_A(c, c)$$

$$(1 - P_A) r_B(d, d) + P_A r_B(d, c) = (1 - P_A) r_B(c, d) + P_A r_B(c, c)$$