# CS 440/ECE448 Lecture 35: Game Theory

Mark Hasegawa-Johnson, 4/2020

Including slides by Svetlana Lazebnik



| Prisoner A \ Prisoner B | **Prisoner B stays silent (*cooperates*)** | **Prisoner B betrays (*defects*)** |
|---|---|---|
| **Prisoner A stays silent (*cooperates*)** | Each serves 1 year | Prisoner A: 3 years<br>Prisoner B: goes free |
| **Prisoner A betrays (*defects*)** | Prisoner A: goes free<br>Prisoner B: 3 years | Each serves 2 years |

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

# Game theory

- **Game theory** deals with systems of interacting agents where the outcome for an agent depends on the actions of all the other agents
  - Applied in sociology, politics, economics, biology, and, of course, AI
- **Agent design:** determining the best strategy for a rational agent in a given game
- **Mechanism design:** how to set the rules of the game to ensure a desirable outcome

The Economist

Modelling behaviour

# Game theory in practice

**Computing: Software that models human behaviour can make forecasts, outfox rivals and transform negotiations**

Sep 3rd 2011 | from the print edition


http://www.economist.com/node/21527025


http://www.spliddit.org

CADE METZ BUSINESS 09.21.15 7:00 AM

# FACEBOOK DOESN'T MAKE AS MUCH MONEY AS IT COULD—ON PURPOSE

YOU CAN THINK of John Hegeman as Facebook's chief economist. He spends his days thinking about the economics of Facebook advertising.

That's an enormous thing. Facebook pulled in $4.04 billion in the second quarter of this year. And the overall economy of Facebook advertising, as Hegeman describes it, is far larger. Advertising, you see, is very much a part of everything else on the world's largest social network. Hegeman doesn't just think about ads. He thinks about how ads fit with the rest of Facebook.

When he joined Facebook in 2007, after getting a master's in economics at Stanford University, Hegeman helped build the online auction that drives the company's advertising system. Auctions are the standard way that online services accept ads from advertisers and place them on web pages and inside smartphone apps. That's what Google uses with AdWords, the system that serves up all those ads when you look for stuff on the company's Internet search engine. Advertisers bid (in dollars) for placement on the results page when you key in a particular word or group of words. But in building Facebook's advertising system, Hegeman and team took online auctions in a new direction.

http://www.wired.com/2015/09/facebook-doesnt-make-much-money-couldon-purpose/

# Outline of today’s lecture

- What is a game?
- What are the questions you can ask?
- Situations with different types of payout matrices
  - Prisoners’ Dilemma: Betrayal Games
  - Stag Hunt: Coordination Games
  - Chicken: Anti-Coordination Games
- What types of strategy are possible?
  - Without knowing the other player’s strategy: Dominant strategy
  - Knowing the other player’s strategy: Nash equilibrium, Pareto optimality
  - Mixed strategies

# What is a game?

Assume that the environment is:

- Fully observable. You can't see thoughts, but you can see actions.
- Deterministic. Actions determine rewards, no randomness.
- Episodic (we'll talk about sequential games next time).
- Static. The environment doesn't change.
- Discrete. You have a small finite set of possible actions.
- Known: all the rules are known in advance.

Despite choosing the simplest type of environment in all six of those categories, rational decision-making is extremely challenging because the environment is:

- Multi-agent: there are two players, each trying to maximize benefit.

# Recall: non-zero-sum games

Each player tries to maximize their own benefit.

Outcome of the game can be predicted using an algorithm similar to minimax: each player makes the best decision for the situation in which they find themselves.

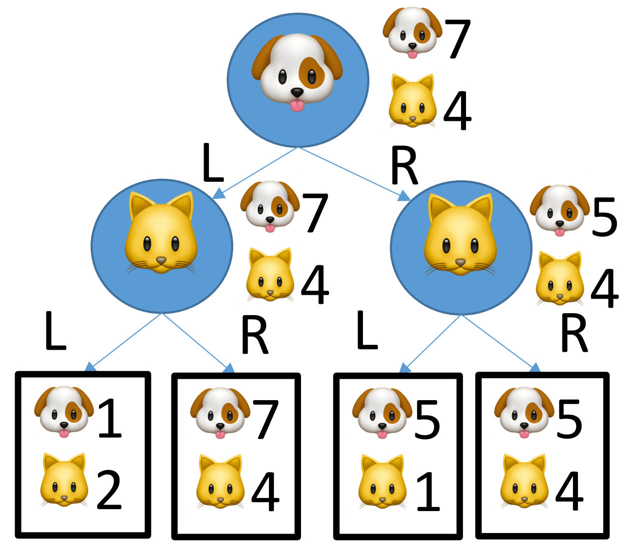

# Payoff matrix

In Game Theory, it's useful to summarize the possible outcomes of the game using a ***payoff matrix***: a list of all possible outcomes, indexed by the actions of each player.

This is also called a normal-form representation of the game.

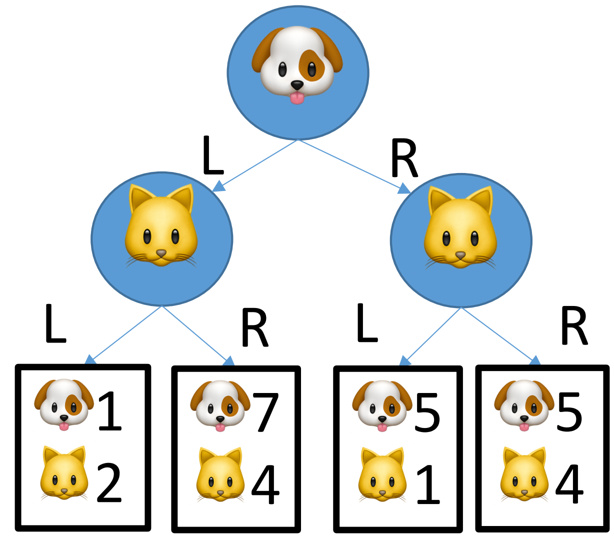


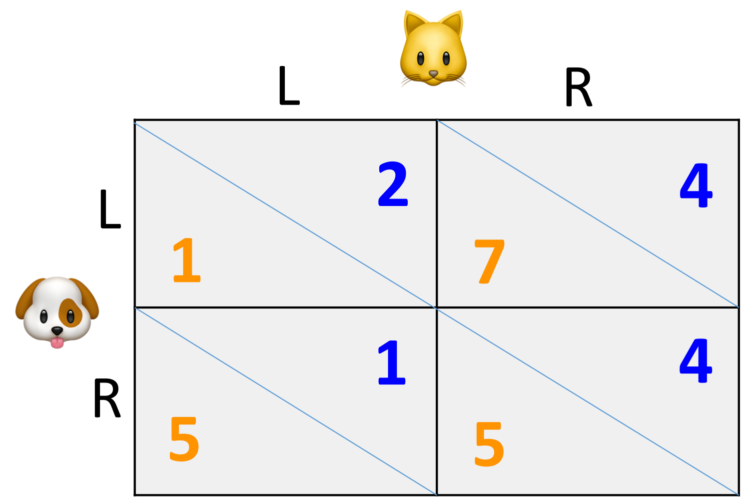


**Payoff matrix**

# The types of questions that Game Theory asks

- What happens if you don't know what the other player will do?
- Are there games that have an optimal strategy even when you don't know what the other player will do?
- If you knew the other player's action in advance, under what circumstances would that cause you to change your own action?

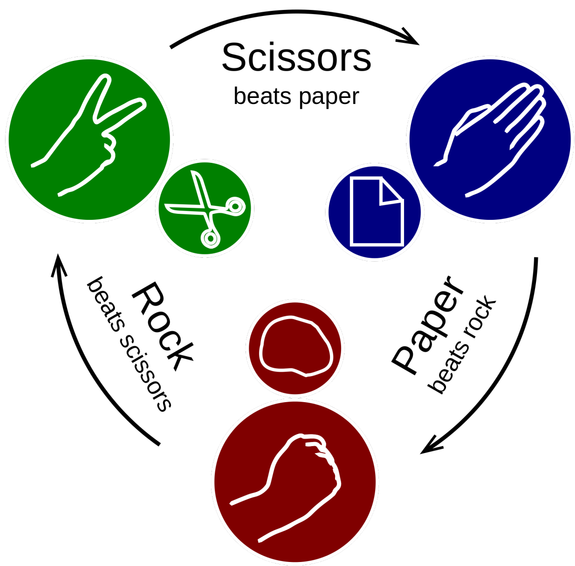


By Enzoklop - Own work, CC BY-SA 3.0,
https://commons.wikimedia.org/w/index.php?curid=27958688

***Normal form*** representation:

**Player 1**

**Player 2**

| |  |  |  |
|---|---|---|---|
|  | 0 / 0 | 1 / -1 | -1 / 1 |
|  | -1 / 1 | 0 / 0 | 1 / -1 |
|  | 1 / -1 | -1 / 1 | 0 / 0 |

**Payoff matrix**

# Outline of today's lecture

- What is a game?
- What are the questions you can ask?
- Situations with different types of payout matrices
  - Prisoners' Dilemma: Betrayal Games
  - Stag Hunt: Coordination Games
  - Chicken: Anti-Coordination Games
- What types of strategy are possible?
  - Without knowing the other player's strategy: Dominant strategy
  - Knowing the other player's strategy: Nash equilibrium, Pareto optimality
  - Mixed strategies

# Payoff matrices

- Working for RAND (a defense contractor) in 1950, Flood and Dresher formalized the “Prisoner’s Dilemma” (PD): a class of payoff matrices that encourages betrayal.
- Jean-Jacques Rosseau (Swiss philosopher, 1700s) invented the “Stag Hunt” (SH): a class of payoff matrices that reward cooperation, but don’t force it. Has been used as a model of climate-change treaties.
- Both PD and SH have stable Nash equilibria. The “Game of Chicken” is a popular subject in movies (*Rebel Without a Cause*, *Footloose*, *Crazy Rich Asians*) because of its inherent instability: the only way to win is by convincing your opponent to lose.

# 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 | 0 / 10 |
| Bob: Refuse | 10 / 0 | 1 / 1 |

# Questions that can be asked

- If you were permitted to discuss options with the other player, but if one of you is more persuasive than the other, what are the different possible outcomes that might result from that discussion?
- If you knew in advance what your opponent was going to do, what would you do?
- If you didn’t know in advance what your opponent was going to do, what would you do?

# Pareto optimality

If you were permitted to discuss options with the other player, but if one of you is more persuasive than the other, what are the different possible outcomes that might result from that discussion?

- If Bob was most persuasive, the (10,0) outcome might result.
- If Alice was most persuasive, the (0,10) outcome might result.
- If equally persuasive, 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.

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


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 ***<u>Nash equilibrium</u>*** 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 | 5 / 5 | 10 / 0 |
| Bob: Refuse | 0 / 10 | 1 / 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 minimizes cost, for one player, regardless of what the other player does.

| | Alice: Testify | Alice: Refuse |
|---|---|---|
| Bob: Testify | 5 5 | 10 0 |
| Bob: Refuse | 0 10 | 1 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 in real life

- Price war
- Arms race
- Steroid use
- Diner's dilemma
- Collective action in politics

| | Defect | Cooperate |
|---|---|---|
| Defect | Lose / Lose | Lose Big / Win |
| Cooperate | Win / Lose Big | Draw / Draw |

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

# How do we avoid Prisoners’ Dilemma situations?

Repeated games.

More next time.

| | Defect | Cooperate |
|---|---|---|
| Defect | Lose<br>Lose | Lose Big<br>Win |
| Cooperate | Win<br>Lose Big | Draw<br>Draw |

# The Stag Hunt: Coordination Games

# Stag hunt

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

| | Defect | Cooperate |
|---|---|---|
| **Defect** | 10, 10 | 10, 0 |
| **Cooperate** | 0, 10 | 100, 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 cooperate in hunting for the stag → each gets to take home half a stag (100kg)
- If one hunts for the stag, while the other wanders off and bags a hare → the defector gets a hare (10kg), the cooperator gets nothing.
- If both hunters defect → each gets to take home a hare.

# Stag hunt

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

| | Defect | Cooperate |
|---|---|---|
| Defect | 10<br>10 | 0<br>10 |
| Cooperate | 10<br>0 | 100<br>100 |

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

- What is/are the Pareto Optimal solution(s)?
- What is/are the Nash Equilibrium/a?
- Is there a Dominant Strategy for either player?
- Model for cooperative activity under conditions of incomplete information (the issue: trust)

# 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

# Chicken: Anti-Coordination Games, Mixed Strategies

# Game of Chicken

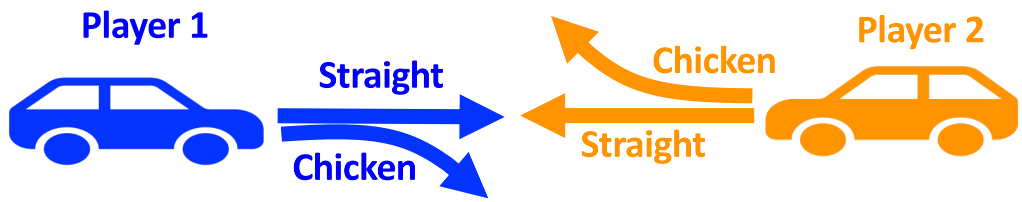


| | Straight | Chicken |
|---|---|---|
| Straight | -10 / -10 | -1 / 1 |
| Chicken | 1 / -1 | 0 / 0 |

- Two players each bet $1000 that the other player will chicken out
- Outcomes:
  - If one player chickens out, the other wins $1000
  - If both players chicken out, neither wins anything
  - 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

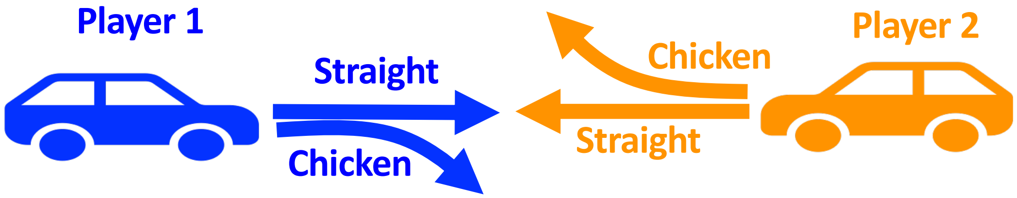


| | Straight | Chicken |
| --- | --- | --- |
| Straight | -10 / -10 | 1 / -1 |
| Chicken | -1 / 1 | 0 / 0 |

- 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?
  - Pre-commitment or threats
  - Different roles: the "hawk" is the territory owner and the "dove" is the intruder, or vice versa

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

# Game of Chicken

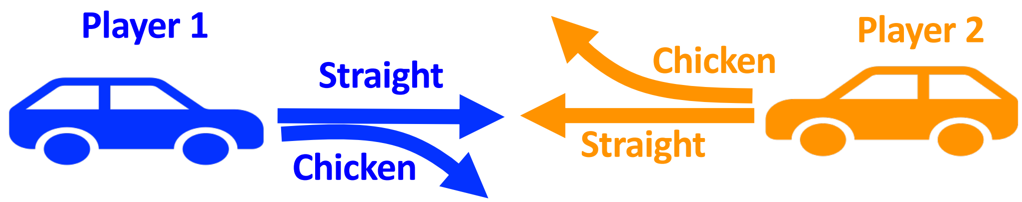


| | Straight | Chicken |
|---|---|---|
| Straight | -10 / -10 | -1 / 1 |
| Chicken | 1 / -1 | 0 / 0 |

- **Mixed strategy:** a player chooses between the different possible actions according to a probability distribution.
- For example, suppose that each player chooses to go straight (S) with probability **1/10**. Is that a Nash equilibrium?

# Game of Chicken

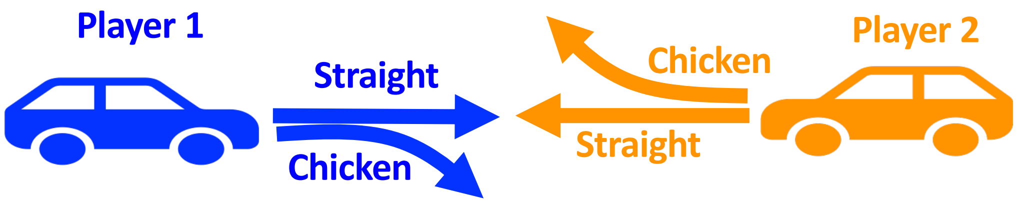


| | Straight | Chicken |
|---|---|---|
| Straight | -10, -10 | 1, -1 |
| Chicken | -1, 1 | 0, 0 |

The expected payoff, to player P1, for choosing to go Straight is:

$$E[\text{Payoff}] = Pr(P2\ chooses\ S)\times\text{Payoff}(to\ P1\ if\ S,S) + Pr(P2\ chooses\ C)\times\text{Payoff}(to\ P1\ if\ S,C)$$

$$= \left(\frac{1}{10}\right)\times(-10) + \left(\frac{9}{10}\right)\times(1) = -\frac{1}{10}$$

The expected payoff, to player P1, for choosing to Chicken Out is:

$$E[\text{Payoff}] = Pr(P2\ chooses\ S)\times\text{Payoff}(to\ P1\ if\ C,S) + Pr(P2\ chooses\ C)\times\text{Payoff}(to\ P1\ if\ C,C)$$

$$= \left(\frac{1}{10}\right)\times(-1) + \left(\frac{9}{10}\right)\times(0) = -\frac{1}{10}$$

So Player P1 has no preference between actions S and C: he's free to choose between them according to a random number generator.

# Finding mixed strategy equilibria

| | | Alice: Defect w/ Prob. $1-p$ | Alice: Coop. w/ Prob. $p$ |
|---|---|---|---|
| Bob | Defect w/ Prob. $1-q$ | $w$, $a$ | $x$, $b$ |
| | Coop. w/ Prob. $q$ | $y$, $c$ | $z$, $d$ |

Here's the trick: for Bob, random selection is rational only if he can't improve his winnings by definitively choosing one action or the other. So, for Bob to decide whether a mixed strategy is rational, he needs to know:

- His own reward for each possible outcome ($w$, $x$, $y$, and $z$), and …
- the probability ($p$) of Alice cooperating.

# Finding mixed strategy equilibria

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

For Bob, random selection is rational only if he can't improve his winnings by definitively choosing one action or the other.

- If Bob defects, he expects to win $(1-p)w + px$.
- If Bob cooperates, he expects to win $(1-p)y + pz$.

So

- it's only logical for Bob to use a mixed strategy if $(1-p)w + px = (1-p)y + pz$.

# 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 either Prisoner's Dilemma or Stag Hunt.

- 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).
- 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).

# 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.

# Outline of today's lecture

- Prisoner's Dilemma
  - Nash equilibrium = both players play their dominant strategy
  - Nash equilibrium ∉ Pareto optimal
- Stag Hunt
  - called a "coordination game" because the fixed-strategy Nash equilibria occur when both players play the same way
  - no dominant strategy for either player
- Game of Chicken
  - called an "anti-coordination game" because the two fixed-strategy Nash equilibria occur when the players act in opposite ways
  - no dominant strategy for either player

# Outline of today's lecture

- 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 $0 \leq p \leq 1$ and $0 \leq q \leq 1$ such that:

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