

# Today

A2 out last week, due Friday, Nov 21

# Today

**Game Theory** 

## First: a game!

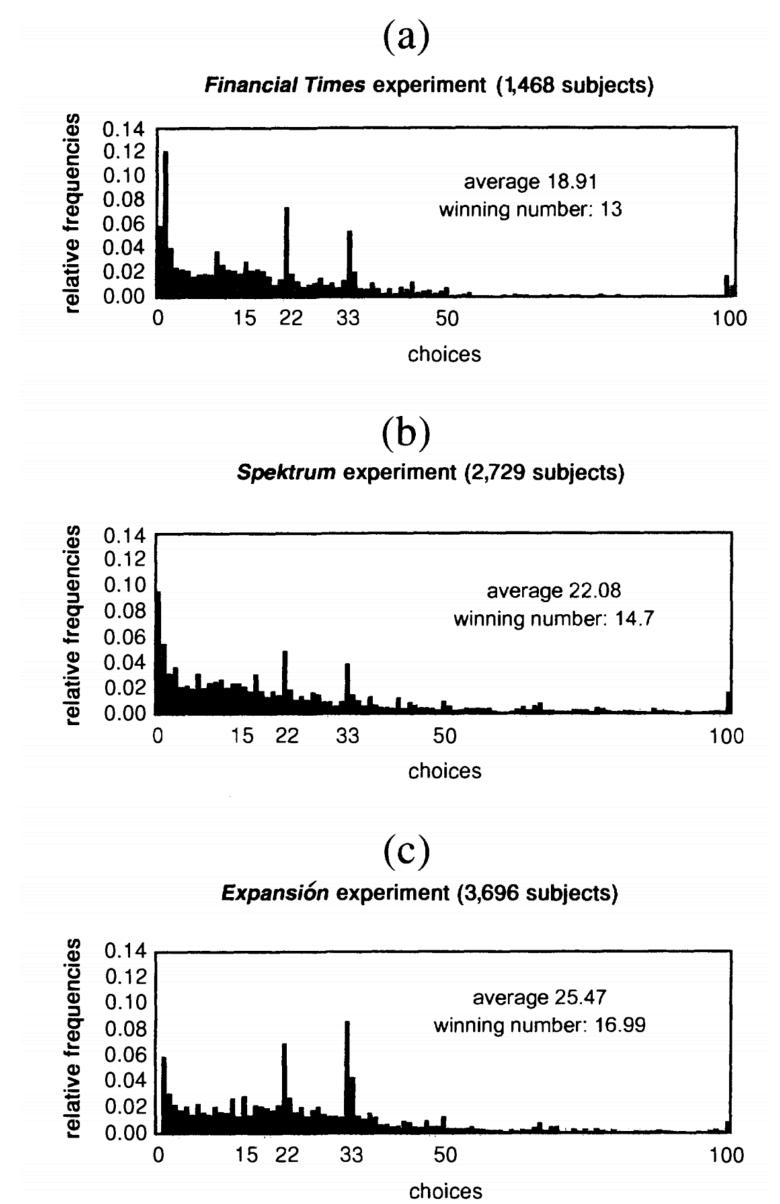
Everyone will guess a number between 0 and 100 (inclusive), and whoever's number is closest to 2/3 of the average guess will win!

No speaking

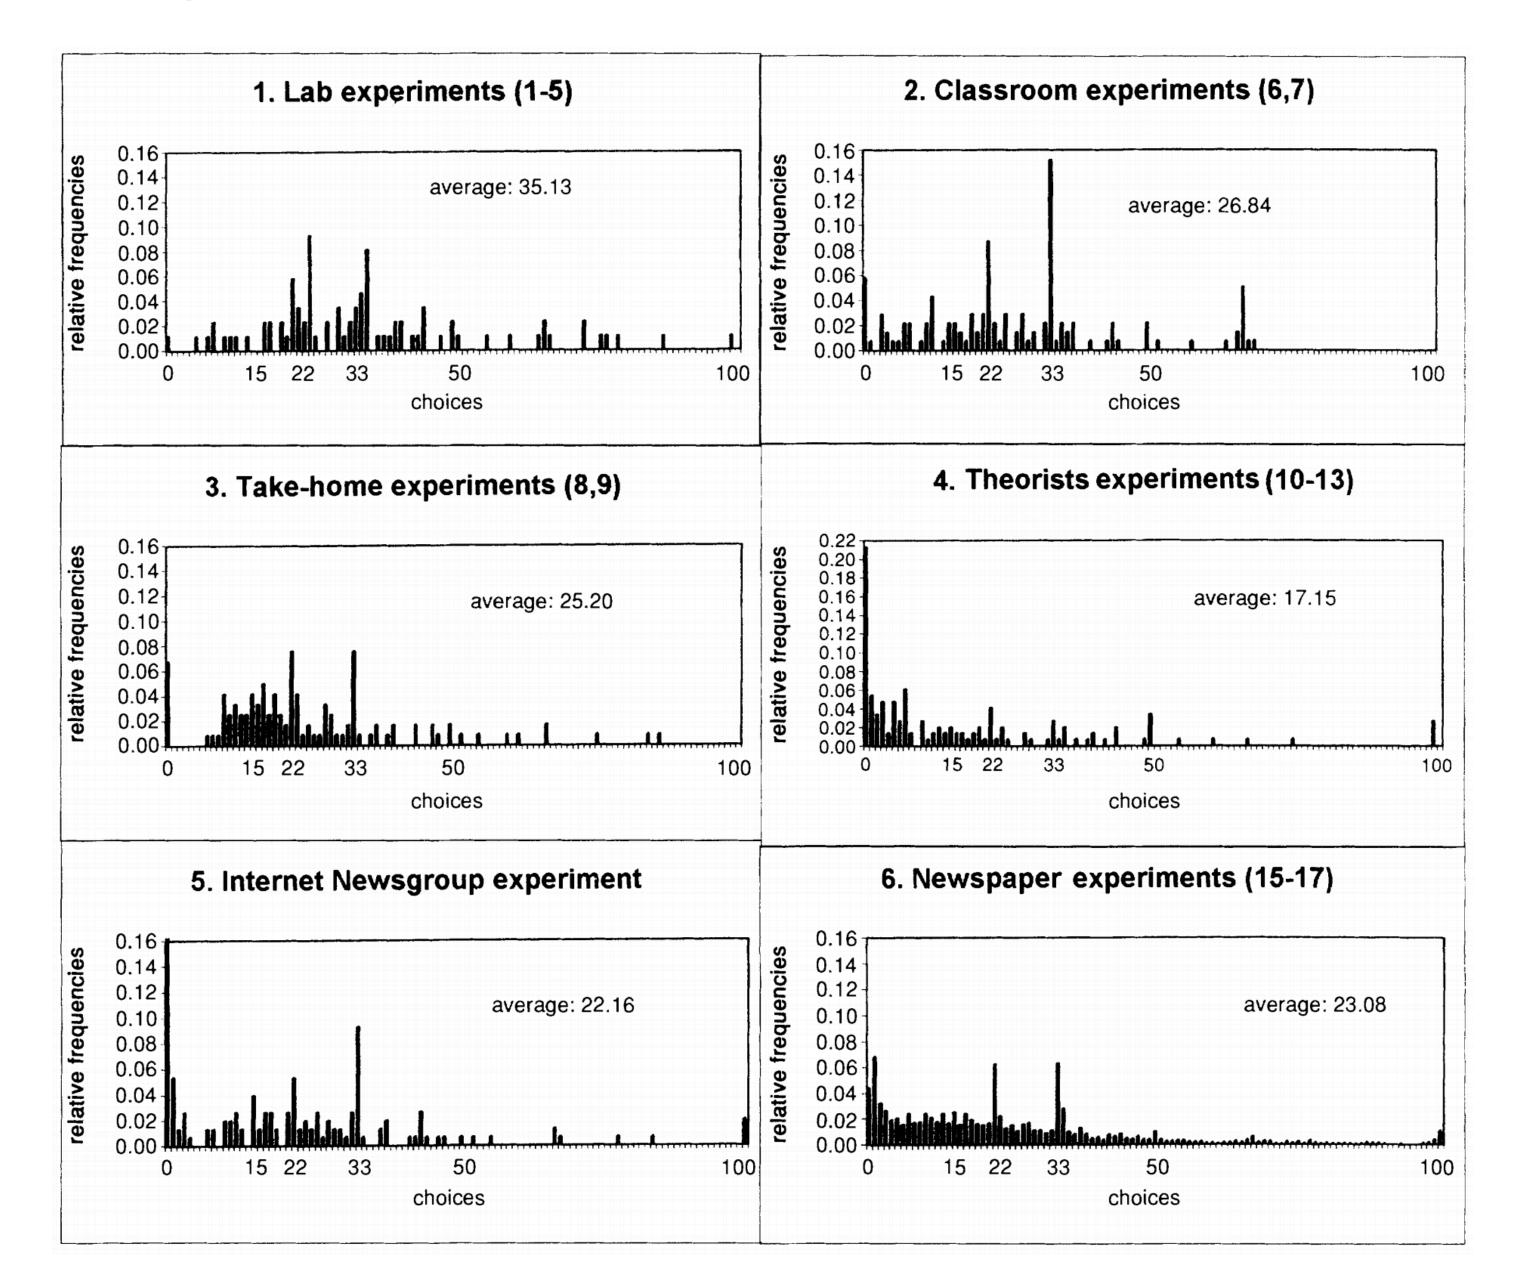
Write down your UTORid along with a single guess

http://tiny.cc/c46game

### "Beauty contest" experiment in newspapers



### "Beauty contest" experiment in newspapers

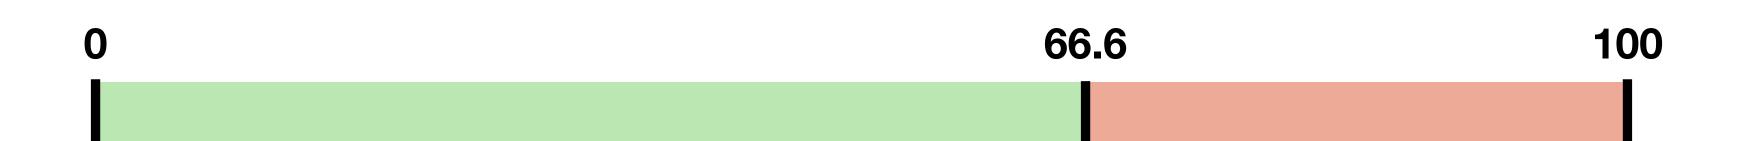


## What is "rational" play?

Assume everyone is rational ("common knowledge of rationality")

Notice: anything between 66.7 and 100 can never win!

Even if everyone guessed 100, 100\*2/3 = 66.6, so 66.6 is a better guess than anything above it



## What is "rational" play?

What now?

#### 66.6 is the new 100!

By the same reasoning, if everyone is rational, no one will guess above 66.6

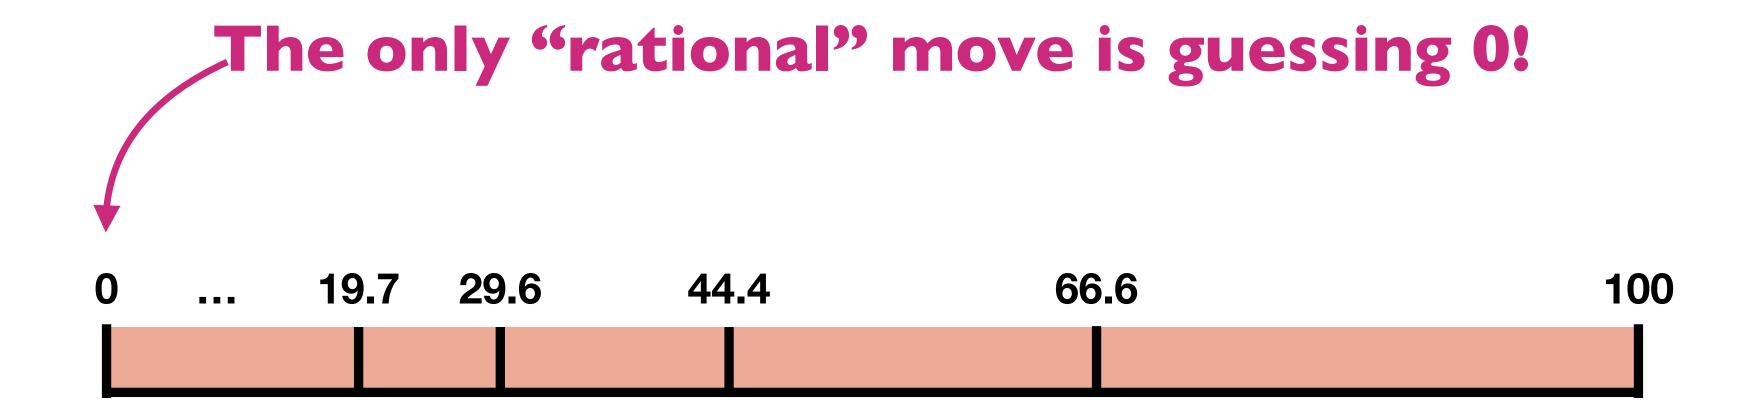
If that's true, then a rational person should never guess anything between 44.4 and 66.6



## What is "rational" play?

Repeat!

44.4 is the new 66.6, and so on

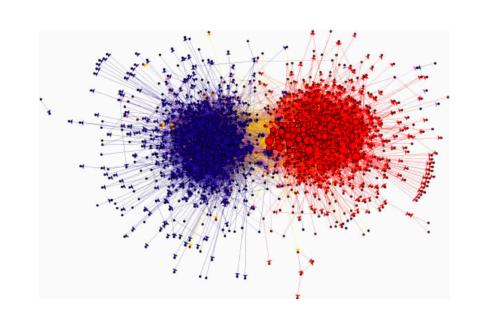


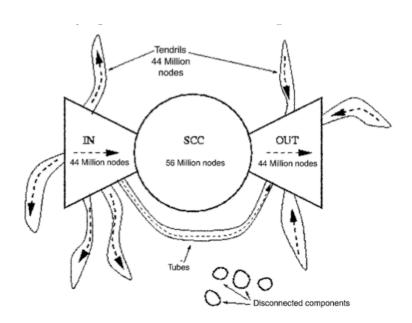
(of course, in real life not everyone is rational)

## Today: Game theory

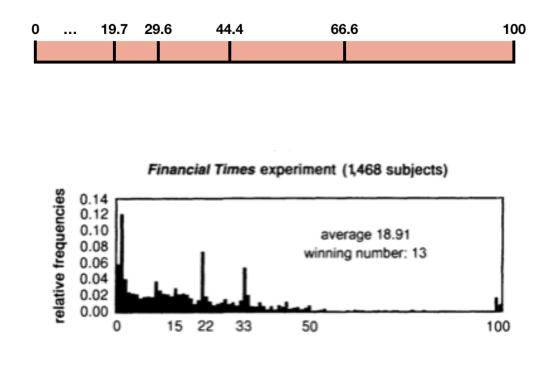
#### Networks: interconnected structure

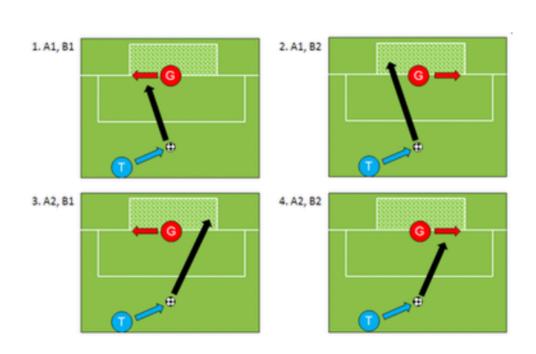


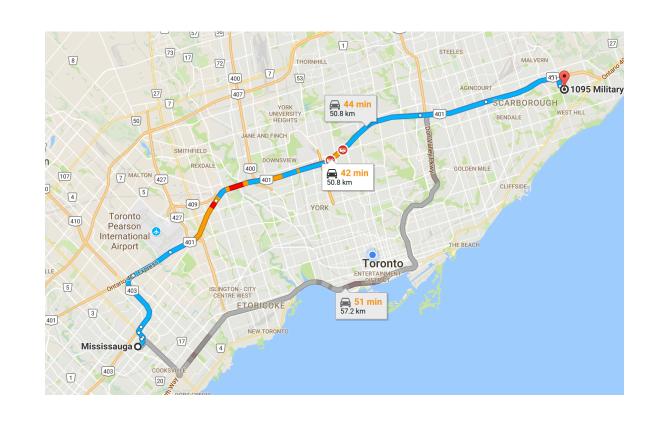




#### Game theory: interconnected behaviour







### **Exam or Presentation?**

A class has two grades: individual exam and a two-person presentation

- Overall grade is the average of your exam and your presentation
- Can't fully prepare for both (sound familiar?)

#### Exam:

- If you study for the exam you'll do well (92%)
- If you don't study then you'll do less well (80%)

[And same for your partner!]

#### Presentation:

- If you both prepare for the presentation you'll do extremely well (100%)
- If just one person prepares then medium (92%), if neither then bad (84%)

What should you do?

### **Exam or Presentation?**

We can summarize the situation in a 2x2 table Your choices are the rows, and your partner's choices are the columns Each box gives the grades: first you, then your partner

Both work on presentation: Avg(100,80)=90

One works on presentation: Avg(92,80)=86, other studies for exam: Avg(92,92)=92

Both study for exam: Avg(84,92)=88

Your score depends not only on your choice but your partner's choice too!

#### Your Partner

	Presentation	Exam
n	90,90	86, 92
n	92,86	88,88

You Presentation Exam

### **Exam-Presentation Game**

#### What should you do?

If you knew your partner would study for the **exam**, what should you do? You'd choose **exam** (88 > 86)

If you knew your partner would work on the **presentation**, what should you do?

You'd choose exam (92 > 90)

No matter what, you should choose exam!

		Your Partner			
		Presentation	Exam		
You $P7$	resentation	90, 90	86,92		
10u	Exam	92,86	88,88		

### **Exam-Presentation Game**

The situation is totally symmetric for your partner, they should choose the **exam** no matter what too

But you'd both be better off preparing for the presentation!

<b>T</b> 7	
Volum	Partner
YOUR	

			Presentation	Exam
$V_{\Omega 11}$	$P_{1}$	resentation	90, 90	86,92
10u		Exam	92,86	88,88

### **Basic Definitions**

Players: you and your partner

Strategies: prepare presentation or study for final

Payoff: grade as a function of everyone's strategy

Payoff matrix: see below

This is a **game** (as in game theory)

Played once, and players select strategies simultaneously and without consulting one another

#### Your Partner

### **Basic Definitions**

#### A game G is a tuple (P,S,O):

**P** = set of Players

**S** = a set of strategies for every player

• = for every outcome (where every player is picking one strategy), a payoff for each player

Payoff matrix summarizes all of these (each dimension is a player, every row/column/etc is a strategy for one player, every cell expresses payoffs for each player)

## Underlying Assumptions

Payoffs summarize everything a player cares about

Every player knows everything about the structure of the game: who the players are, the strategies available to everyone, payoffs for each player and strategy

Every player is **rational**: wants to maximize payoff and succeeds in doing so

Your Partner

## Underlying Assumptions

Weird conclusions? Assumptions are probably to blame!

Your Partner

 $You \ {\begin{tabular}{c} Presentation \\ Exam \end{tabular}}$ 

Presentation	Exam
90, 90	86, 92
92,86	88,88

Two bank robbery suspects are held in separate chambers

Not enough evidence to convict them, but they resisted arrest



Two bank robbery suspects are held in separate chambers. Not enough evidence to convict them, but they did resist arrest

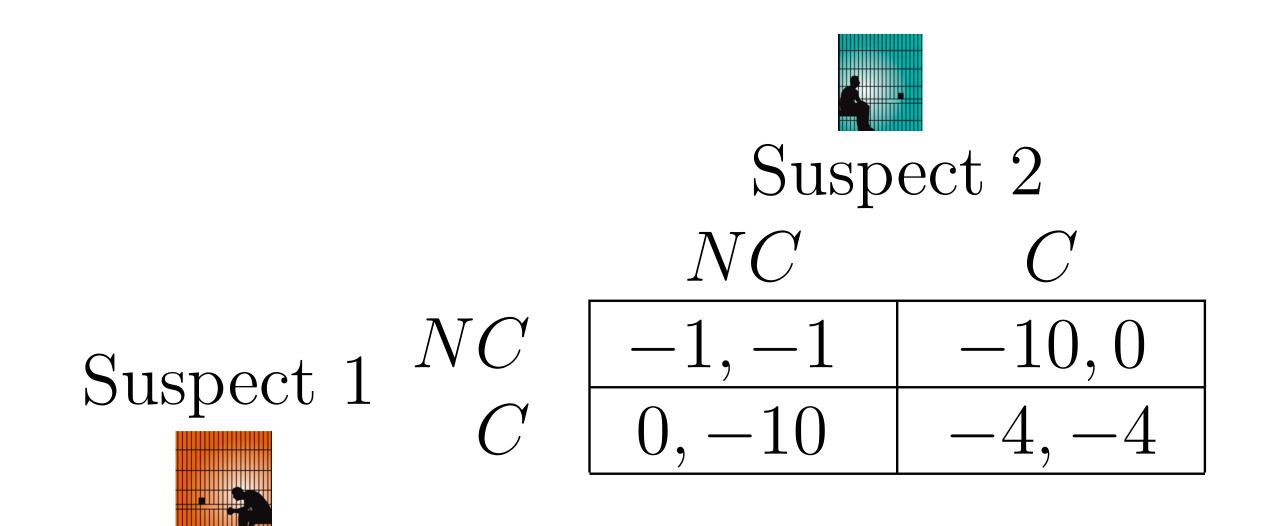


Police take both aside **separately**, and tell each one:

- If you confess, and your partner doesn't confess:
  - You will be released
  - Your partner will be sent to prison for 10 years
- If you both confess, then we don't need either of you to testify against the other, and:
  - You will both be convicted of the robbery
  - Both serve 4 years in prison
- Finally, if neither of you confesses, then we can't convict either of you of the robbery:
  - Both charged with resisting arrest only (I year in prison)
- Your partner is being offered the same deal. Do you want to confess?"

We can represent this situation in a simple matrix: Suspect I's choices are the rows, and Suspect 2's choices are the columns (Confess and Not-Confess)

Each box gives the outcomes: first Suspect 1, then Suspect 2



Similar situation! Confessing is best for both suspects

Suspect 2
$$NC = \frac{NC}{C}$$
Suspect 1
$$C = \frac{NC}{C}$$

$$C = \frac{1}{C}$$

$$C = \frac{1}{C}$$

$$C = \frac{1}{C}$$

$$C = \frac{1}{C}$$

Compare with exam vs. presentation game:

$Presentation \ Exam$	
$Y_{\text{ou}}$ Presentation 90,90 86,92	
Exam 92, 86 88, 88	

### Fundamental Concepts: Strict Dominant Strategy

A strategy that is strictly better than all other options, regardless of what other players do

Exam is a **strictly dominant strategy** for both players Sadly, (90,90) is not achievable with rational play Even if you could commit to preparing for the presentation, your partner would still be better off studying for the final

		Your Partner			
		Presentation	Exam		
You	Presentation	90, 90	86, 92		
IOU	Exam	92,86	88,88		

Varia Daration

### Prisoner's Dilemma in the Real World

Drug doping in professional sports (dope vs. don't dope)

Arms races between countries (build arms vs. don't)

Countries respecting climate change treaties (Do or don't restrict CO2 emissions)

Overfishing (do or don't overfish the seas)

Advertising (advertise or don't)

		Suspect 2			
		NC	C		
Suspect 1	NC	-1, -1	-10, 0		
Duspect 1	C	0, -10	-4, -4		

### Practice Question

Recall the game Rock-Paper-Scissors (paper beats rock, scissors beat paper, rock beats scissors)

Representing win/draw/loss as +1/0/-1, express Rock-Paper-Scissors as a game theory game

### Practice Question

Recall the game Rock-Paper-Scissors (paper beats rock, scissors beat paper, rock beats scissors)

Representing win/draw/loss as +1/0/-1, express Rock-Paper-Scissors as a game theory game

Player 2
----------

	P1\P2	Rock	Paper	Scissors
	Rock	0,0	-1,+1	+1,-1
Player 1	Paper	+1,-1	0,0	-1,+1
	Scissors	-1,+1	+1,-1	0,0

Let's define some more of the fundamental concepts we just used

Best response is just what it sounds like: if player 2 plays **T**, then the best thing I can do is play **S** 

Suspect 2 NC CSuspect 1 NC -1, -1 -10, 0 0, -10 -4, -4

S1's best response to NC is: ? S1's best response to C is: ?

Let's define some more of the fundamental concepts we just used

Best response is just what it sounds like: if player 2 plays **T**, then the best thing I can do is play **S** 

Suspect 2 NC CSuspect 1 NC -1, -1 -10, 0 0, -10 -4, -4

S1's best response to NC is: C S1's best response to C is: C

Let's define some more of the fundamental concepts we just used Strategy  $\bf S$  by  $P_1$  is a **best response** to strategy  $\bf T$  by  $P_2$  if the payoff from  $\bf S$  as at least as good as anyone other strategy against  $\bf T$ 

$$P_1(S,T) \ge P_1(S',T)$$
 for all other S' by  $P_1$ 

It's a strict best response if:

$$P_1(S,T) > P_1(S',T)$$
 for all other S' by  $P_1$ 

S1's best response to NC is: C S1's best response to C is: C

P1\P2	A	В	C	D	E
A	3, 5	-2, 1	4,3	1,6	9,2
В	2,2	1,10	3,6	4,2	5,3
C	8,-1	-2,6	-3,1	9,2	1,3

### Fundamental Concepts: Dominant Strategy

A dominant strategy for  $P_1$  is a strategy that is a best response every strategy by  $P_2$ 

A strict dominant strategy for  $P_1$  is a strategy that is a strict best response every strategy by  $P_2$ 

(Note: In Prisoner's Dilemma, P<sub>I</sub> has a strict dominant strategy, so we expect PI to play it. There can be several dominant strategies, and it'd be unclear which one to expect)

### Fundamental Concepts: Dominant Strategy

P1\P2	A	В	C	D	E
A	3, 5	-2, 1	4,3	1,6	9,2
B	2,2	1,8	3,6	4,9	5,3
C	8,-1	-2,2	-3,1	9,4	1,3

### Dominant Strategies Don't Always Exist

Prisoner's Dilemma was relatively easy to analyze because every player has a strictly dominant strategy

However, dominant strategies don't always exist!

P1\P2	A	В	C
A	3, 5	-2, 1	4,3
В	2,2	1,10	3,6
C	8,-1	-2,6	-3,1

## Marketing Game

Consider a marketing scenario: two firms, Firm 1 and Firm 2

Firm I is more popular and gets 80% of profits when they compete They can each either make an upscale product or a low-priced one 60% of the population prefers a low-priced product

Firm 2
Firm 1

What are the strategies? Payoffs?

## Marketing Game

Consider a marketing scenario: two firms, Firm I and Firm 2

Firm I is more popular and gets 80% of profits when they compete Two strategies each: make an upscale product or a low-priced one? 60% of population prefers a low-priced product

Does Firm I have a dominant strategy? Does Firm 2?

		Firm 2	
		$Low ext{-}Priced$	Upscale
Firm 1	$Low\mbox{-}Priced$	.48, .12	.60, .40
	Upscale	.40, .60	.32, .08

What happens?

### Marketing Game

Notice Firm I has a strictly dominant strategy: go low-priced

Firm 2 does not have a dominant strategy

But since Firm I has a strictly dominant strategy, expect to play it. Firm 2's best response to Low-Priced is to play *Upscale* 

Although we're reasoning in two steps, remember that the game itself is still plays the same way: both firms play their strategies simultaneously Intuitive prediction: Firm I ignores Firm 2, Firm 2 steers clear of directly competing with Firm I

		Firm 2	
		$Low ext{-}Priced$	Upscale
Firm 1	$Low ext{-}Priced$	.48, .12	.60, .40
	Upscale	.40, .60	.32, .08

#### What about no strictly dominant strategies?

What happens when neither player in a two-player game has a strictly dominant strategy?

Need another way to predict what will happen

A more intricate marketing game:

Players: Firm 1, Firm 2

Strategies: Approach client A, B, C

Payoff matrix:

		Firm 2		
		A	B	C
	A	$\boxed{4,4}$	0, 2	0, 2
Firm 1	B	0,0	1,1	0, 2
	C	0,0	0, 2	1,1

#### A Three-Client Marketing Game

Neither firm has a dominant strategy

#### For Firm I:

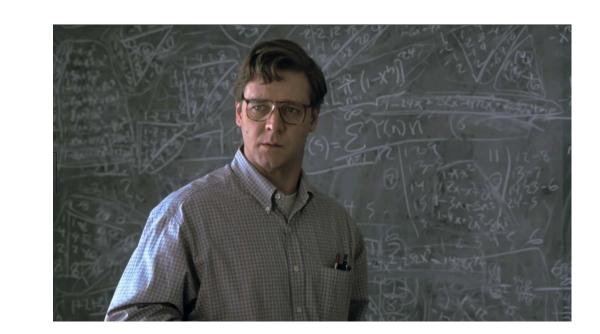
- A is a strict best response to strategy A by Firm 2
- B is a strict best response to B
- C is a strict best response to C

#### For Firm 2:

- A is a strict best response to strategy A by Firm I,
- C is a strict best response to B,
- B is a strict best response to C

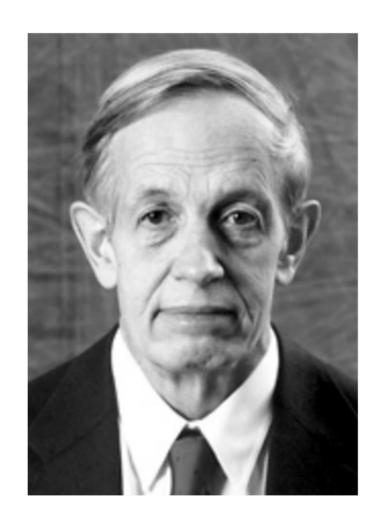
		Firm 2		
		A	B	C
	A	$\boxed{4,4}$	0, 2	0,2
Firm 1	B	0,0	1, 1	0,2
	C	0,0	0, 2	$\boxed{1,1}$

In 1950, John Nash proposed a **simple** and **powerful** principle for reasoning about behaviour in general games (and won the Nobel Prize for it in 1994)



Even when there are no dominant strategies, we should expect players to use strategies that are best responses to each other

A pair of strategies (S,T) is a Nash equilibrium if S is a best response to T and T is a best response to S



Why?

First consider a pair of strategies that **don't** constitute a Nash equilibrium

If both players expected (B,B) as an outcome, would they be happy?

		Firm 2		
		A	B	C
	A	$\boxed{4,4}$	0,2	0,2
Firm 1	B	0,0	1, 1	0,2
	C	0,0	0,2	1, 1

Why?

First consider a pair of strategies that **don't** constitute a Nash equilibrium

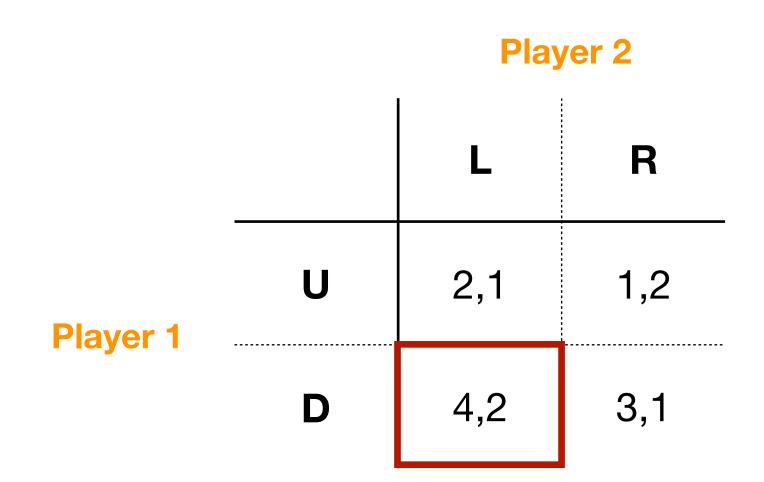
If both firms expected (B,B) as an outcome, would they be happy?

No! Firm 2 would rather play C in response to B.

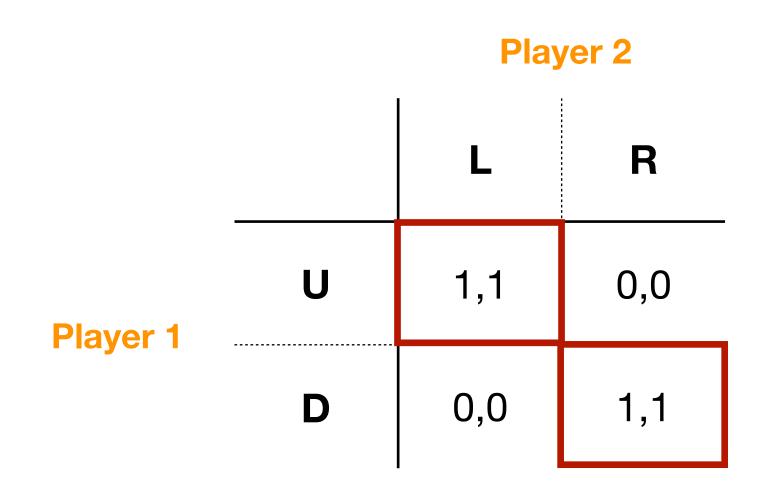
		Firm 2		
		A	B	C
	A	$\boxed{4,4}$	0,2	0,2
Firm 1	B	0,0	1, 1	0,2
	C	0,0	0,2	1, 1

		Player 2		
		L	R	
Player 1	U	1,2	2,3	
	D	2,1	1,2	

		Player 2		
		L	R	
Player 1	U	1,2	2,3	
	D	2,1	1,2	



	Player 2		
		L	R
Player 1	U	1,1	0,0
	D	0,0	1,1



#### Multiple Equilibria

In the case of a single Nash equilibrium, it seems natural to predict that the players will play the strategies in this equilibrium (otherwise someone's not playing a best response)

A lot of games can have more than one equilibrium though

Example: coordination game

Players: you, your partner

Strategies: PowerPoint, Keynote

Payoff matrix:

		Your Partner			
		PowerPoint	Keynote		
$V_{O11}$	PowerPoint	1, 1	0,0		
10u	Key note	0,0	1, 1		

### Multiple Equilibria

This is called a **Coordination game** because all the players care about is playing the same strategy

Lots of coordination games in real life: what side of the street to walk on, what side of the road to drive on, what hand to shake with

		Your Partner			
		PowerPoint	Keynote		
$V_{011}$	PowerPoint	1, 1	0,0		
10u	Keynote	0,0	1, 1		

#### Multiple Equilibria

How does society deal with this?

Sometimes there is a **focal point** that causes the players to focus on one strategy over the others ("it's just the way we do things")

Example: what side of the road to drive on

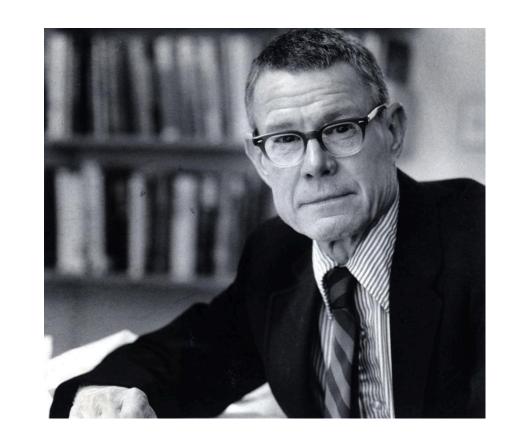
**Social norms, conventions** are often ways of introducing a focal point into coordination games



#### **Unbalanced Coordination Games**

Focal point idea: use a feature intrinsic to the game (rather than an external social convention) to make a prediction

		Your Partner			
		PowerPoint	Keynote		
You	PowerPoint	1, 1	0,0		
10u	Keynote	0,0	2,2		



	Driver 2		
		L	R
Driver 1	L	100,100	-100,-100
	R	0,0	1,1

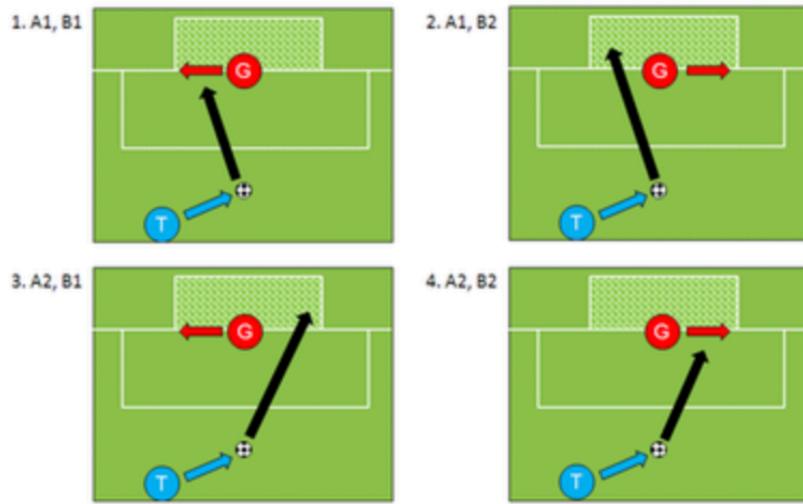
#### **Unbalanced Coordination Games**

But say you and your partner disagree on the best slides software

	Your F	Partner	
	PowerPoint	Keynote	
You PowerPoint	1,2	0,0	("Battles of the Sexes")
Keynote	0,0	2, 1	

# Matching Pennies





### Matching Pennies

Attack-defense structure: interests are in direct conflict

"Zero-sum game"

Players: 1, 2

Strategies: Heads, Tails

Payoff matrix:

Player 2
$$H T$$
Player 1  $H -1, +1 +1, -1$ 

$$T +1, -1 -1, +1$$

What are Nash equilibria of this game?

There are none: no pair of strategies are best responses to each other

#### Mixed strategies

Solution: introduce randomization

Sometimes I'll do this, sometimes I'll do that (randomly)

Intuition: make it harder for my opponent to exploit me

**Strategy:** now corresponds to a choice of mixture probabilities between "pure" strategies.

Payoffs: Expected value under other person's mixture

### Matching Pennies

Players: 1, 2

Strategies:

I: play H probability p

2: play H probability q

Player 2

If P1 chooses p=1 corresponding to pure strategy H: payoff becomes

$$(-1)(q) + (1)(1-q) = 1-2q$$

If P1 chooses p=0 corresponding to pure strategy T: payoff becomes

$$(1)(q) + (-1)(1-q) = 2q - 1$$

Note there pure strategies can't be part of a Nash equilibrium, so p and q must be strictly between 0 and 1

What is Player 1's best strategy to Player 2 choosing q?

Playing H gives him I—2q, and playing T gives him 2q—I

If one was bigger than the other, he should just put all the weight on the bigger one

But no pure strategy Nash equilibrium, so I—2q=2q—I

In any Nash equilibrium, we must have q = 1/2

Similarly for Player I: we must have p = 1/2

Player 2 
$$H \qquad T$$
 Player 1 
$$\frac{H}{T} \begin{bmatrix} -1, +1 & +1, -1 \\ +1, -1 & -1, +1 \end{bmatrix}$$

**Intuitively**: if Player I believes that Player 2 will play H strictly more than T, then she should definitely play T — in which case Player 2 should not be playing H more than half the time.

Make yourself the least exploitable possible

Make the opponent indifferent between their strategies

Player 2 
$$H$$
  $T$ 

Player 1  $H$   $-1, +1$   $+1, -1$   $+1, -1$   $+1, -1$ 

#### Large game-theoretic study of 1400 penalty kicks

Kind of a real-life matching pennies

To make kicker indifferent between shooting L or R, goalie needs to select right q:

$$(.58)(q) + (.95)(1 - q) = (.93)(q) + (.70)(1 - q)$$
  
 $q = 0.42$ 



Amazing fact: goalies dive left exactly 42% of the time!

Every game has a mixed-strategy Nash equilibrium [Nash, 1950]

## Solutions to games

Dominant strategy? Sometimes.

Pure Nash Equilibria? Sometimes.

Mixed Equilibria? Always exists.

#### Mixed Strategies Example: Football

Players: Offense, Defense

Strategies: Run, Pass and Defend Run, Defend Pass

Payoff matrix:

#### Defense

Offense 
$$egin{array}{c|c} Pass & Defend \ Pass \\ Run & 5,-5 & 0,0 \\ \hline \end{array}$$

Mixed Nash:

q = 2/3p = 1/3

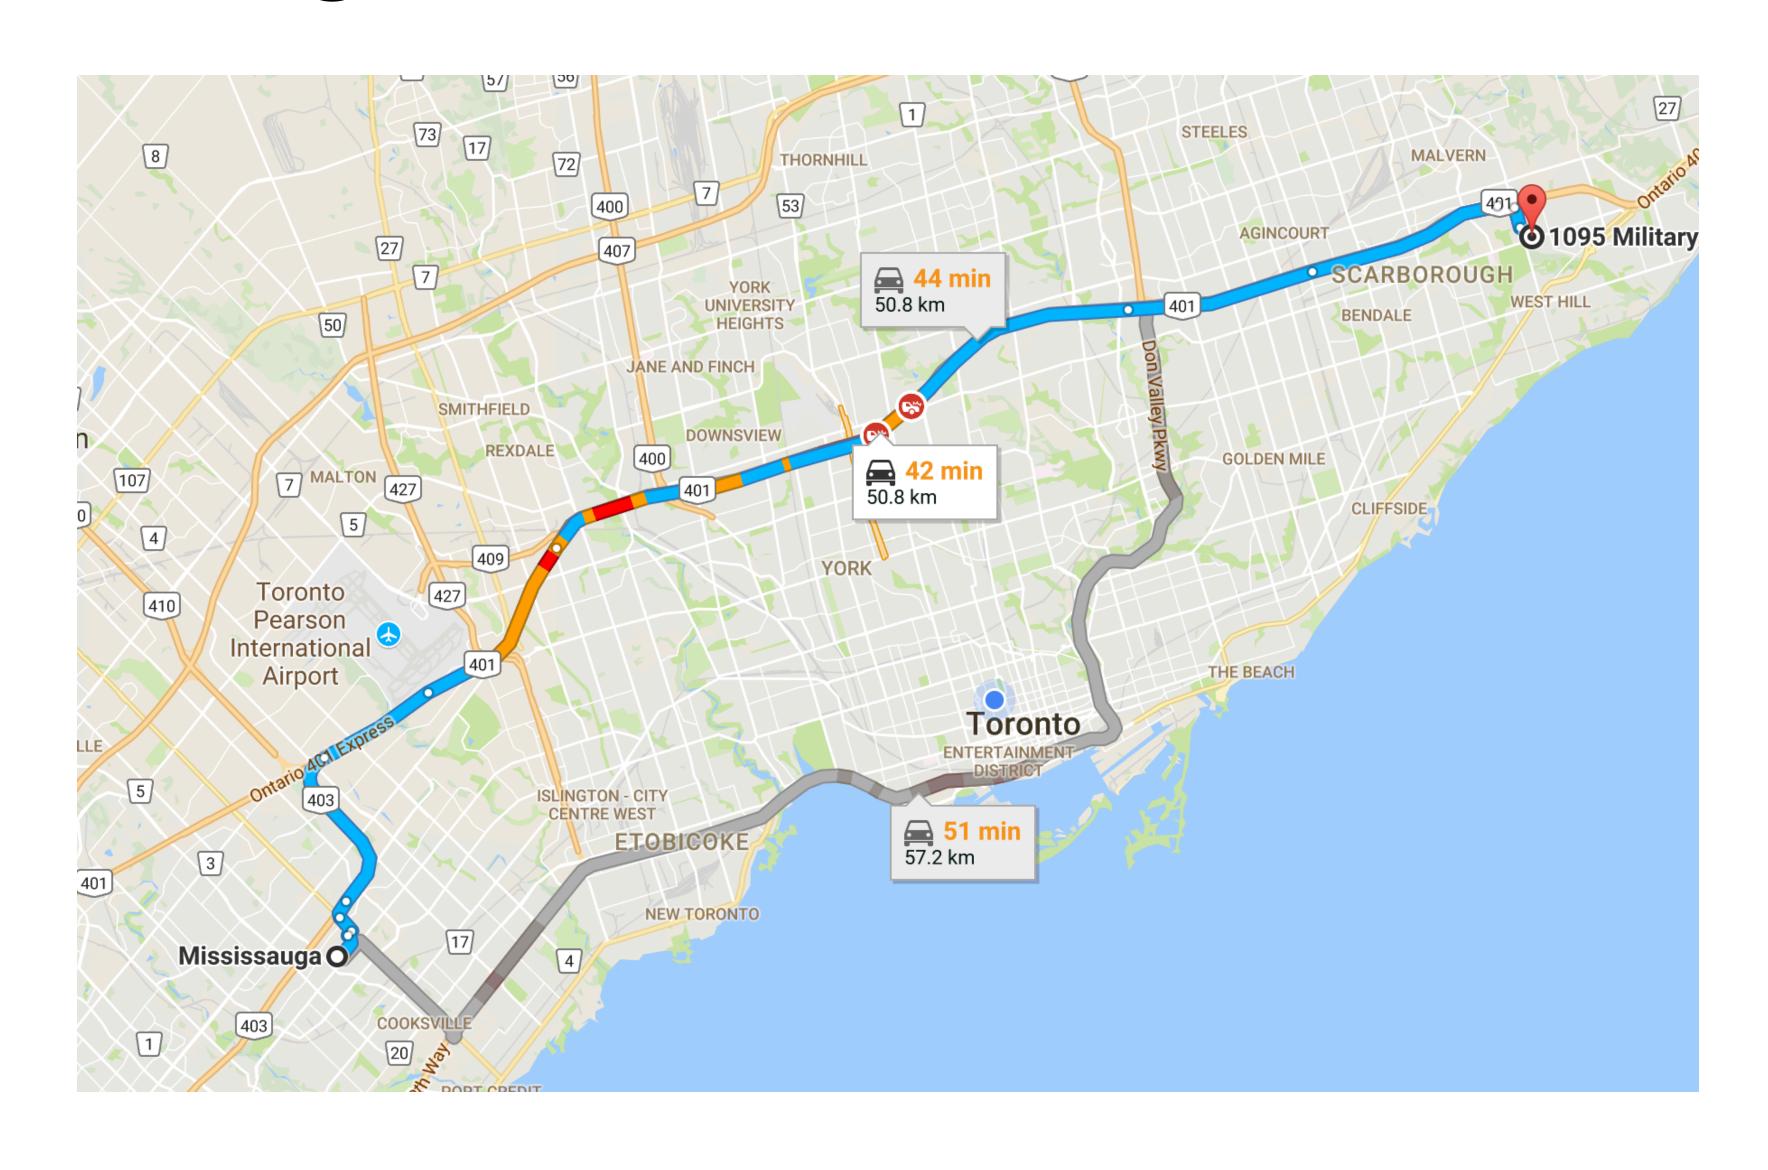
No Nash equilibria in this game

O's expected payoff for **Pass** when D mixes q, l-q: 0\*(q)+10\*(1-q)=10-10q

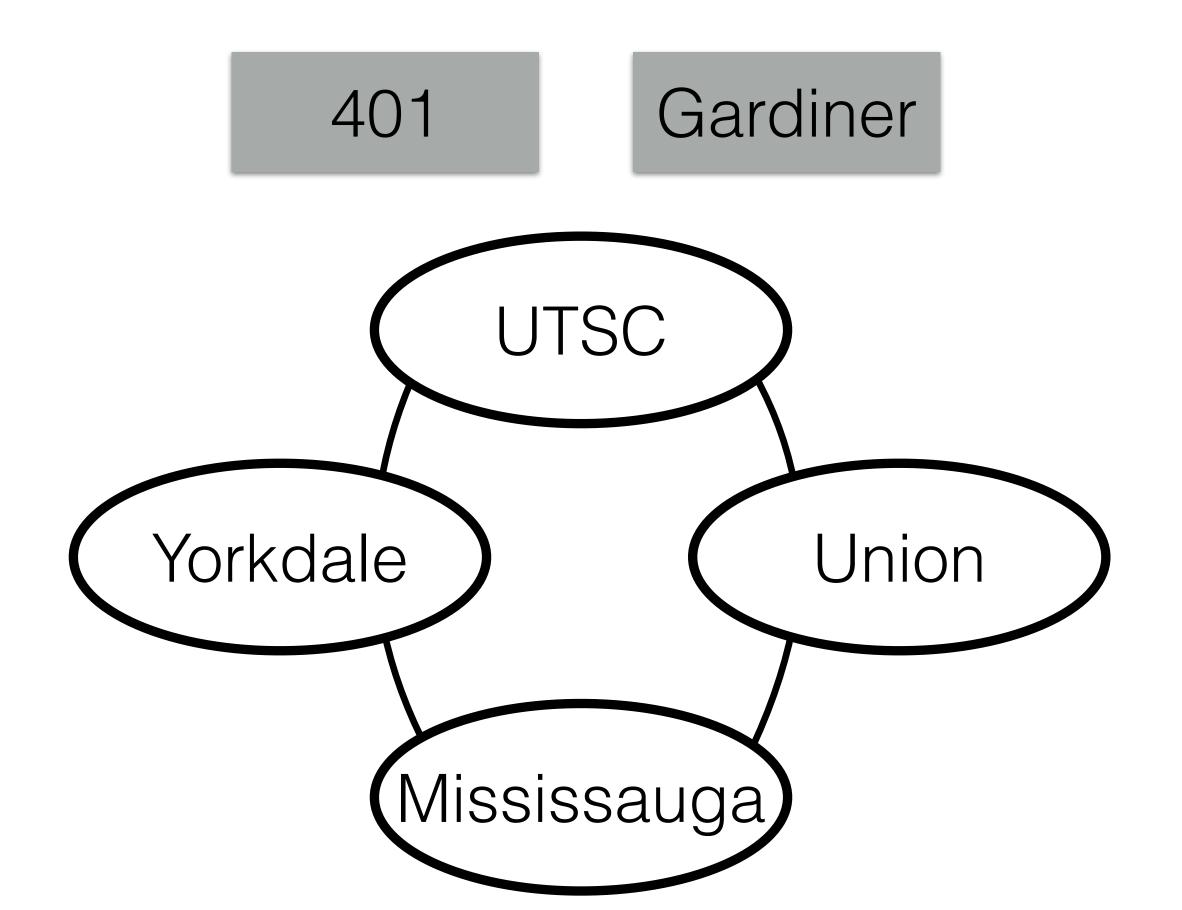
O's expected payoff for **Run** when D mixes q, I-q: 5\*(q)+0\*(I-q) = 5q

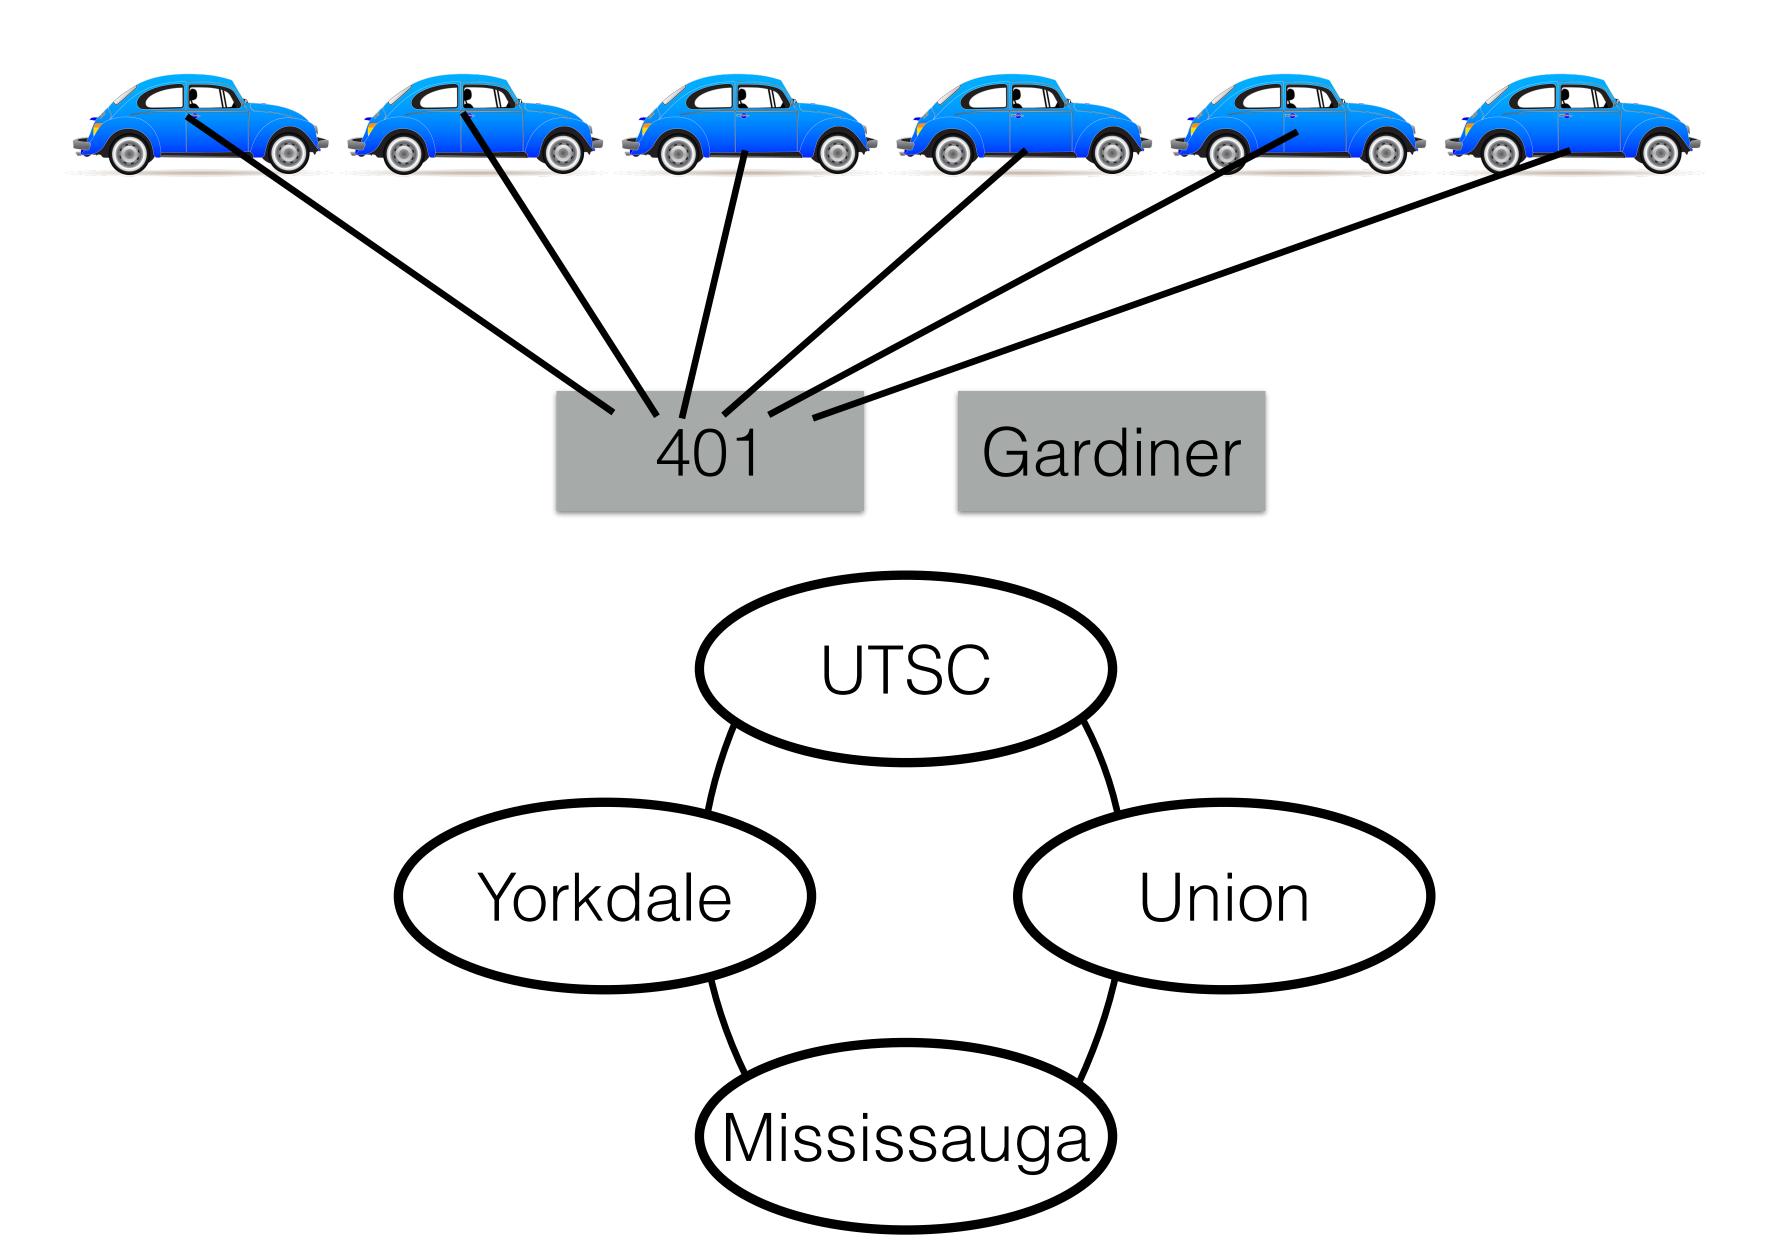
Defense makes Offense indifferent when q=2/3

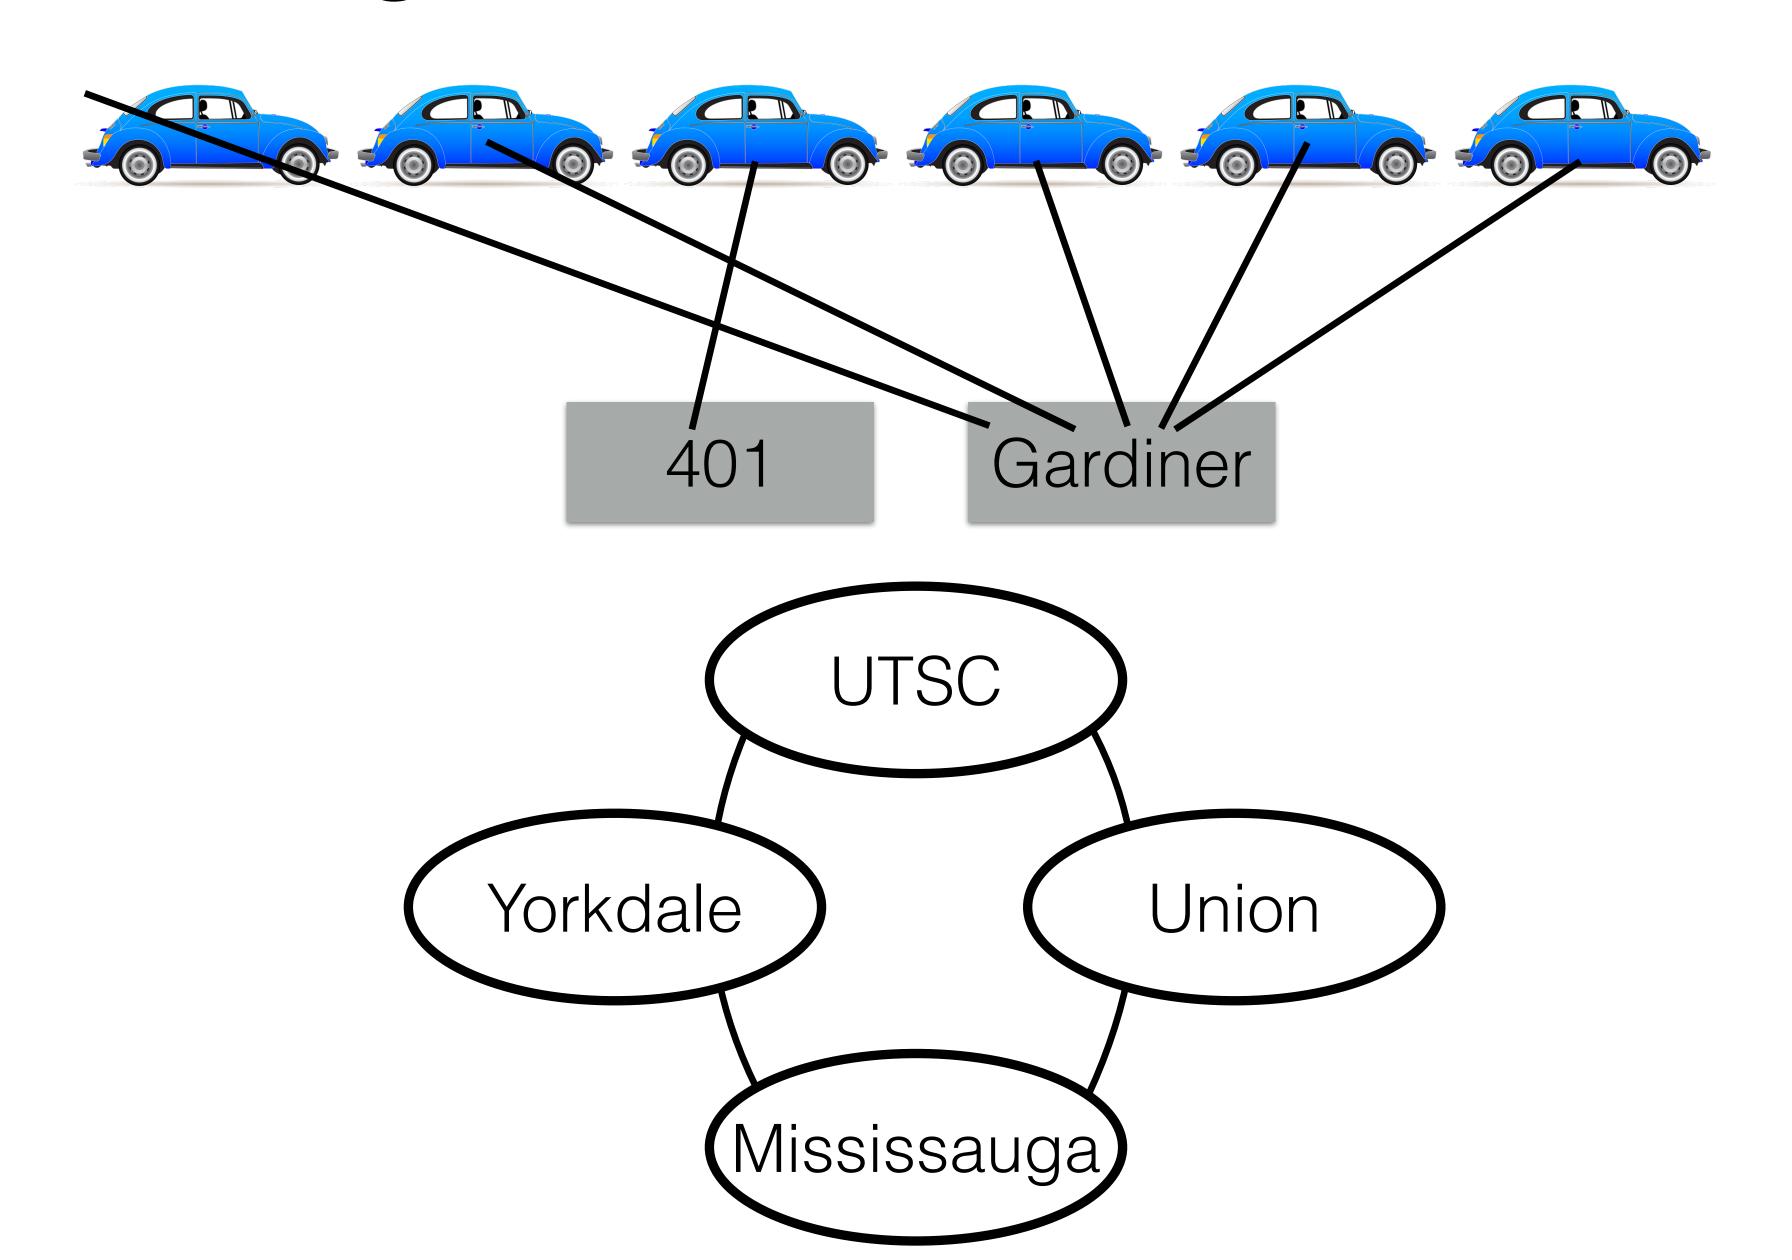
### Game Theory in the World: Traffic

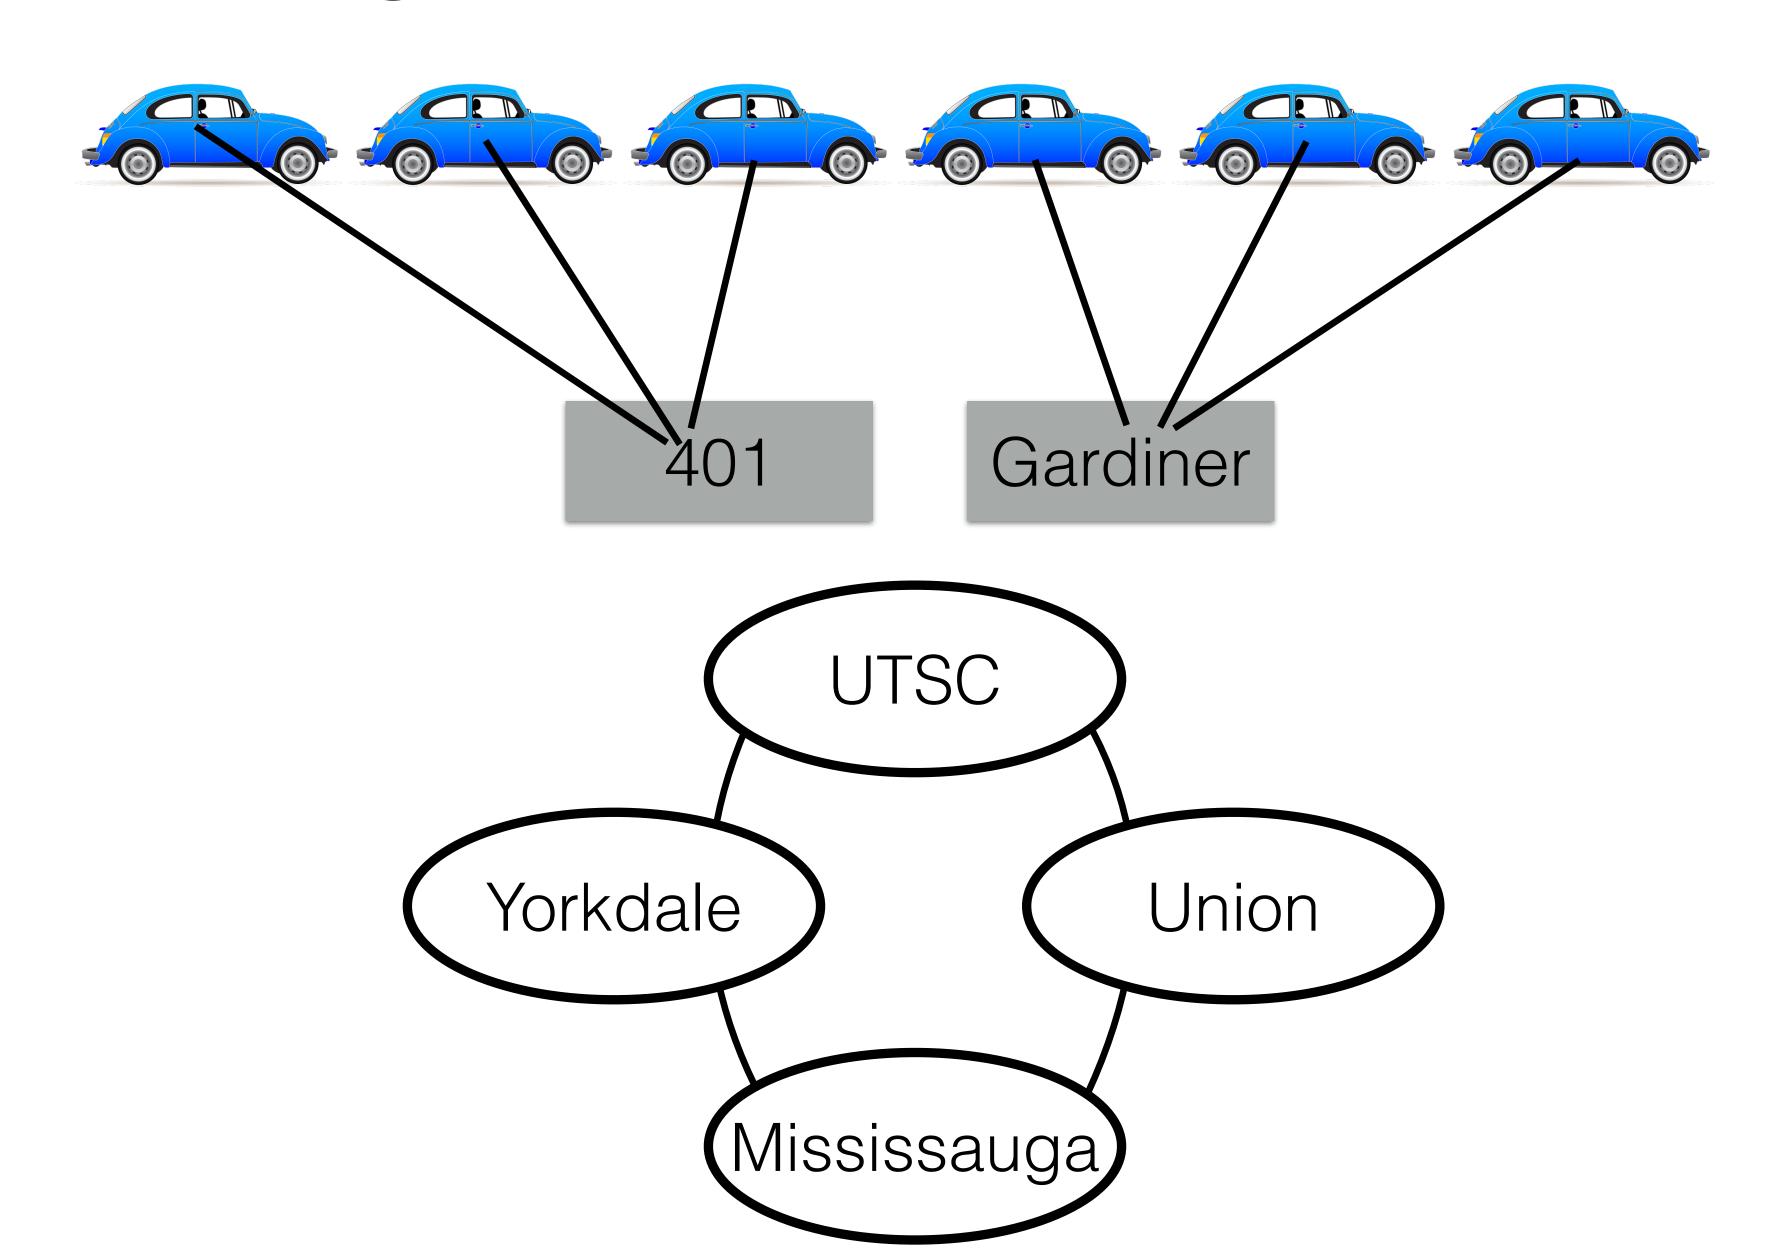






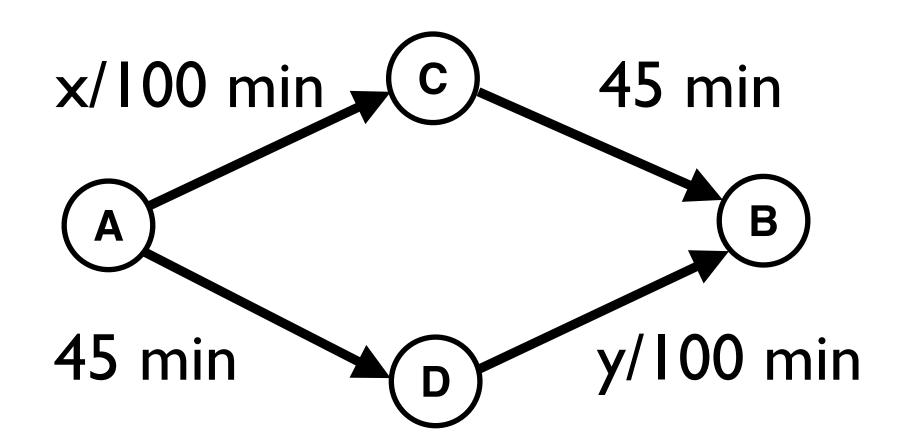






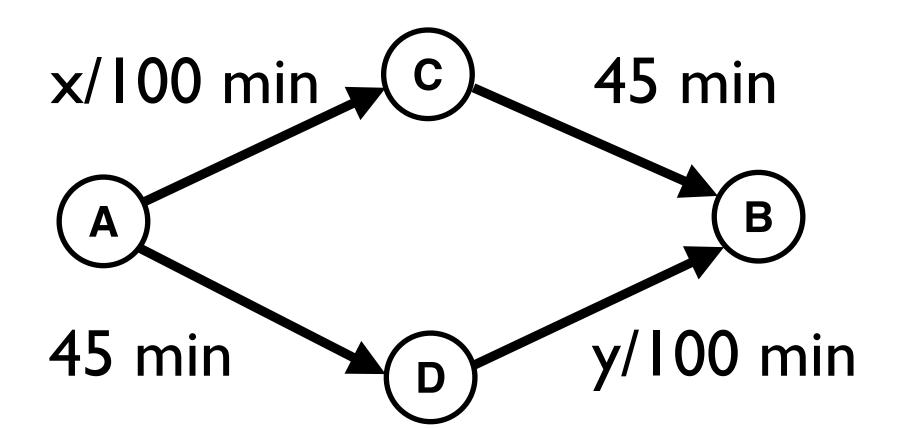
#### Traffic routing

- 4000 drivers
- Two route options: A-C-B or A-D-B.
- 4000 top, 0 bottom:
  - Top: 4000/100 + 45 = 85 min
  - Bottom: 45 + 0/100 = 45 min
- 2000 top, 2000 bottom:
  - Top: 2000/100 + 45 = 65 min
  - Bottom: 45 + 2000/100 = 65 min



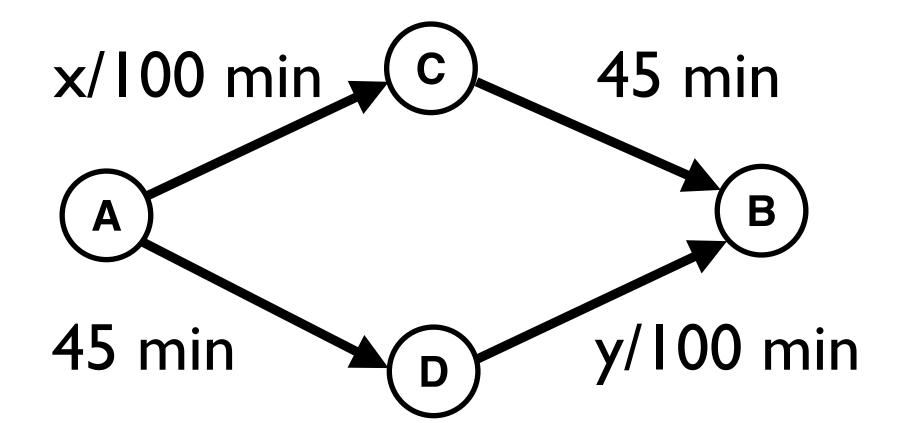
#### Traffic modeled as a game

- **Players:** Drivers 1,2,3...,4000
- Strategies: A-C-B, A-D-B
- Payoffs: ?



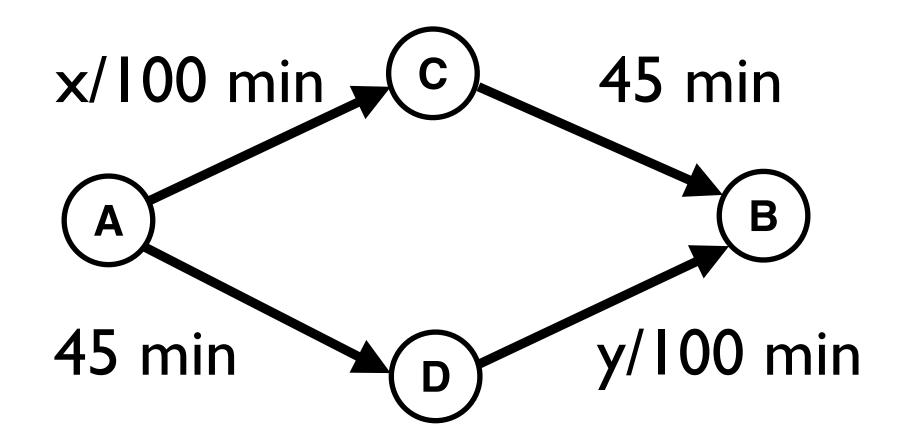
#### Traffic modeled as a game

- **Players:** Drivers 1,2,3...,4000
- Strategies: A-C-B, A-D-B
- Payoffs: Negative drive time
  - A-C-B time: (x/100 + 45)
  - A-D-B time: (45 + y/100)



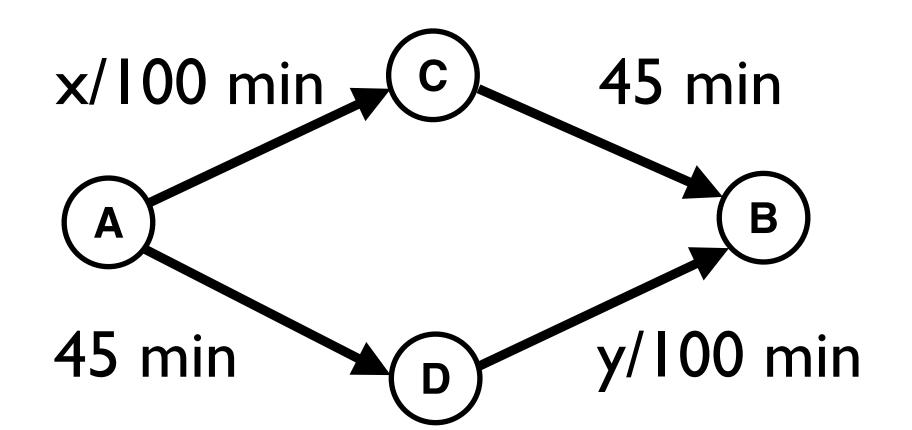
## Equilibrium?

- 4000 drivers
- Two route options: A-C-B or A-D-B.
- 4000 top, 0 bottom:
  - Top: 4000/100 + 45 = 85 min
  - Bottom: 45 + 0/100 = 45 min
- 2000 top, 2000 bottom:
  - Top: 2000/100 + 45 = 65 min
  - Bottom: 45 + 2000/100 = 65 min



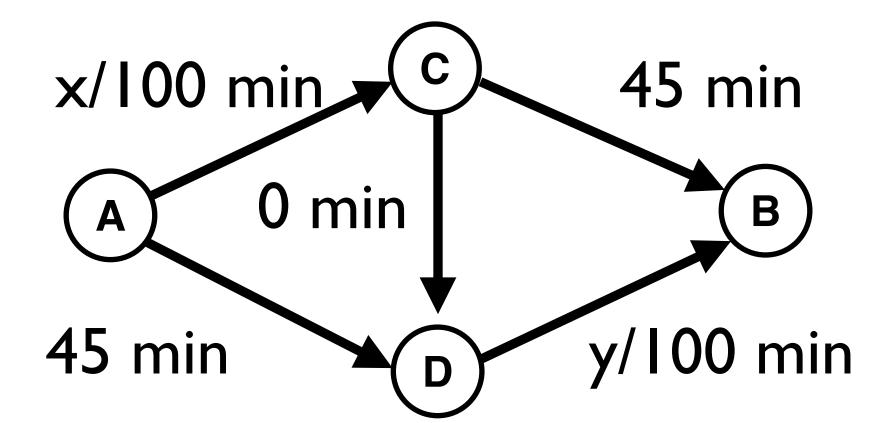
#### Traffic modeled as a game

- **Players:** Drivers 1,2,3...,4000
- Strategies: A-C-B, A-D-B
- Payoffs: Negative drive time
  - A-C-B time: (x/100 + 45)
  - A-D-B time: (45 + y/100)
- Any set of strategies "2000 choose top, 2000 choose bottom" is an equilibrium. Why?
- Any other set of strategies, deviation benefits someone.



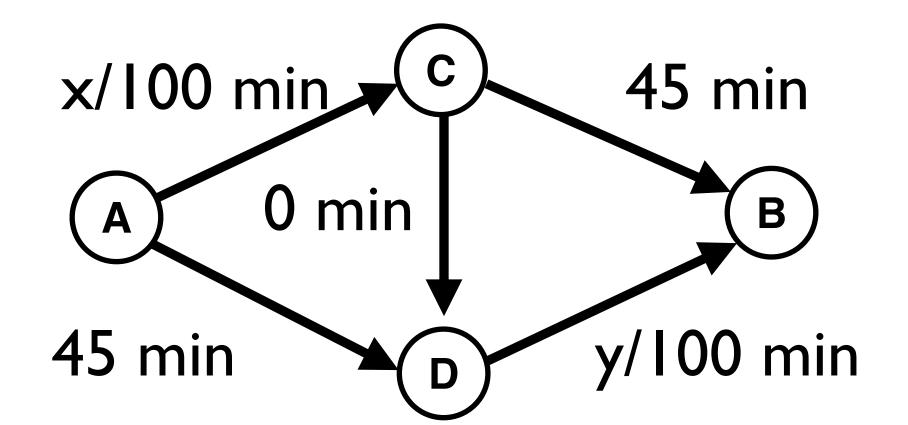
# Traffic modeled as a game

- Add a teleport!
- Players can take it if they want or not



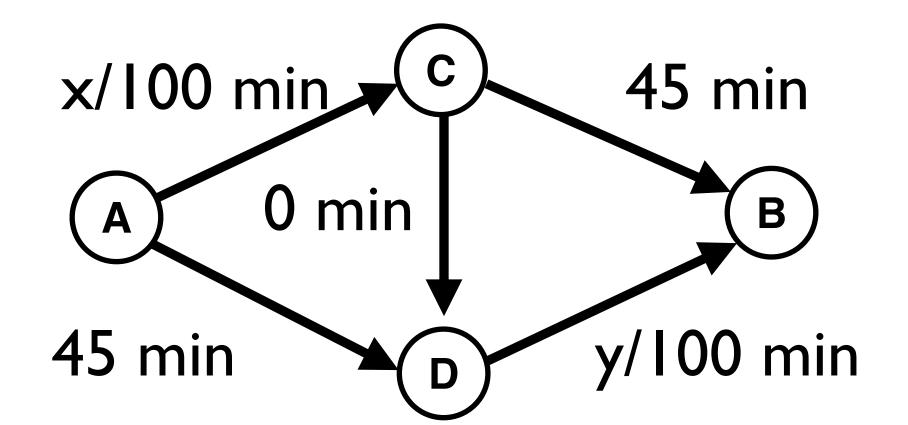
# Traffic modeled as a game

- **Players:** Drivers 1,2,3...,4000
- Strategies: A-C-B, A-D-B, A-C-D-B
- Payoffs: Negative drive time
  - A-C-B time: (x/100 + 45)
  - A-D-B time: (45 + y/100)
  - A-C-D-B time: (x/100 + y/100)



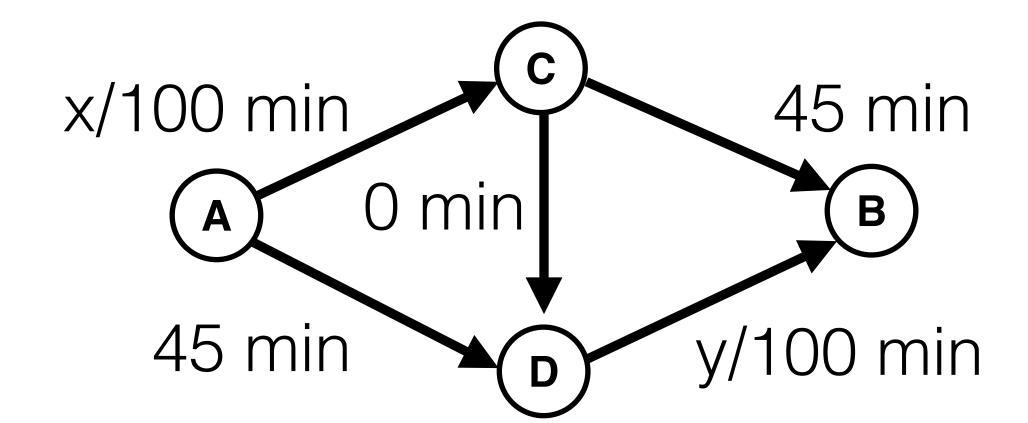
## Would you teleport?

- Payoffs when 2000 ACB, 2000 ADB, 0 ACDB
  - A-C-B time: (x/100 + 45)
    - -2000/100 + 45 = 65 minutes
  - A-D-B time: (45 + y/100)
    - -2000/100 + 45 = 65 minutes
  - A-C-D-B time: (x/100 + y/100)
    - $\bullet$  2000/100 + 2000/100 = 40 minutes



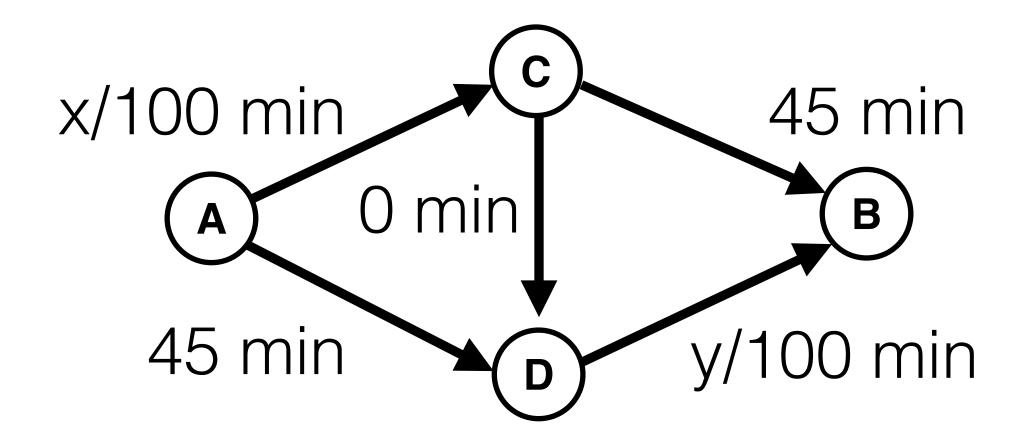
# New equilibrium?

- Payoffs when 0 ACB, 0 ADB, 4000 ACDB
  - A-C-B time: (x/100 + 45)
  - A-D-B time: (45 + y/100)
  - A-C-D-B time: (x/100 + y/100)



# New equilibrium?

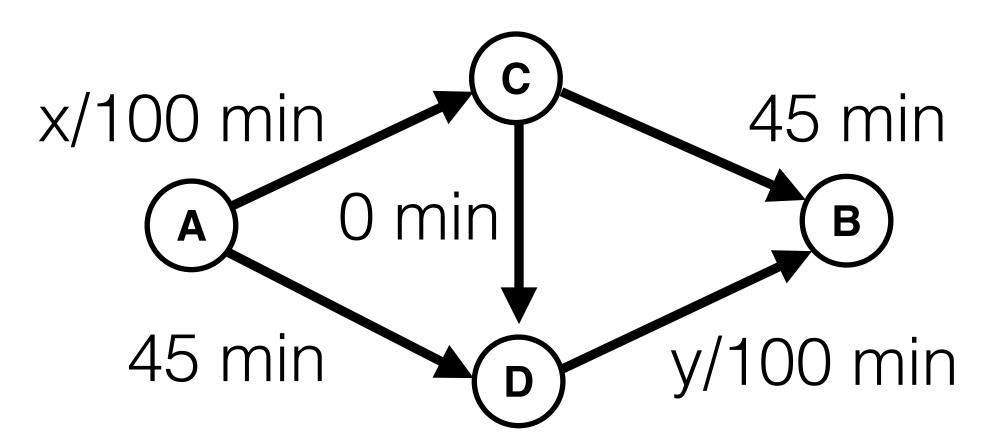
- Payoffs when 0 ACB, 0 ADB, 4000 ACDB
  - A-C-B time: (x/100 + 45)
    - -4000/100 + 45 = 85 minutes
  - A-D-B time: (45 + y/100)
    - -45 + 4000/100 = 85 minutes
  - A-C-D-B time: (x/100 + y/100)
    - -4000/100 + 4000/100 = 80 minutes



# New equilibrium?

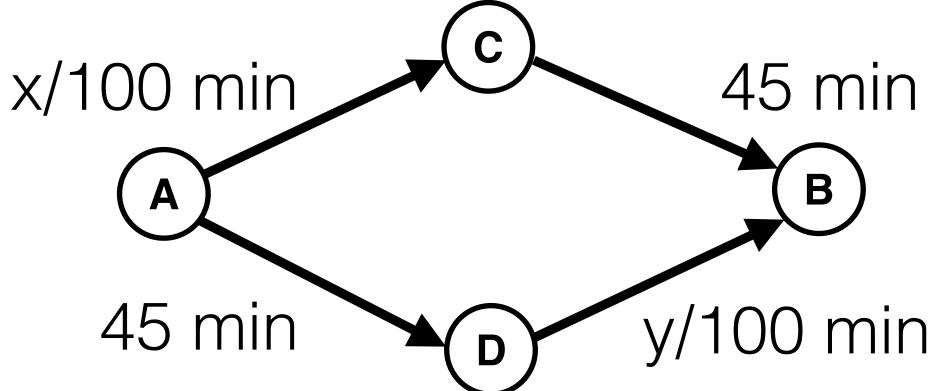
- · Payoffs when 0 ACB, 0 ADB, 4000 ACDB
  - A-C-B time: -(x/100 + 45)4000/100 + 45 = 85 minutes
  - A-D-B time: -(45 + y/100)45 + 4000/100 = 85 minutes
  - A-C-D-B time: -(x/100 + y/100)4000/100 + 4000/100 = 80 minutes

## **Dominant strategy = only equilibrium!**

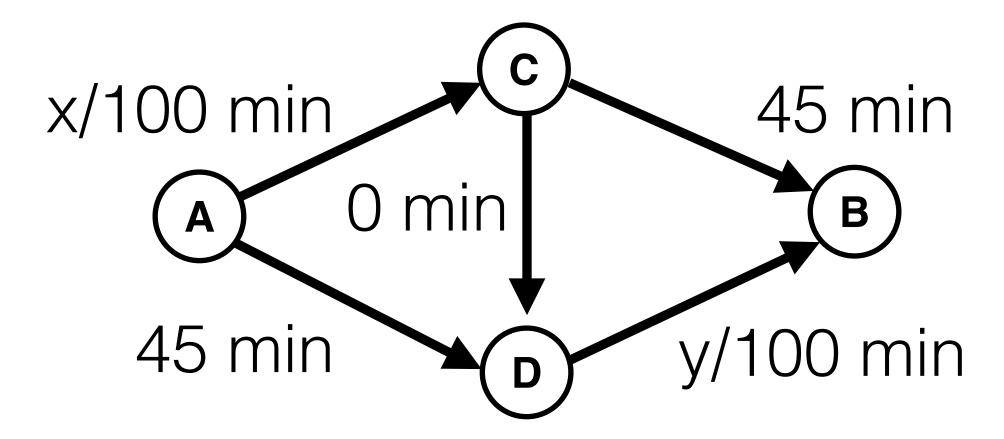


# What just happened?

• Equilibrium: 65 minutes

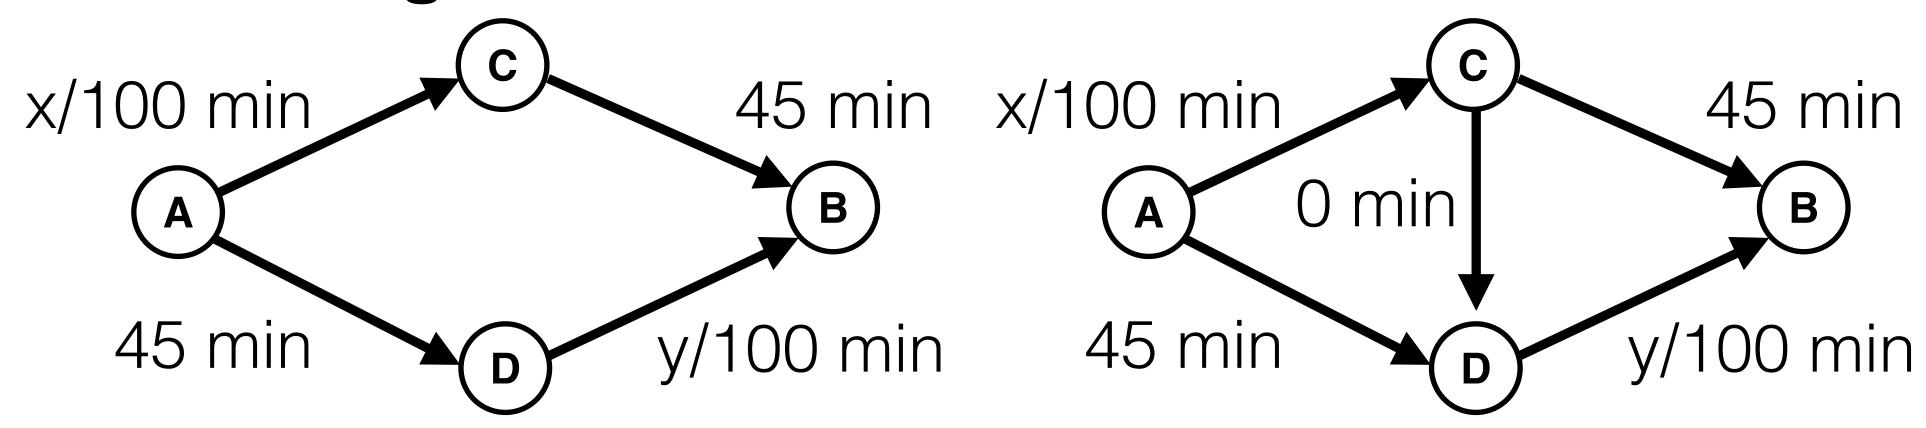


• Equilibrium: 80 minutes



# What does "paradox" mean?

## Routing:



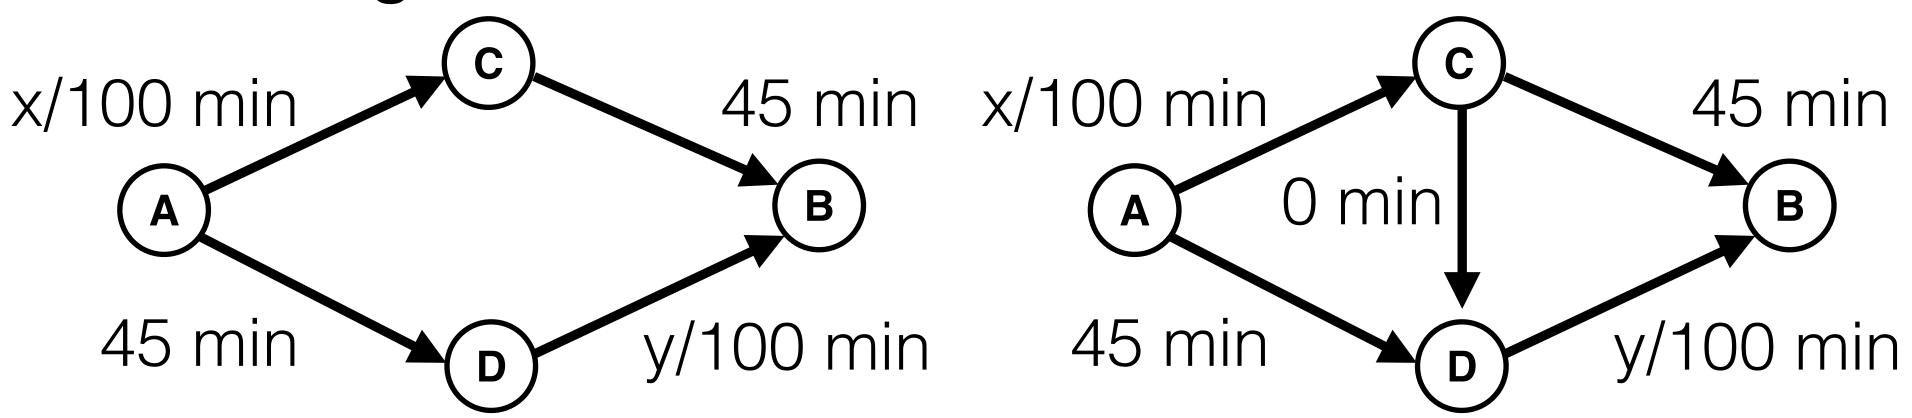
#### Prisoner's Dilemma:

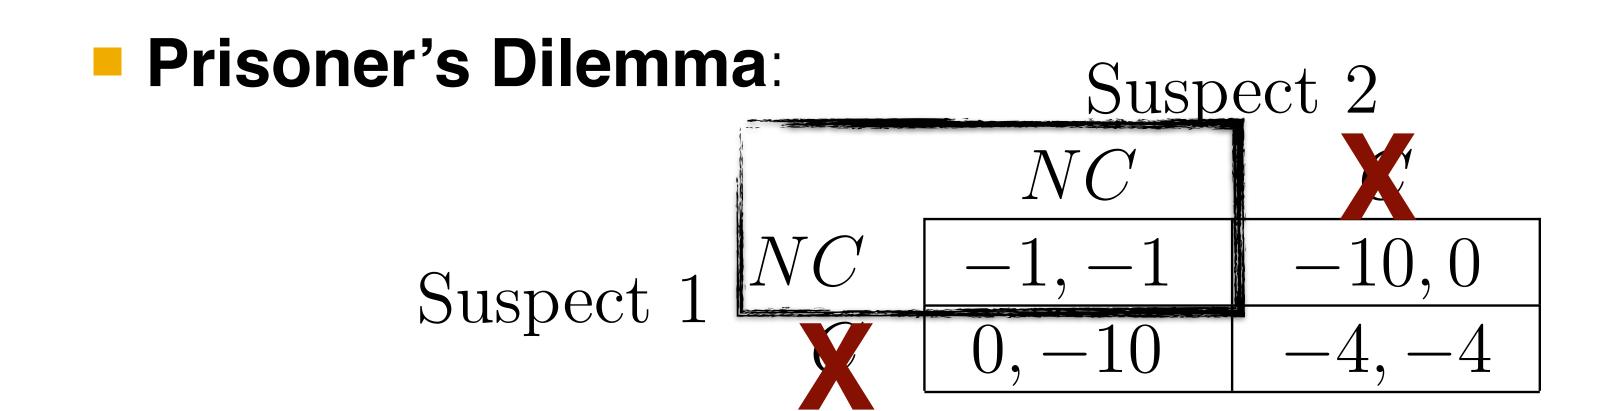
Suspect 2

Suspect 1 
$$NC = \begin{bmatrix} NC & C \\ -1, -1 & -10, 0 \\ C & 0, -10 & -4, -4 \end{bmatrix}$$

# What does "paradox" mean?

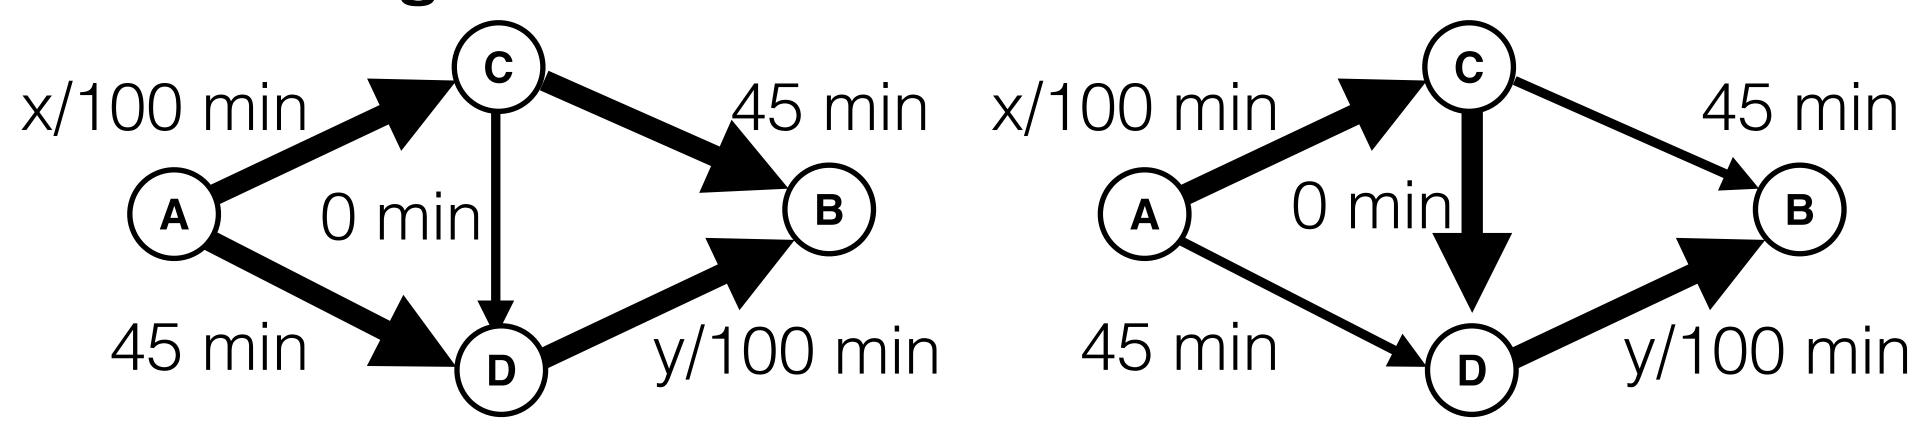
## Routing:





# How bad can it get?

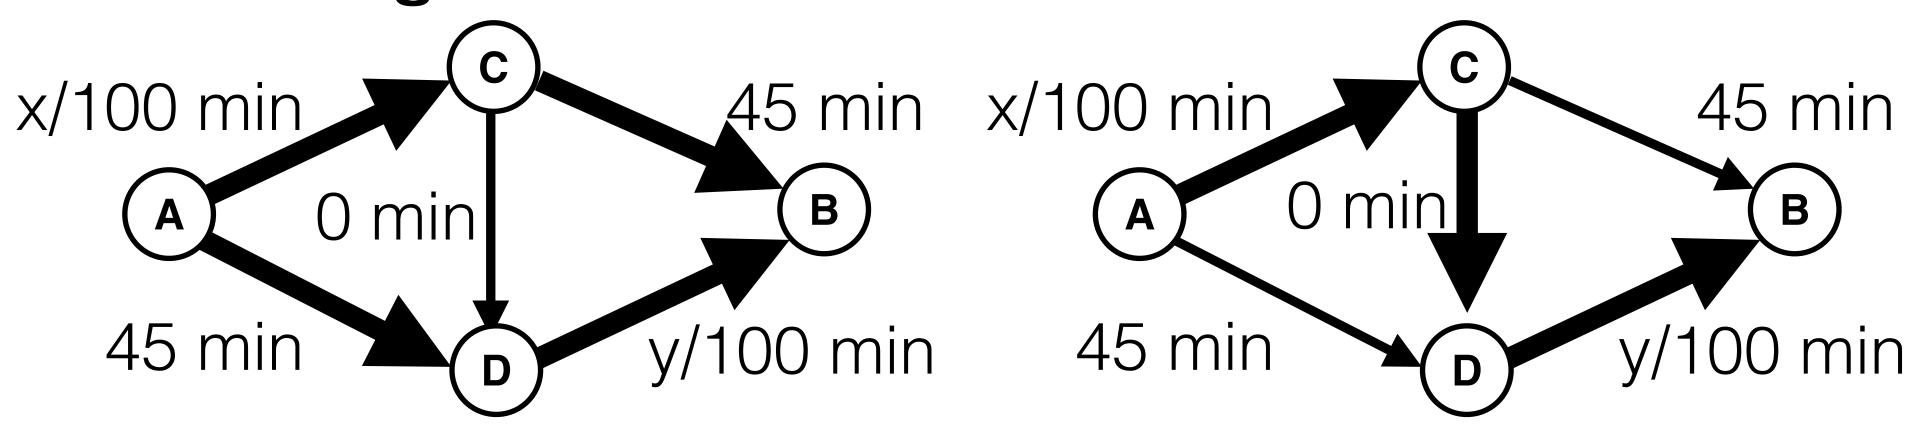
### Routing:



- Ratio between socially optimal and selfish routing?
  - This example: 85/65 = 1.23x worse
  - Worst case: 4/3 = 1.33 worse

# How bad can it get?

### Routing:



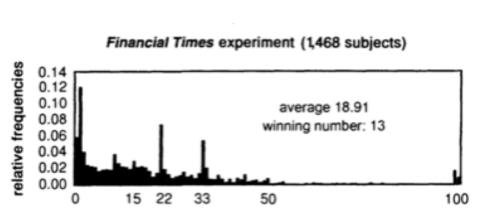
- Ratio between socially optimal and selfish routing?
  - This example: 85/65 = 1.23x worse
  - Worst case: 4/3 = 1.33 worse (!)

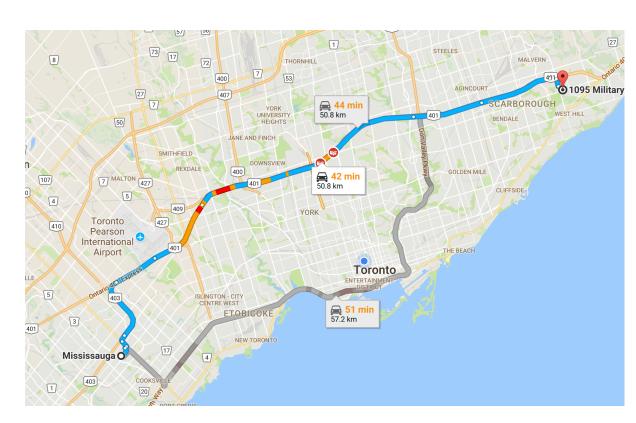
For selfish routing, "Price of Anarchy" = 4/3

# Today: game theory

## Mathematical framework to analyze strategic behaviour







- A game is characterized by players, strategies, and payoffs
- Captures a wide variety of strategic situations
- Best response, (strict) dominant strategies, mixed strategies, Nash equilibrium