CS 580

Algorithmic Game **Theory**

Instructor: Ruta Mehta

Game Theory

Multiple self-interested agents interacting in the same environment

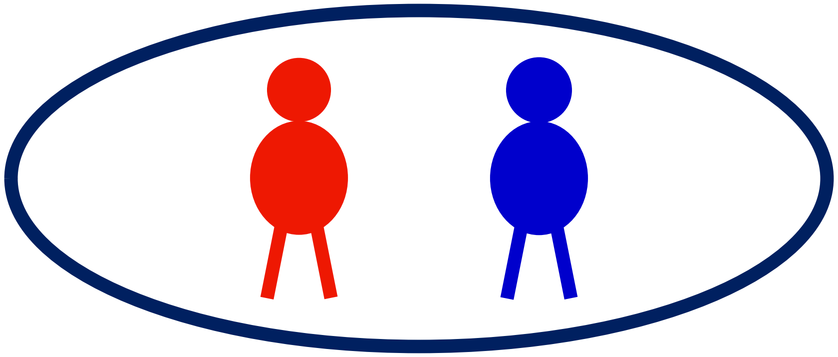

Deciding what to do.

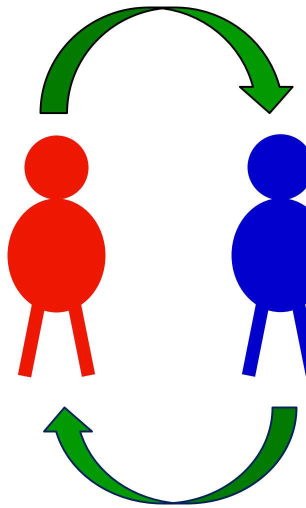

Q: What to expect? How good is it? Can it be controlled?

Game of Chicken (Traffic Light)

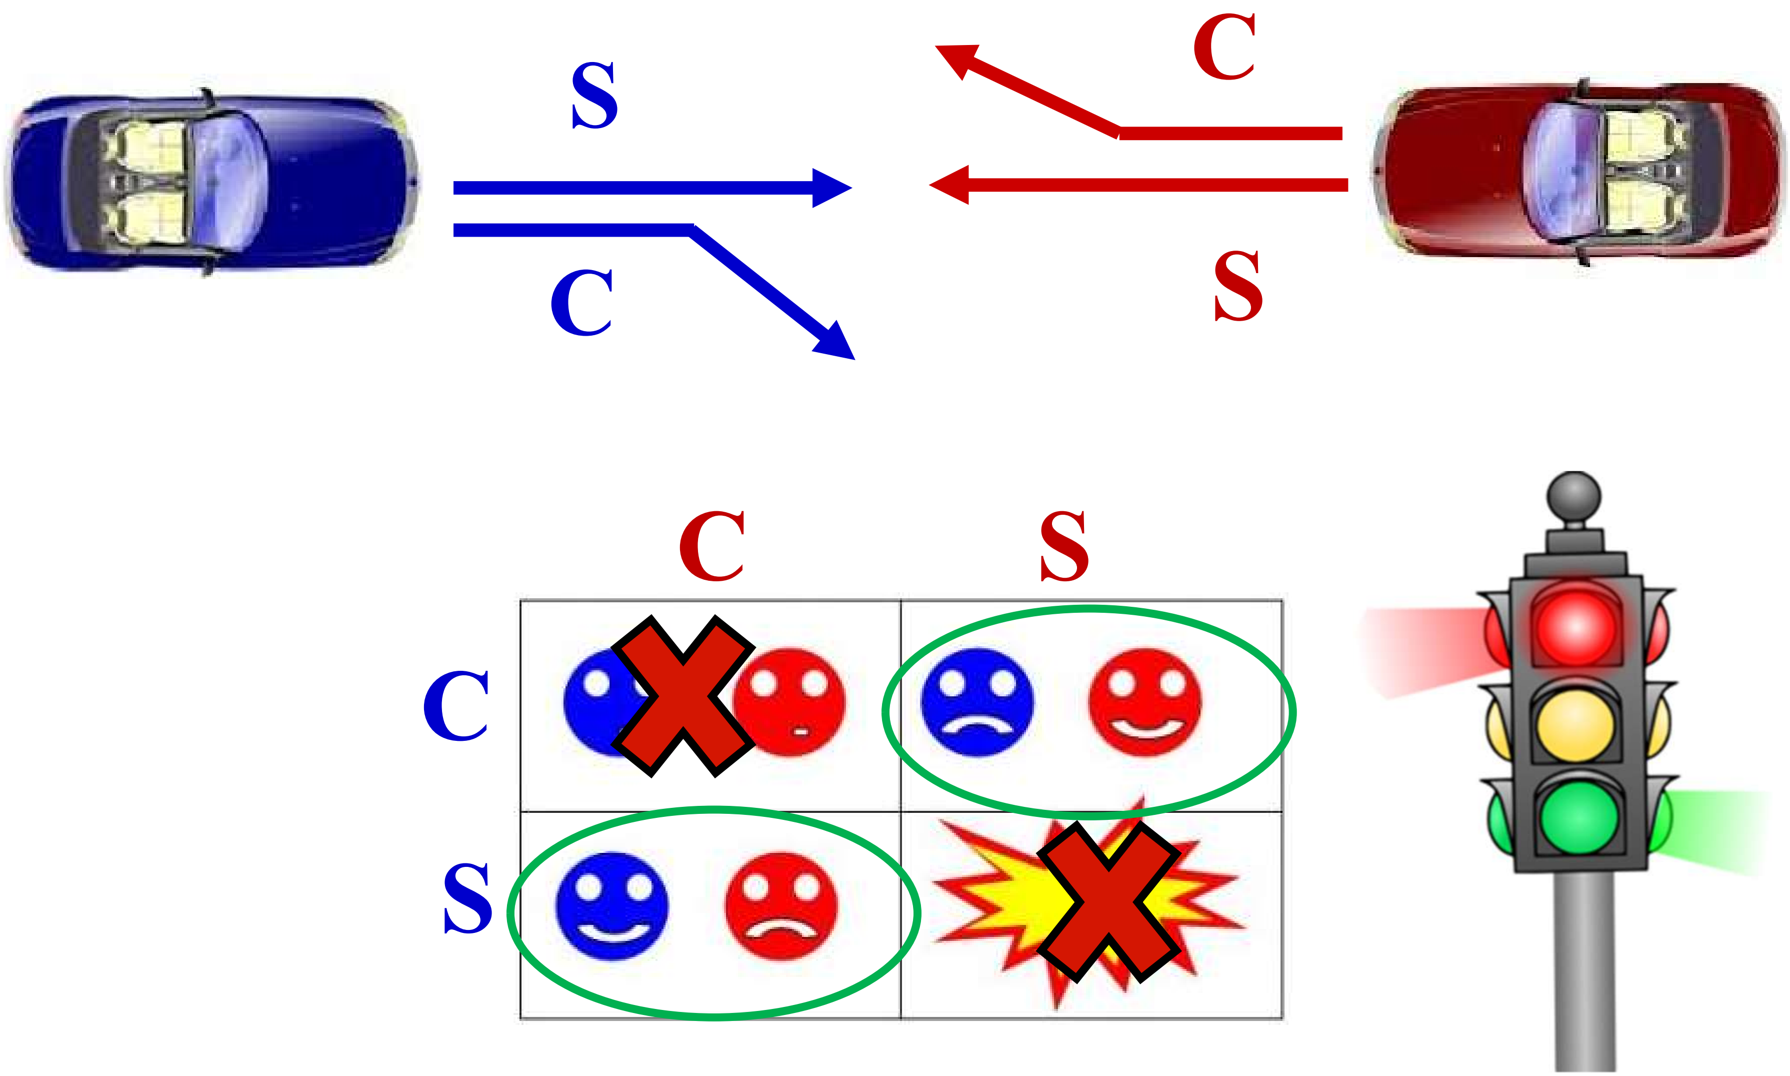

Algorithmic Game Theory

AGT, in addition, focuses on designing efficient algorithms to compute solutions that are crucial (e.g., to make accurate prediction).

- **What to expect**

  Research-oriented Course

  - Exposure to key concepts and proof techniques from AGT
  - Explore research problems and novel questions

- **What is expected from you**
  - Pre-req: Basic knowledge of linear-algebra, linear programming, probability, algorithms.
  - Energetic participation in class
  - Research/Survey Project (individually or in a group of two).

- Instructor: Ruta Mehta (Me)
- TA: Vasilis Livanos
- Office hours:
  - Ruta: Tue 2-3pm in Siebel 3218
  - Vasilis: Thu 10-11am in TBD

Useful links

- Webpage: https://courses.engr.illinois.edu/cs580/fa2022
- Piazza Page: piazza.com/illinois/fall2022/cs580
- Slack: FA22 - Algorithmic Game Theory CS 580
- Gradescope for grading

Check webpage/piazza at least twice a week for the updates.

HW0 will be posted today.

- Grading:
  - 3 homeworks – 30% (10,10,10)
  - Research/Survey Project – 45%
    - Work – 20%
    - Presentation – 12.5%
    - Report – 12.5%
  - Final Exam – 22%
  - Class participation – 3%

HW0 is for self-study (not to be submitted).

References

- T. Roughgarden, Twenty Lectures on Algorithmic Game Theory, 2016.
- N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani (editors), Algorithmic Game Theory, 2007. (Book available online for free.)
- R. Myerson, Game Theory: Analysis of conflict, 1991.

Recent papers, and other lecture notes that we will post on course website.

3 Broad Goals

Goal #1

Understand outcomes arising from
interaction of intelligent and self-interested agents.

Games and Equilibria

Prisoner's Dilemma

Two thieves caught for burglary.

Two options: {confess, not confess}

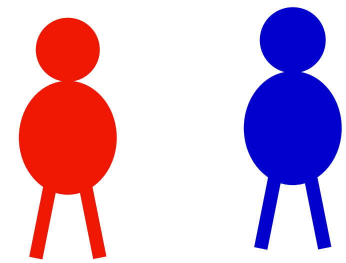

	N	C
N	-1 -1	-6 0
C	0 -6	-5 -5

Prisoner's Dilemma

Two thieves caught for burglary.

Two options: {confess, not confess}

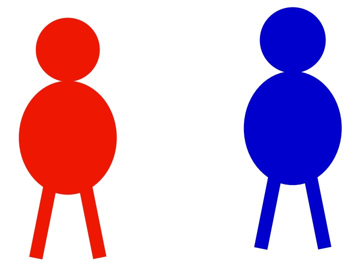

	N	C
N	-1 -1	-6 0
C	0 -6	-5 -5

Only stable state!

Rock-Paper-Scissors

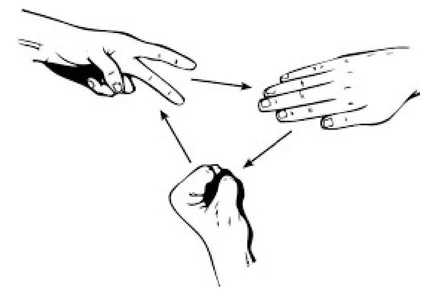

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

No pure stable state!

Both playing (1/3,1/3,1/3) is a NE.

Nash Eq.: No player gains by deviating individually

Why?

Rock-Paper-Scissors

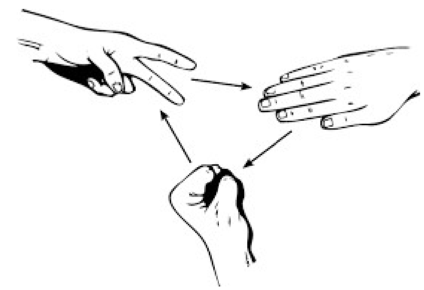

	R	P	S
R	0 0	-1 1	1 -1
P	1 -1	0 0	-1 1
S	-1 1	1 -1	0 0

No pure stable state!

Nash Eq.: No player gains by deviating individually

Both playing (1/3,1/3,1/3) is **the only** NE.

Why?

- Finite (normal form) games and Nash equilibrium existence
- Computation:
  - Zero-sum: minmax theorem,
  - General: (may be) Lemke-Howson algorithm
- Complexity: PPAD-complete
- Other equilibrium notions – correlated, markets, security games
- Incomplete information, Bayesian Nash
- Collusion, Core, Nash bargaining

Food for Thought

You and your friend choose a number …

Food for Thought

You and your friend choose a number …

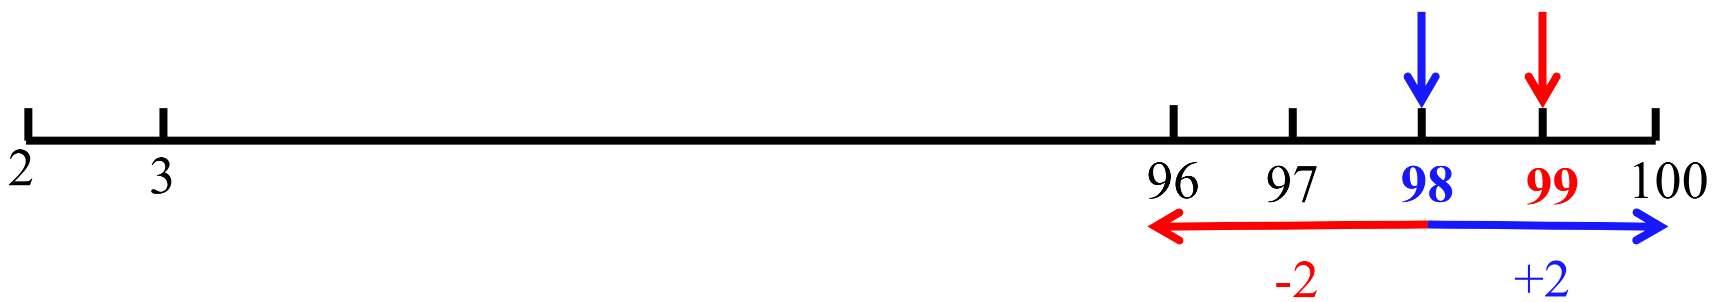

What will you choose? **What if +/- 50?**

What are Nash equilibria?

Tragedy of commons

Limited but open resource shared by many.

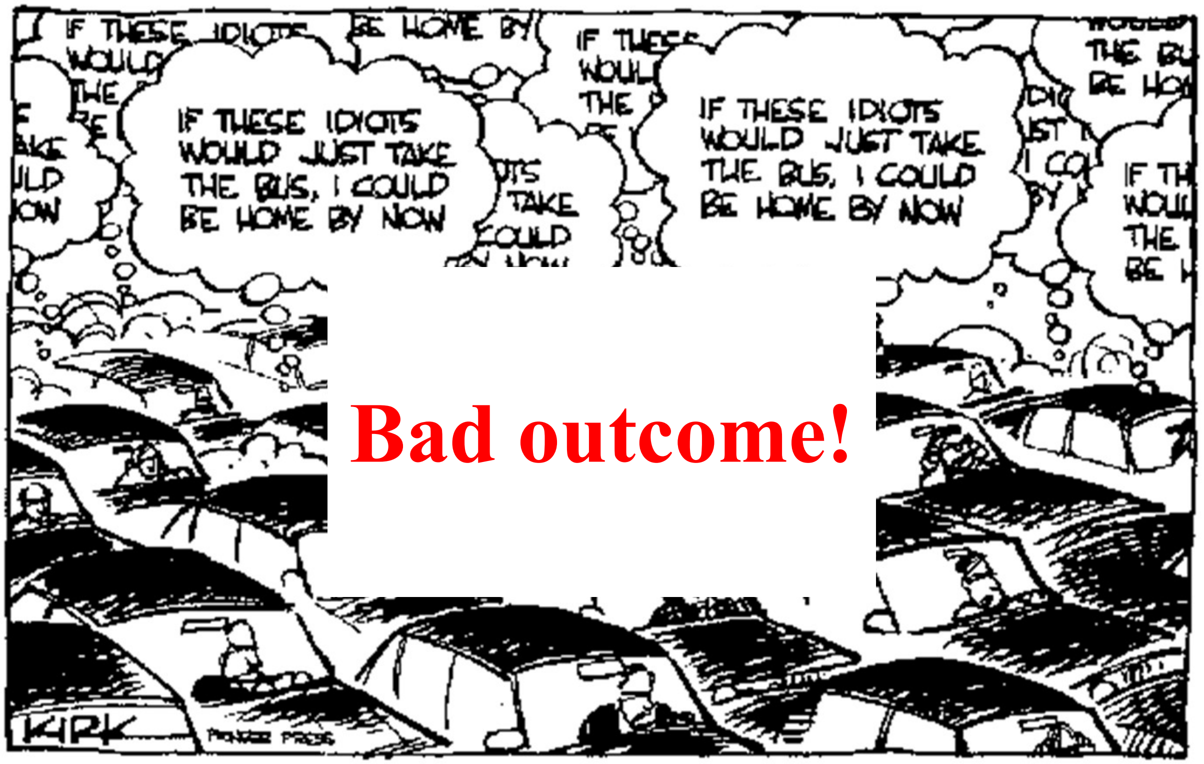

Bad outcome!

Stable: Over use => Disaster

Goal #2

Analyze quality of the outcome arising from strategic interaction, i.e. OPT vs NE.

Price of Anarchy

Braess' Paradox

60 commuters

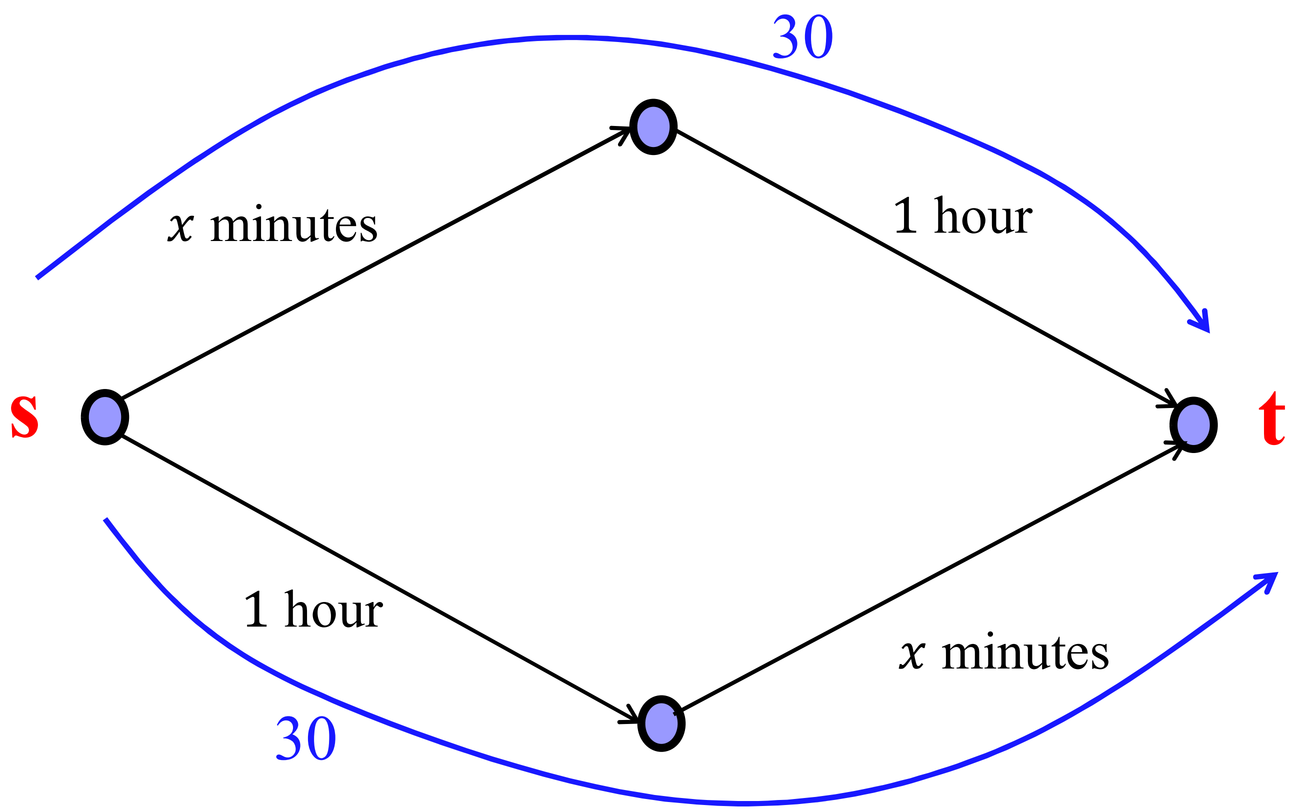

Commute time: 1.5 hours

Braess' Paradox

60 commuters

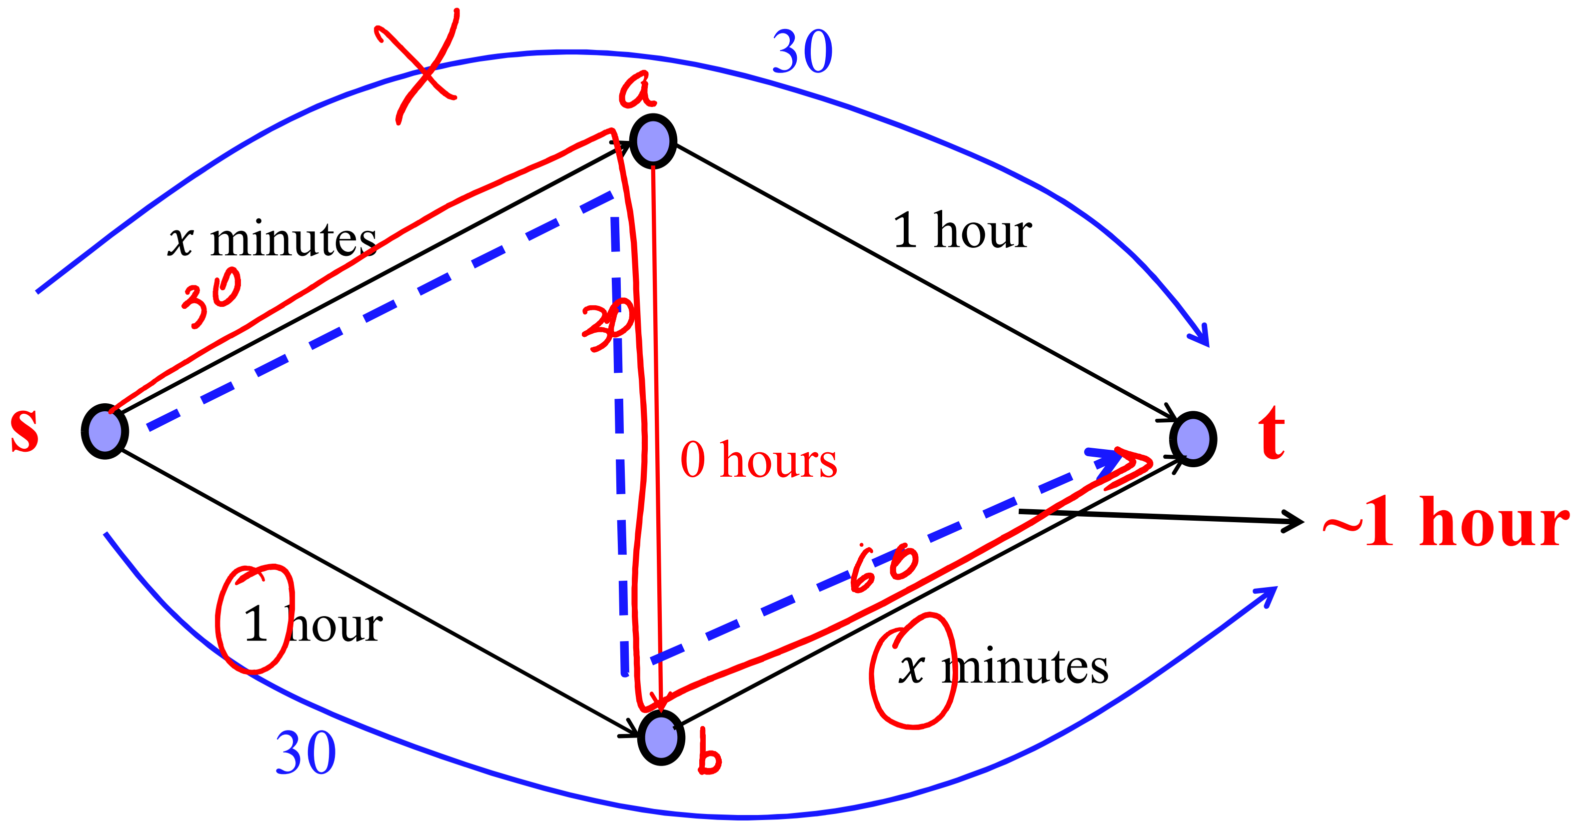

Commute time: 1.5 hours

Braess' Paradox

60 commuters

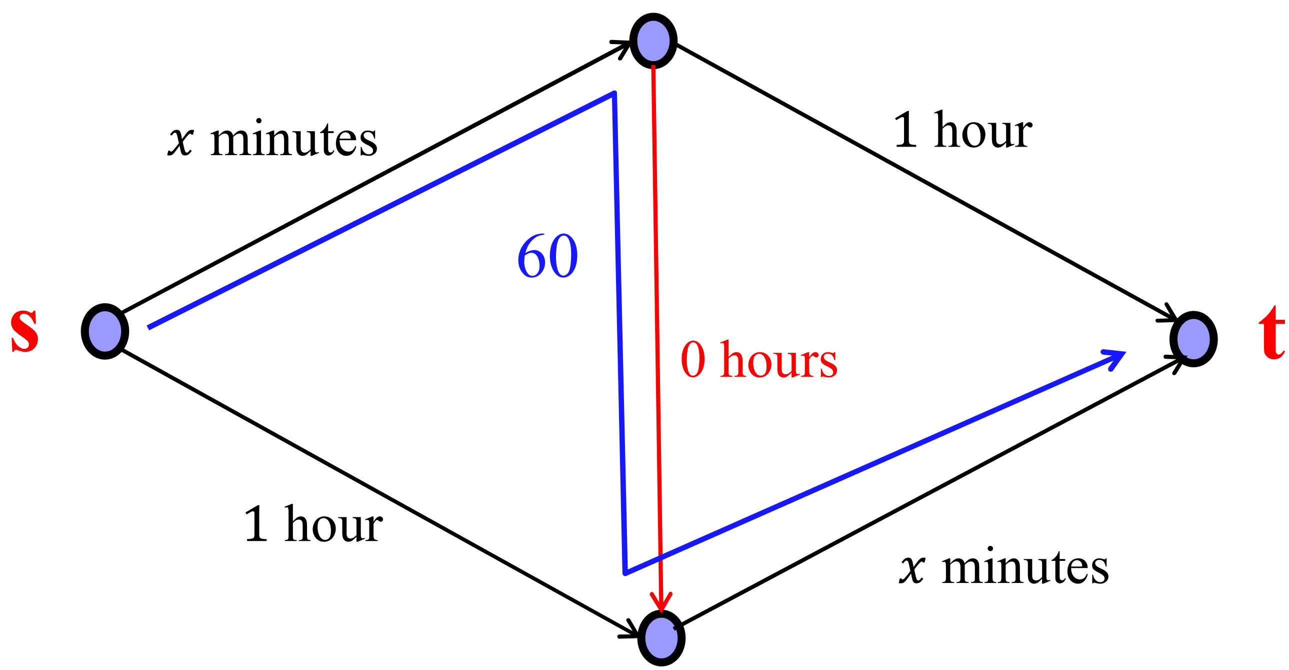

Commute time: 2 hours!

[Braess' Paradox in real life]

Braess' Paradox

60 commuters

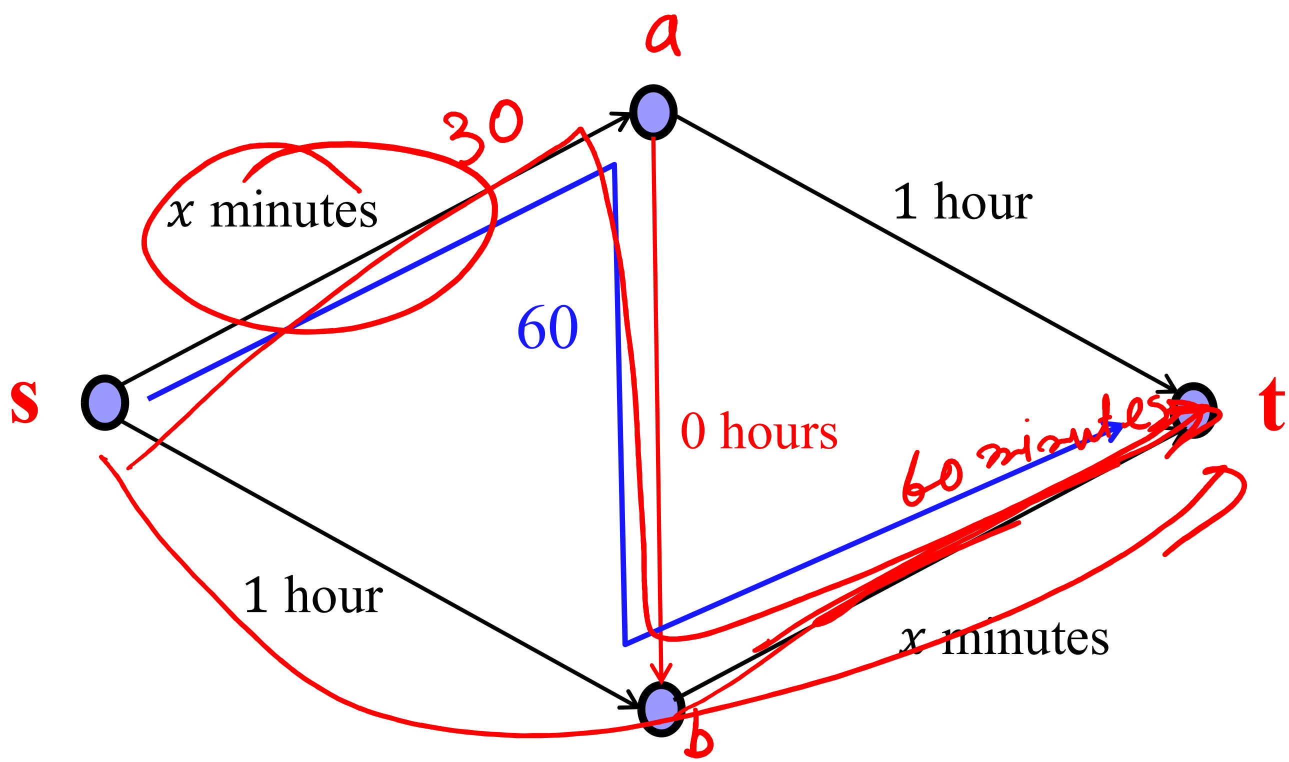

Price of Anarchy (PoA): $\frac{\textit{worst NE}}{\textit{OPT}} = \frac{2}{1.5} = \frac{4}{3}$

Can not be worse!

- Network routing games
- Congestion (potential) games
- PoA in linear congestion games
  - Smoothness framework
- Iterative play (dynamics) and convergence

Goal #3

Designing rules to ensure "good" outcome under strategic interaction among selfish agents.

Mechanism Design

At the core of large industries

Online markets – eBay, Uber/Lyft, TaskRabbit, cloud markets

Spectrum auction – distribution of public good. enables variety of mobile/cable services.

Search auction – primary revenue for google!

Tons of important applications

Fair Division – school/course seats assignment, kidney exchange, air traffic flow management, …

Matching residents to hospitals,
Voting, review, coupon systems.

So on …

- MD without money
  - Fair division
    - Divisible items: Competitive equilibrium
    - Indivisible items: EF1, EFX, MMS, Max. Nash Welfare, …
  - Stable matching, Arrow's theorem (voting)

- MD with money
  - First price auction, second price auction, VCG
  - Generalized second price auction for search (Google)
  - Optimal auctions: Myerson auction and extensions
  - Prophet inequalities and simple auctions
  - Fair MD (may be)

Fun Fact!

Olympics 2012 Scandal

Check out Women's doubles badminton tournament

Video of the fist controversial match

Example: How to divide fairly?

How to divide among the two so that both are happy with their share, and the division seems "fair" to both?

Sol'n: I-**Cut-You-Choose**

PS: Finds mention in the Bible, in the Book of Genesis (chapter 13).

Example: How to divide fairly?

Sol'n:
I-Cut-you-Choose

Envyfree: No one envies other's share

Proportional: Each gets at least half the value
(assuming $v(A \cup B) \leq v(A) + v(B)$, for $A, B \subseteq Cake$)

PS: Finds mention in the Bible, in the Book of Genesis (chapter 13).

Divisible goods

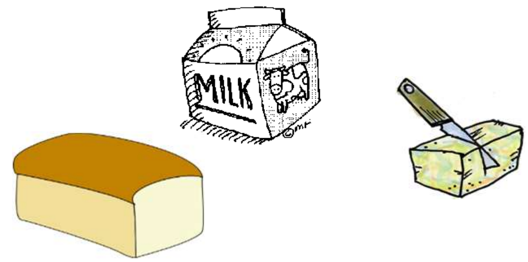

Goal: Find *fair* and *efficient* allocation


ASSETS
LIABILITIES
U.S.ARMY
U.S.AIR FORCE
UNITED STATES COAST GUARD 1790
MARINE CORPS
NAVY

Model

- A: set of n agents 
- M: set of m divisible goods (manna)

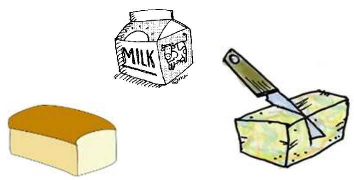

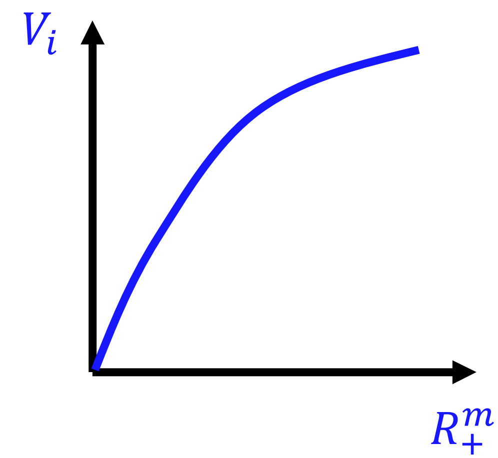

- Each agent i has
  - Concave valuation function $V_i: R_+^m \to R_+$ over bundles of items
  - Captures ***decreasing marginal returns***.

Goal: Find *fair* and *efficient* allocation

Agreeable (Fair)

Non-wasteful (Efficient)

Allocation: Bundle $X_i \in R_+^m$ to agent i

Envy-free: **No agent *envies* other's allocation over her own.**

For each agent i,
$V_i(X_i) \geq V_i(X_j), \forall j \in [n]$

Proportional: **Each agent i gets value at least $\frac{v_i(M)}{n}$**

For each agent $i, V_i(X_i) \geq \frac{v_i(M)}{n}$

Pareto-optimal: **No other allocation is better for all.**

Welfare Maximizing
$(max: \sum_i V_i)$

Example: Half moon cookie

(i)

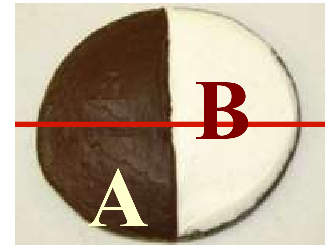

(ii)

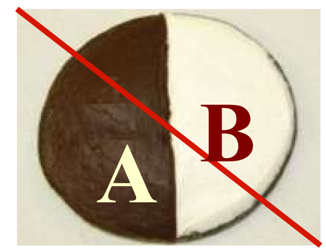

(iii)

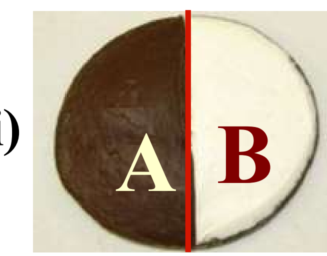

Agreeable (Fair)

Envy-free: No agent *envies* other's allocation over her own.

Proportional: Each agent i gets value at least $\frac{v_i(M)}{n}$

Non-wasteful (Efficient)

Allocation in red

[3, 2, 2]
[0, 0, 0]

[20, 20, 30]
[0, 0, 0]

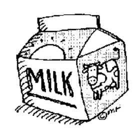

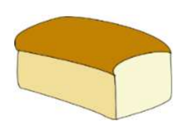

Agreeable (Fair)

Non-wasteful (Efficient)

Envy-free: No agent *envies* other's allocation over her own.

Proportional: Each agent i gets value at least $\frac{v_i(M)}{n}$

Pareto-optimal: No other allocation is better for all.

[3, 2, 2]
[1/2, 1/2, 1/2]

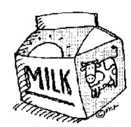

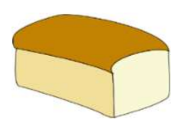

Allocation in red

[20, 20, 30]
[1/2, 1/2, 1/2]

Agreeable (Fair)

Envy-free: No agent *envies* other's allocation over her own.

Proportional: Each agent i gets value at least $\frac{v_i(M)}{n}$

Non-wasteful (Efficient)

Pareto-optimal: No other allocation is better for all.

Welfare Maximizing

$(max: \sum_i V_i)$

Allocation in red

[3, 2, 2]
[1, 1/2, 0]

[20, 20, 30]
[0, 1/2, 1]

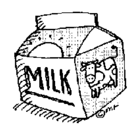

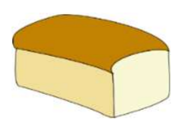

Agreeable (Fair)

Envy-free: No agent *envies* other's allocation over her own.

Proportional: Each agent i gets value at least $\frac{v_i(M)}{n}$

Non-wasteful (Efficient)

Pareto-optimal: No other allocation is better for all.

Welfare Maximizing $(max: \sum_i V_i)$

Allocation in red

[3, 2, 2]
[0, 0, 0]

[20, 20, 30]
[1, 1, 1]

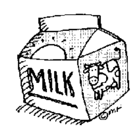

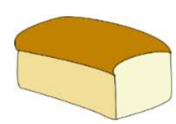

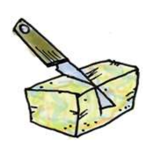

Agreeable (Fair)

Envy-free: No agent *envies* other's allocation over her own.

Proportional: Eac agent i gets value at least $\frac{v_i(M)}{n}$

Non-wasteful (Efficient)

Pareto-optimal: No other allocation is better for all.

(Nash) Welfare Maximizing ($\Pi_i V_i$)

Allocation in red

[3, 2, 2]
[1, 1/2, 0]

[20, 20, 30]
[0, 1/2, 1]

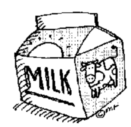

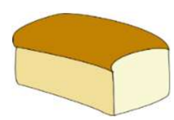

Agreeable (Fair)	Non-wasteful (Efficient)
Envy-free	Pareto-optimal
Proportional	(Nash) Welfare Maximizing

Competitive Equilibrium
(with equal income)

Competitive (market) Equilibrium (CE)

traditional setting…

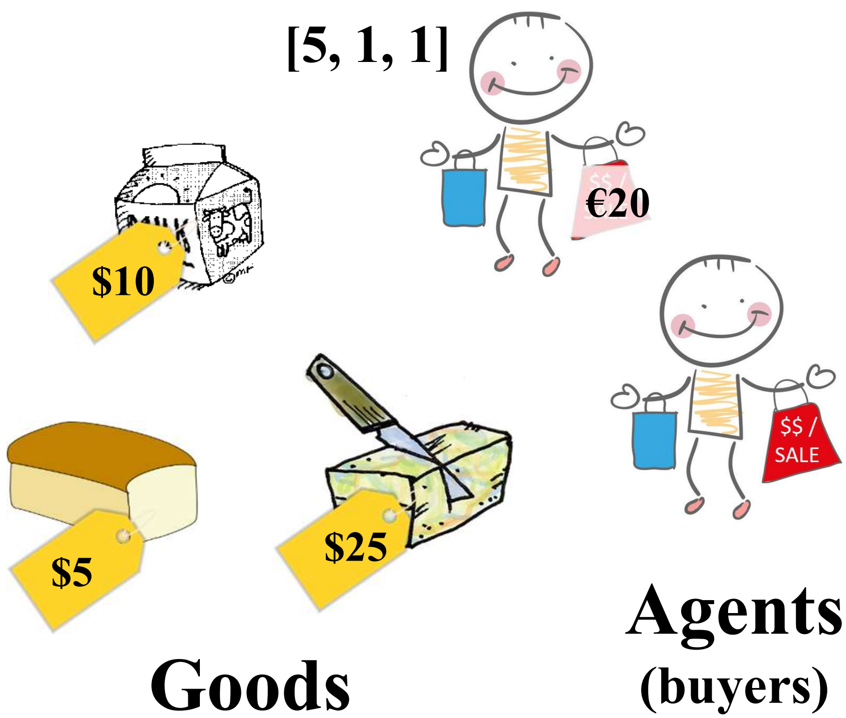

Demand optimal bundle

Competitive (market) Equilibrium (CE)

traditional setting…

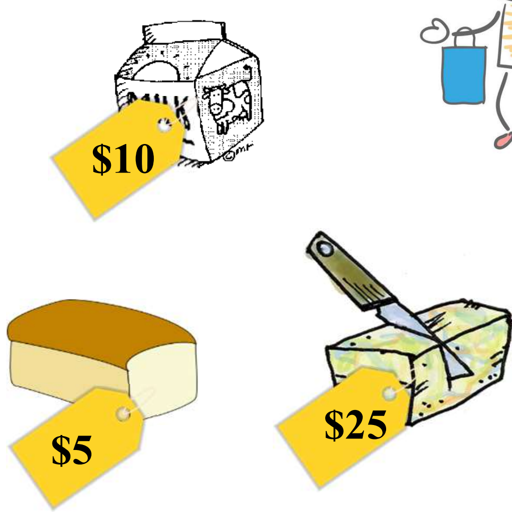

Goods

Agents
(buyers)

Buy optimal bundle → Demand

Competitive Equilibrium:
Demand = Supply

CE Example

traditional setting…

[2, 0]

[5, 1]

Demand ≠ Supply!

[1, 4]

[0, 1]

CE Example

traditional setting…

[1, 0]

[5, 1]

Demand = Supply
CE

[1, 4]

[0, 1]

w/ equal income (CEEI):
Agents have the same amount of money

CEEI: Properties

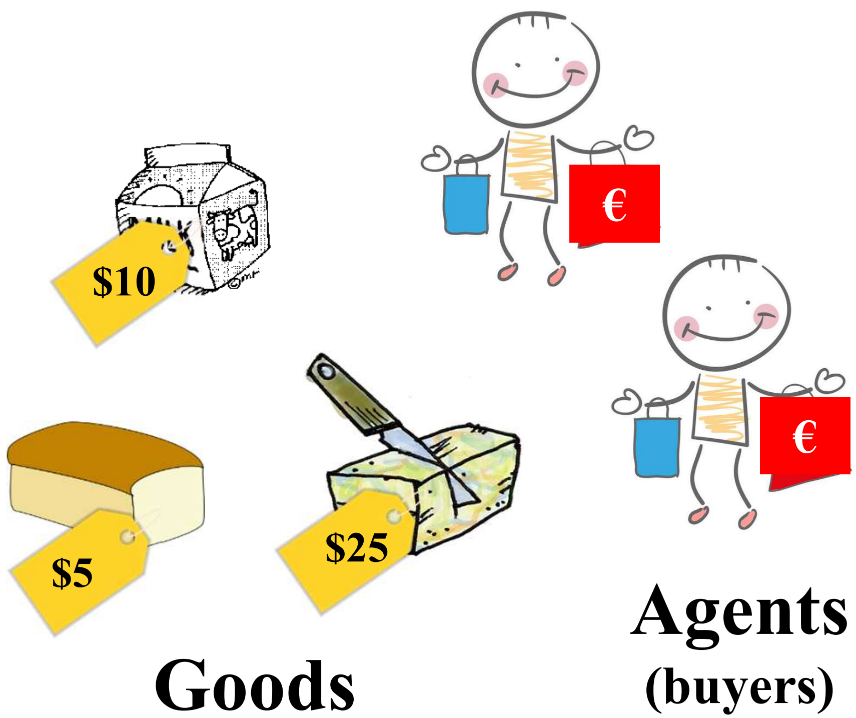

An agent can afford anyone's bundle, but demands hers
⇒ Envy-free

Envy-free + everything allocated
⇒ Proportional

1^{st} welfare theorem
⇒ Pareto-optimal

Demand optimal bundle

Competitive Equilibrium:
Demand = Supply

CE History

Adam Smith
(1776)

Leon Walras
(1880s)

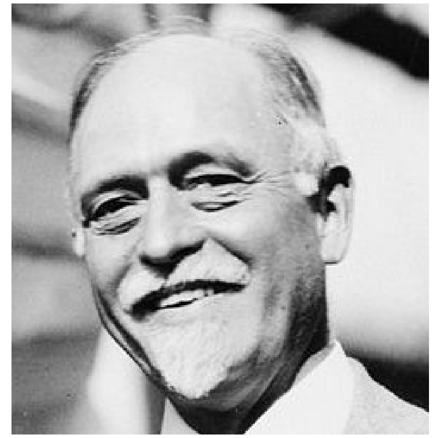

Irving Fisher (1891)

Arrow-Debreu (1954)

(Nobel prize)

(Existence of CE in the exchange model w/ firms)

...

Computation of CE (w/ goods)

Algorithms

- Convex programming formulations
  - Eisenberg-Gale (1959): CEEI w/ 1-homogeneous valuations
  - Shmyrev (2009), DGV (2013), CDGJMVY (2017) …
- (Strongly) Poly-time algorithms (linear valuations)
  - DPSV (2002), Orlin (2010), DM (2015), GV (2019) …
- Simplex-like algorithms: Eaves (1976), GM.SV (2011), GM.V (2014), …

Complexity

- PPAD: Papadimitrou'92, CDDT'09, VY'11, CPY'17, Rubinstein'18, …
- FIXP: EY'09, GM.VY'17, F-RHHH'21 …

Learning: RZ'12, BDM.UV'14, …, FPR'22, …

Matching/mechanisms: BLNPL'14, …, KKT'15, …, FGL'16, …, AJT'17, …, BGH'19, BNT-C'19, …

***Alaei, Bei, Branzei, Chen, Cole, Daskalakis, Deng, Devanur, Duan, Dai, Etessami, Feldman, Fiat, Filos-Ratsikas, Garg, Gkatzelis, Hansen, Hogh, Hollender, Jain, Jalaly, Hoefer, Kleinberg, Lucier, Mai, Mehlhorn, Mehta, Mansour, Morgenstern, Nisan, Paes, Lee, Leme, Papadimitriou, Paparas, Parkes, Roth, Saberi, Sohoni, Talgam-Cohen, Tardos, Vazirani, Vegh, Yazdanbod, Yannakakis, Zhang,… … …**