CS 440/ECE448 Lecture 21: Game Theory

Mark Hasegawa-Johnson, 4/2021

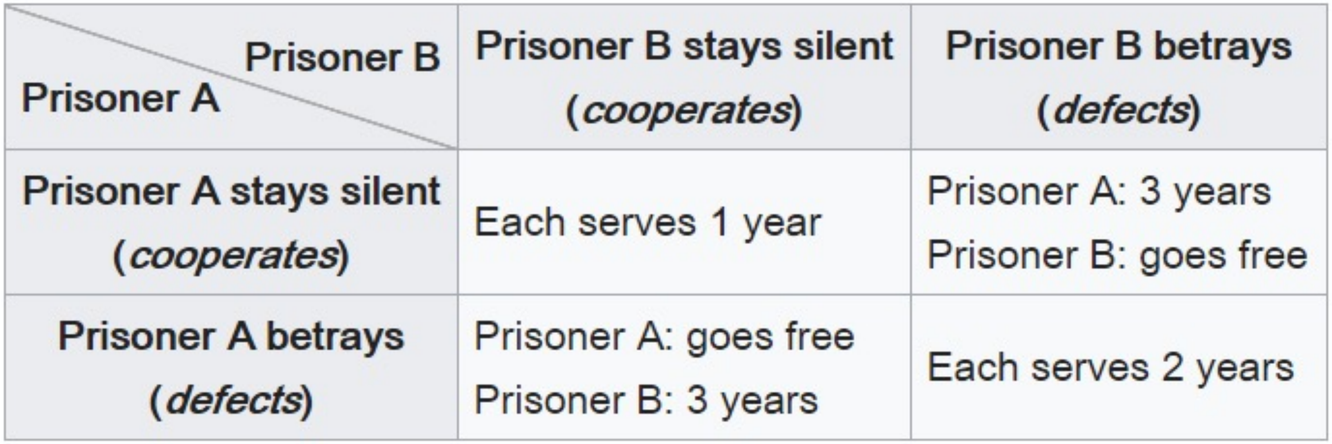

Prisoner B **Prisoner A**	**Prisoner B stays silent** **(*cooperates*)**	**Prisoner B betrays** **(*defects*)**
Prisoner A stays silent **(*cooperates*)**	Each serves 1 year	Prisoner A: 3 years Prisoner B: goes free
Prisoner A betrays **(*defects*)**	Prisoner A: goes free Prisoner B: 3 years	Each serves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

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

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

Christian Montenegro

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

Games so far: Zero-sum two-player games...

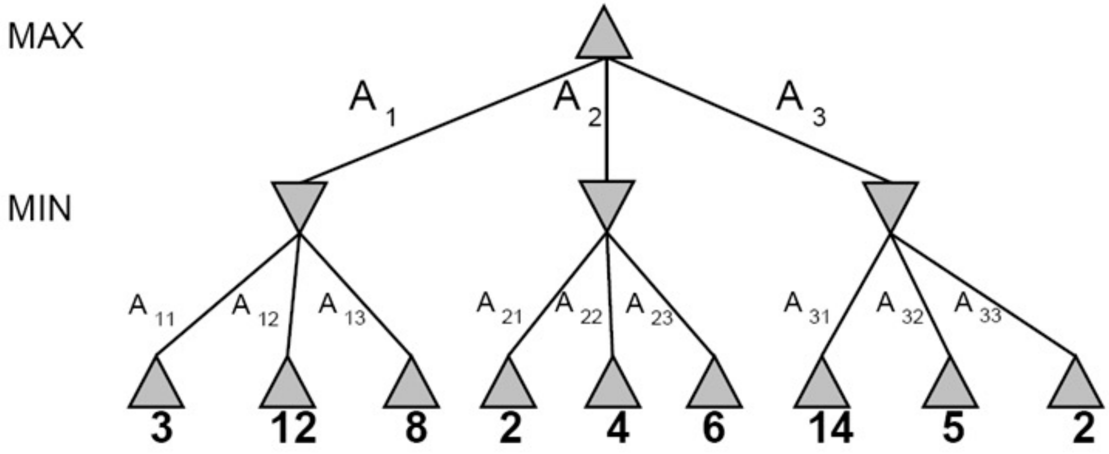

▲ = game state from which MAX can play

▼ = game state from which MIN can play

number = value of that game state for MAX

Games so far: multi-player & non-zero-sum games…

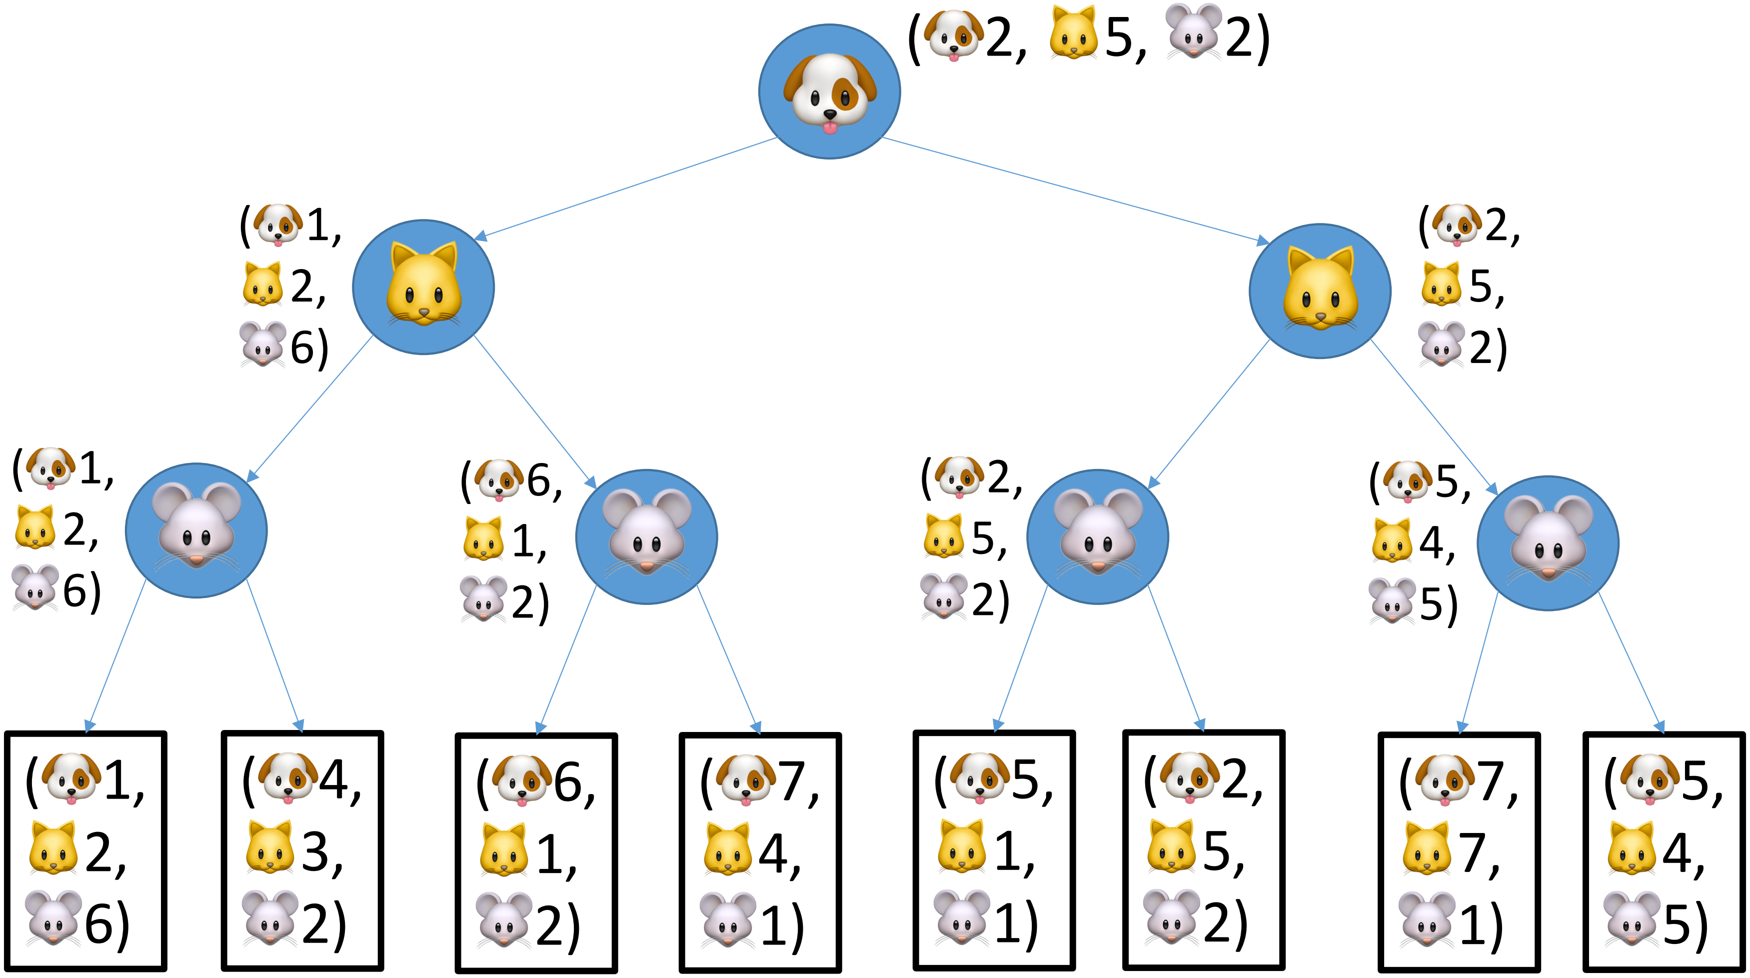

Characteristics of games we’ve seen so far

- Rational vs. Irrational opponent ✓
- Two-player vs. Multi-player ✓
- Zero-sum vs. Non-zero-sum ✓
- Deterministic vs. Stochastic ✓
- Sequential vs. Simultaneous moves **X**

Games with Simultaneous Moves

- Assume: two-player game, rational players, deterministic environment, but NOT necessarily zero-sum.
  - These assumptions are not necessary, but they simplify the problem.
- Both players play at the same time.
- What is the rational thing to do:
  1. If you know in advance what the other player will do?
  2. If you can negotiate your move with the other player?
  3. If you DON’T know in advance what the other player will do?
  4. If it is rational to behave randomly?

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: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Networks (GAN)

Notation: sequential games

- Terminal node is marked with the value for each player
- Non-terminals node inherits its value from its minimax-optimal descendant

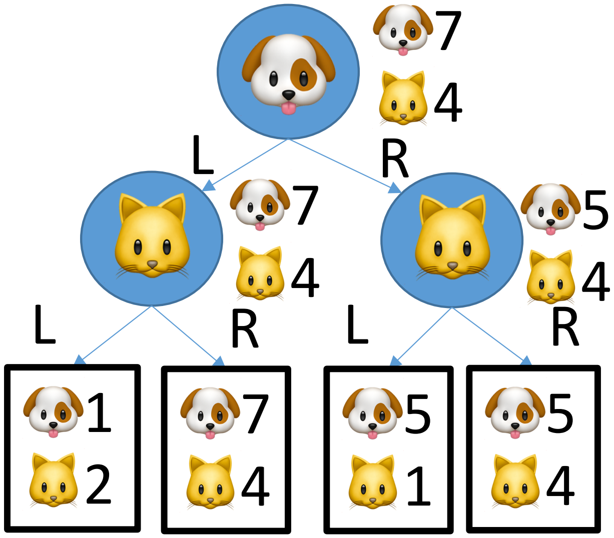

Notation: simultaneous games

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.

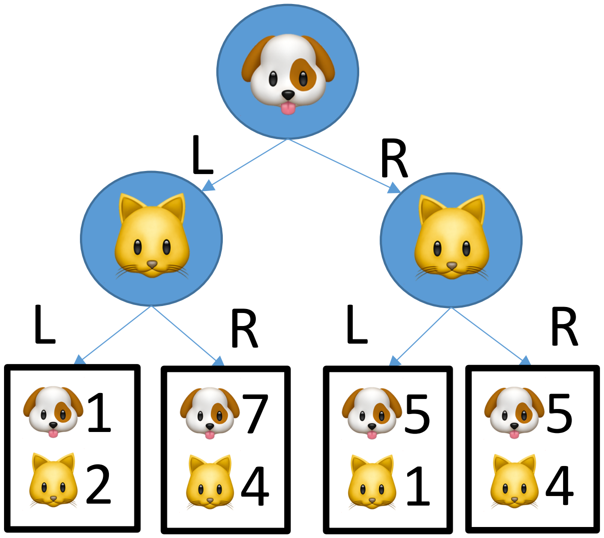

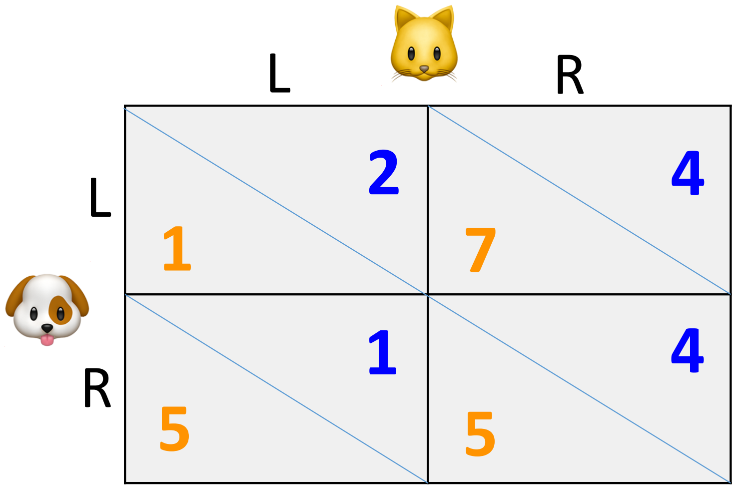

	L	R
L	2 1	4 7
R	1 5	4 5

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: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Network (GAN)

Stag hunt

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

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

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.

Nash Equilibrium

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

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

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	Cooperate
Bob	Defect	10 10	0 10
	Cooperate	10 0	100 100

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.

Pareto-optimal outcome

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

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

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

What if the players talk to each other in advance, and make promises, and trust one another's promises?

- Then they will both choose to cooperate.

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: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Network (GAN)

Asymmetric Coordination Games

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

		Alice	
		Stag	Alligator
Bob	Stag	10 20	0 0
	Alligator	0 0	20 10

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

Alice prefers alligator. Bob prefers stag.

If they don't cooperate, they each get nothing.

Asymmetric Coordination Games

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

		Alice	
		Stag	**Alligator**
Bob	**Stag**	10 20	0 0
	Alligator	0 0	20 10

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

The Nash equilibria are (Stag,Stag) and (Gator,Gator).

- If Bob knows that Alice will hunt gator, then it’s rational for him to do the same.
- If Alice knows that Bob will hunt stag, then it’s rational for him to do the same.

What happens if they trust one another?

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

		Alice	
		Stag	Alligator
Bob	Stag	10 20	0 0
	Alligator	0 0	20 10

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

What happens if they discuss their actions, and make promises, and trust one another?

It depends: whose needs are considered more important?

- If Bob’s needs are more important, then they will hunt stag.
- If Alice’s needs are more important, then they will hunt alligator.

Pareto optimal outcome

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

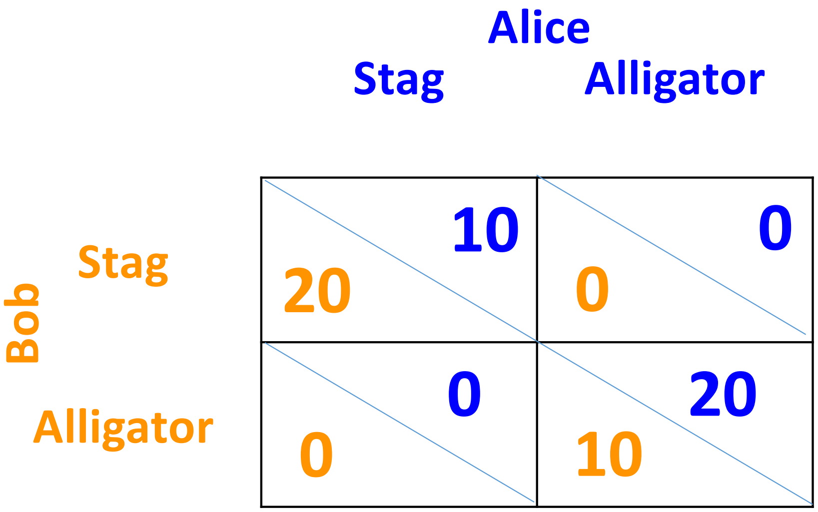

		Alice	
		Stag	Alligator
Bob	Stag	10 20	0 0
	Alligator	0 0	20 10

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

An outcome is Pareto-optimal if the only way to increase value for one player is by decreasing value for the other.

- (Stag,Stag) is Pareto-optimal: one could increase Alice's value, but only by decreasing Bob's value.
- (Alligator,Alligator) is Pareto-optimal: one could increase Bob's value, but only by decreasing Alice's value.

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: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Network (GAN)

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

	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

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

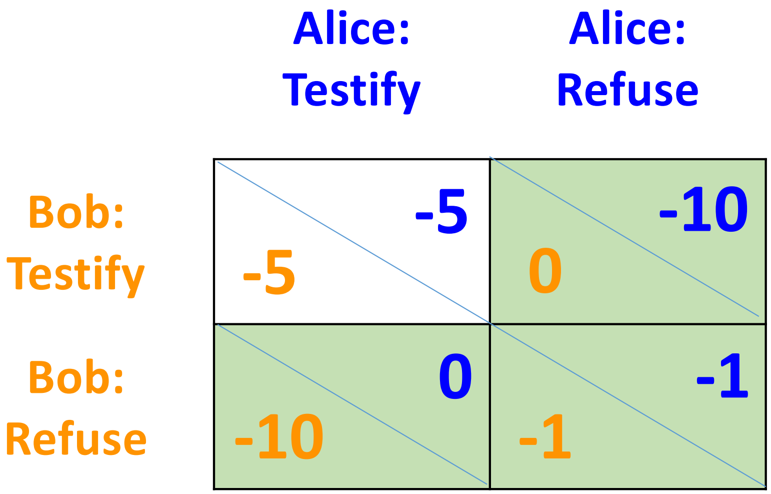

	Alice: Testify	**Alice: Refuse**
Bob: Testify	-5 -5	-10 0
Bob: Refuse	0 -10	-1 -1

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.

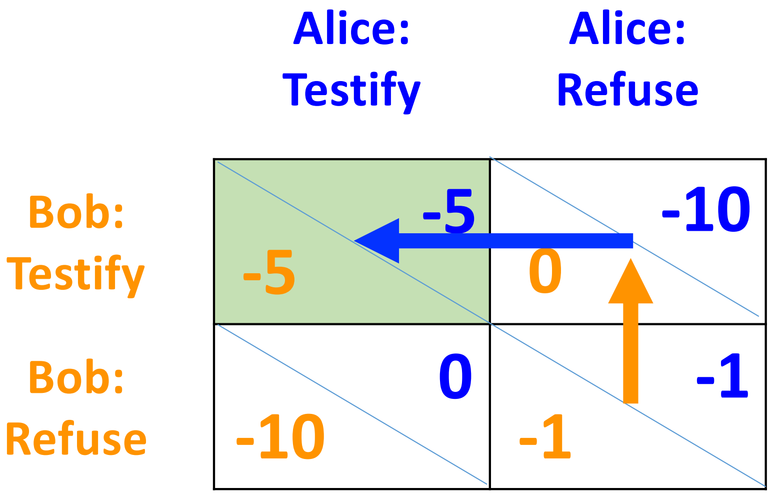

	Alice: Testify	**Alice: Refuse**
Bob: Testify	-5 -5	-10 0
Bob: Refuse	0 -10	-1 -1

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.

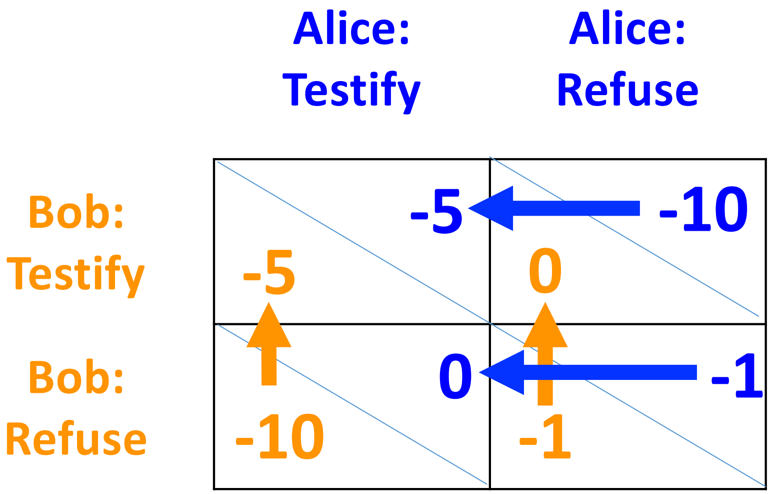

	Alice: Testify	**Alice: Refuse**
Bob: Testify	-5 -5	-10 0
Bob: Refuse	0 -10	-1 -1

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 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

Prisoner’s dilemma in real life

- Price war
- Arms race
- Steroid use
- [Diner’s dilemma]

	Defect	**Cooperate**
Defect	**Lose** **Lose**	**Lose Big** **Win**
Cooperate	**Win** **Lose Big**	**Draw** **Draw**

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

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: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Network (GAN)

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.

Game of Chicken

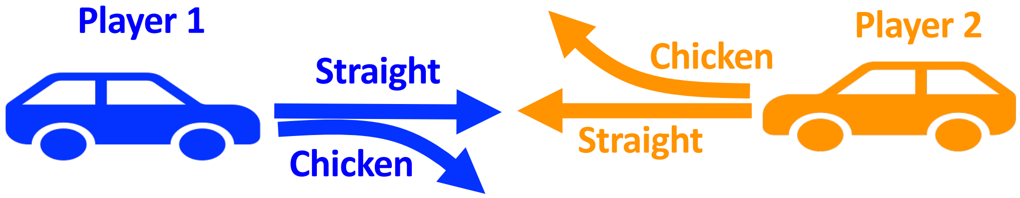

		Alice	
		Straight	Chicken
Bob	Straight	-10 -10	-1 1
	Chicken	1 -1	0 0

- 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

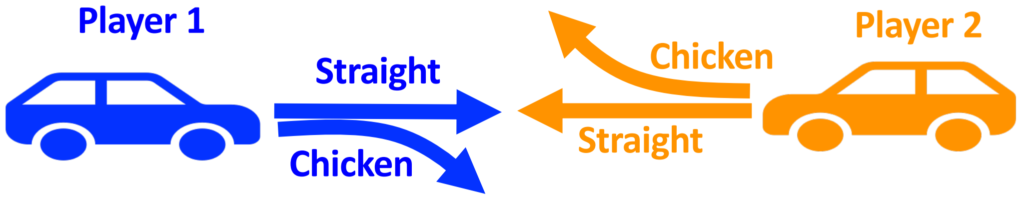

		Alice	
		Straight	Chicken
Bob	Straight	-10 -10	-1 1
	Chicken	1 -1	0 0

- 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?
  - 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

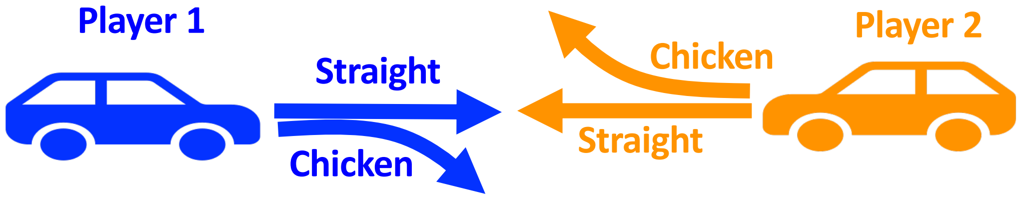

		Alice	
		Straight	Chicken
Bob	Straight	-10 -10	-1 1
	Chicken	1 -1	0 0

- 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?
  - 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:

- Bluff. One player convinces the other that he will behave irrationally. The other player concedes the game. Result: (straight,chicken) or (chicken,straight).
- 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.

Game of Chicken

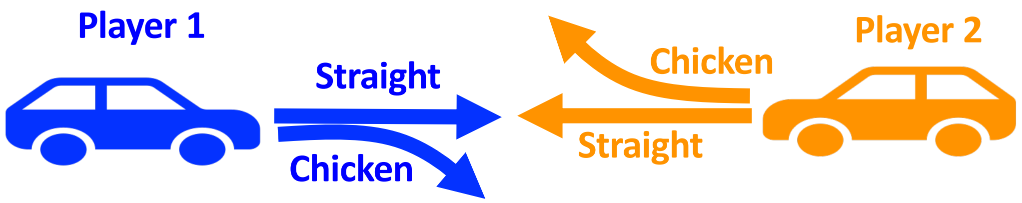

		Alice	
		Straight	Chicken
Bob	Straight	-10 -10	-1 1
	Chicken	1 -1	0 0

- **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

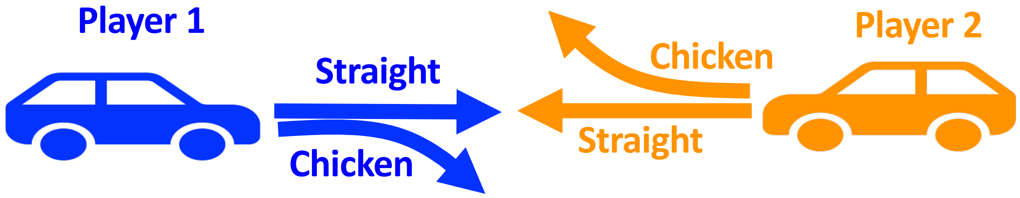

		Alice	
		Straight	Chicken
Bob	Straight	-10 -10	-1 1
	Chicken	1 -1	0 0

The expected payoff, to Alice, for choosing to go Straight is:

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

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

The expected payoff, to Alice, for choosing to Chicken Out is:

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

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

So Alice has no preference between actions S and C. Therefore, it is rational for her to choose between the two actions in any arbitrary way, e.g., using a random number generator.

Finding mixed strategy equilibria

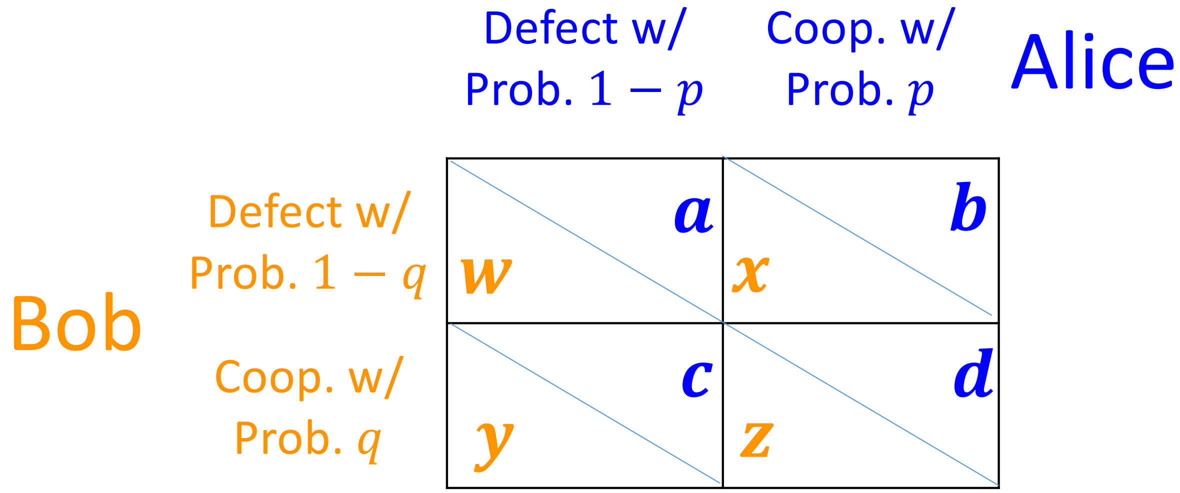

		Alice	
		Defect w/ Prob. $1 - p$	Coop. w/ Prob. p
Bob	Defect w/ Prob. $1 - q$	a w	b x
	Coop. w/ Prob. q	c y	d z

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

Alice

Bob

	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	$\boldsymbol{c}$ $\boldsymbol{y}$	$\boldsymbol{d}$ $\boldsymbol{z}$

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 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).

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

- Games with simultaneous moves: Notation
- Example: Stag Hunt (Coordination Games)
  - Nash Equilibrium: Each player knows what the other will do, and responds rationally
- Example: Asymmetric Coordination 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)
  - Factors external to the game: How well can you bluff?
  - Rational action within the game: Mixed Nash Equilibrium
- Example: Generative Adversarial Network (GAN)

Using a neural net to generate synthetic data

- A neural network can be trained to generate images, speech, text…
- The usual training criterion: mean-squared error. If y is the target vector, network $\hat{y}$ in order to minimize $E\left[\left\|\hat{Y}-Y\right\|^2\right]$.
- The minimum-MSE solution is $\hat{y}=E[Y|X=x]$
- But what if you don't want the network to ALWAYS generate the AVERAGE vector? What if you want a natural variety of different outputs?

Generative Adversarial Network

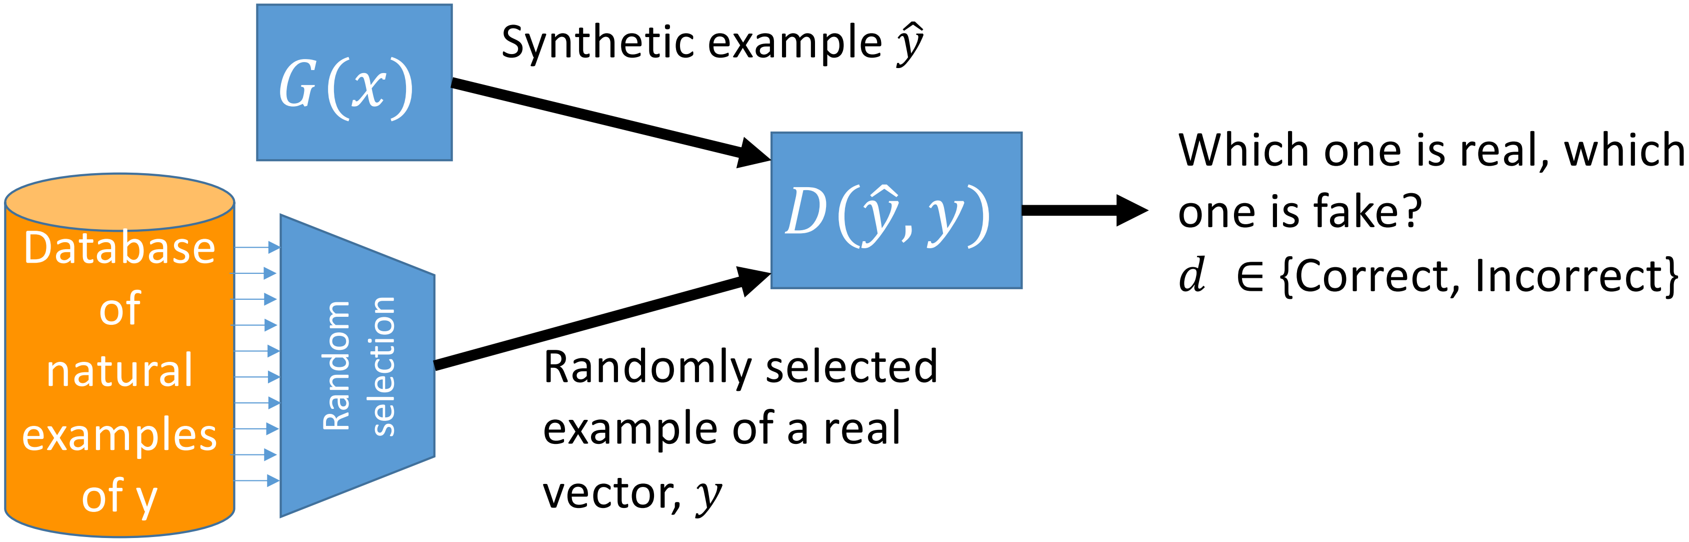

Goodfellow et al. proposed training TWO networks:

- GENERATOR: $\hat{x} = G(z)$ generates synthetic data
- DISCRIMINATOR: $d = D(\hat{y}, y)$ is presented with $\hat{y}$, and with a randomly chosen real data vector, y, and decides which of the two is synthetic.

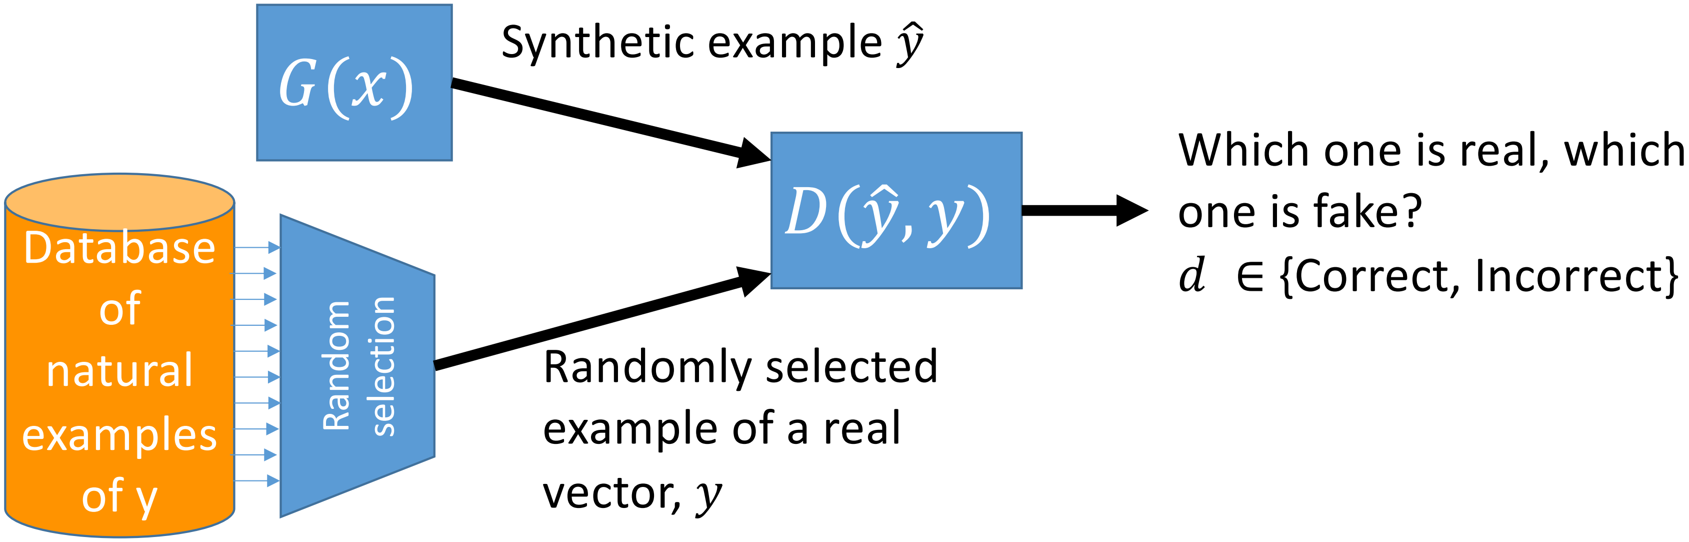

Synthetic example $\hat{y}$

Which one is real, which one is fake?
d ∈ {Correct, Incorrect}

Randomly selected example of a real vector, y

The discriminator is trained to maximize accuracy.

$$D = \text{argmax} \frac{1}{n} \sum_{i=1}^{n} 1[d_i = \text{Correct}]$$

The generator is trained to minimize the discriminator's accuracy

$$G = \text{argmin} \frac{1}{n} \sum_{i=1}^{n} 1[d_i = \text{Correct}]$$

Dominant strategy

Suppose that the natural data distribution is $p_D(Y)$. Suppose the generator produces examples with a distribution $p_G(Y)$. Then the discriminator has a dominant strategy:

- If it sees an example for which $p_D(Y) > p_G(Y)$, call it “natural.”
- If it sees an example for which $p_D(Y) < p_G(Y)$, call it “synthetic.”
- If is sees an example for which $p_D(Y) = p_G(Y)$, make a random decision.

Reminder: 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.

Generator’s best response to the Discriminator’s dominant strategy

- If the discriminator sees any example for which $p_D(Y) > p_G(Y)$, then it can get a score of better than 0.5 by calling that example “natural.”
- If the discriminator sees any example for which $p_D(Y) < p_G(Y)$, then it can get a score of better than 0.5 by calling that example “synthetic.”
- If the discriminator only sees examples for which $p_D(Y) = p_G(Y)$, then the discriminator will be forced to accept a score of 0.5.
- Therefore, the **best response** is to generate a random assortment of vectors, $\hat{y}$, so that their distribution exactly matches the true data distribution, i.e., $p_D(Y = \hat{y}) = p_G(Y = \hat{y})$.
- This can be done, for example, by generating a Gaussian random vector x, and then computing its transformation $\hat{y} = G(x)$ so that $p_D(Y = \hat{y}) = p_G(Y = \hat{y})$.

Nash equilibrium for the generative adversarial network

- Generator produces a random assortment of vectors, $\hat{y}$, with a distribution that exactly matches the true data distribution, i.e., $p_D(Y = \hat{y}) = p_G(Y = \hat{y})$.
- Discriminator has no choice but to accept a correctness score of 50%.

Example

Figure 2. Images generated by a VAE and a DCGAN. First row: samples from a VAE. Second row: samples from a DCGAN.

From Yeh et al., “Semantic Image Inpainting with Deep Generative Models.”

- VAE (Variational AutoEncoder) is producing $\hat{y} = E[Y|X = x]$ from a Gaussian random vector x.
- DCGAN is transforming x so that $p_D(Y = \hat{y}) = p_G(Y = \hat{y})$.

Summary

- Prisoner’s Dilemma
  - Nash equilibrium = both players play their dominant strategy
  - Nash equilibrium $\notin$ 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

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

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