15-382 Collective Intelligence – S18

Lecture 26: Game Theory 1

Instructor:
Gianni A. Di Caro

Ice-cream Wars

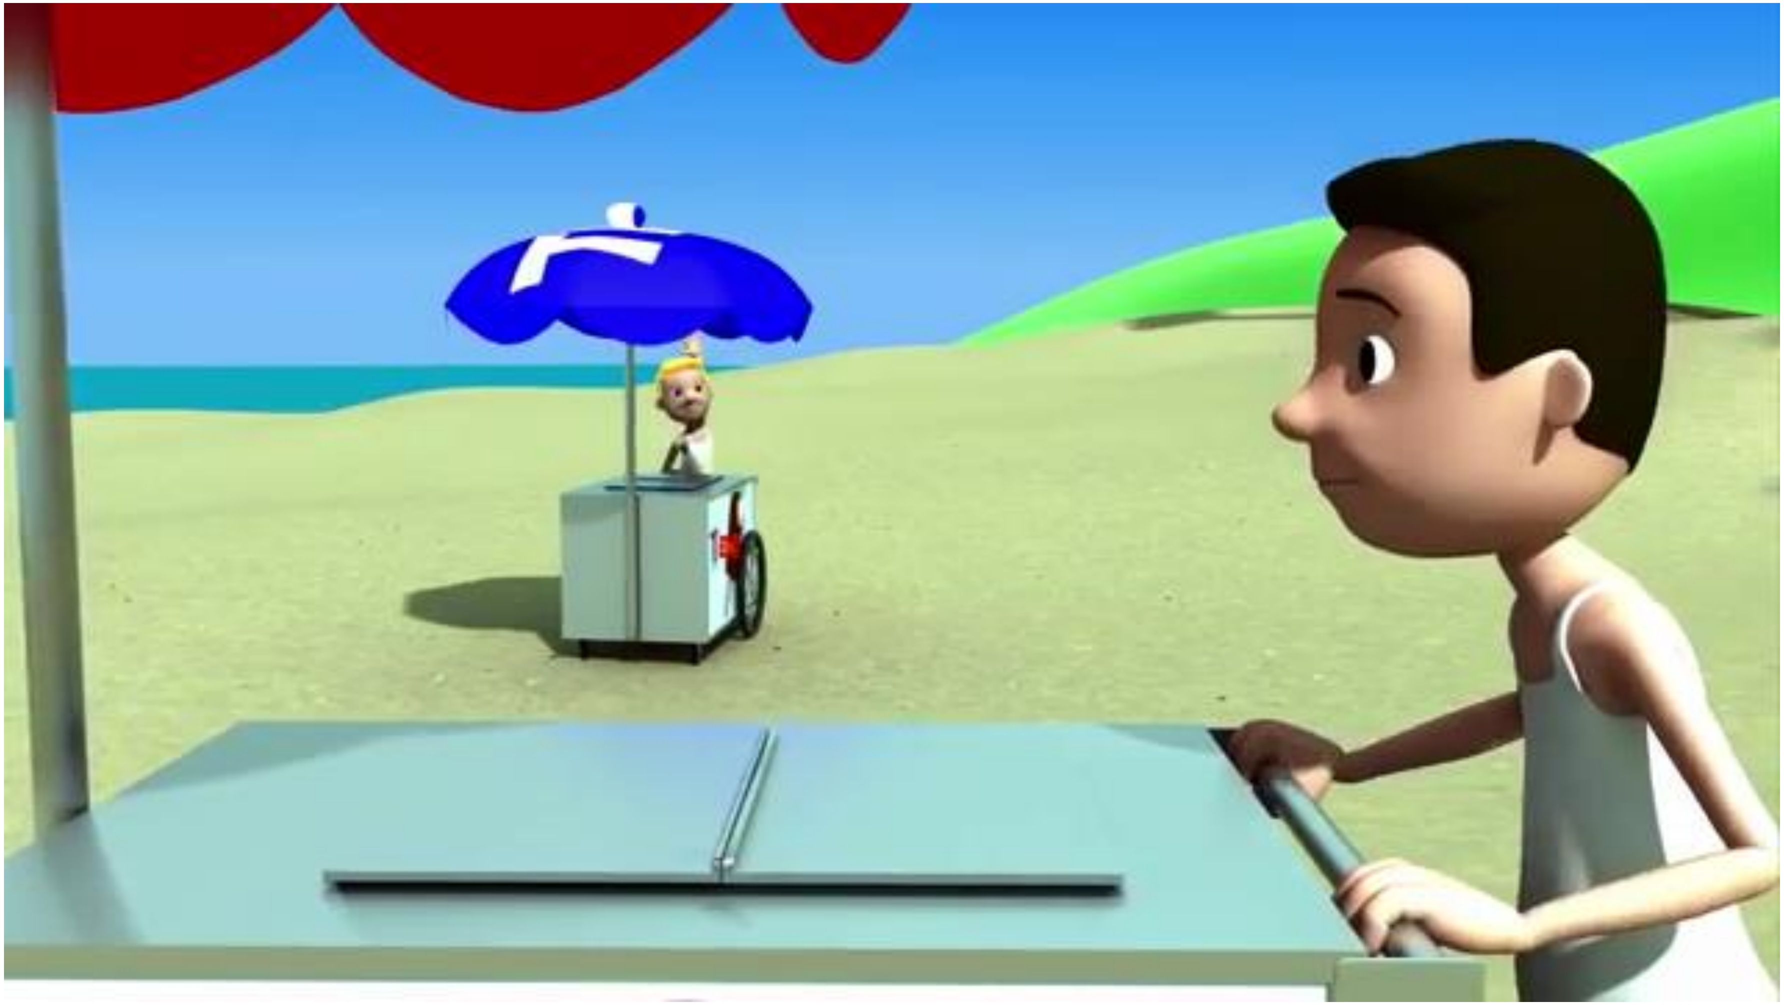

http://youtu.be/jILgxeNBK_8

Game Theory

- Game theory is the formal study of conflict and cooperation in (rational) multi-agent systems
- Decision-making where several players must make choices that potentially affect the interests of other players: **the effect of the actions of several agents are interdependent** (and agents are aware of it)
- Example: Auctioning!

Psychology:
Theory of social situations

Elements of a Game

- The players: how many players are there? Does nature/chance play a role? Players are assumed to be <u>rational</u>

- A complete description of what the players can do: the set of all possible actions.

Elements of a Game

- A description of the payoff / *consequences* for each player for every possible combination of actions chosen by all players playing the game.

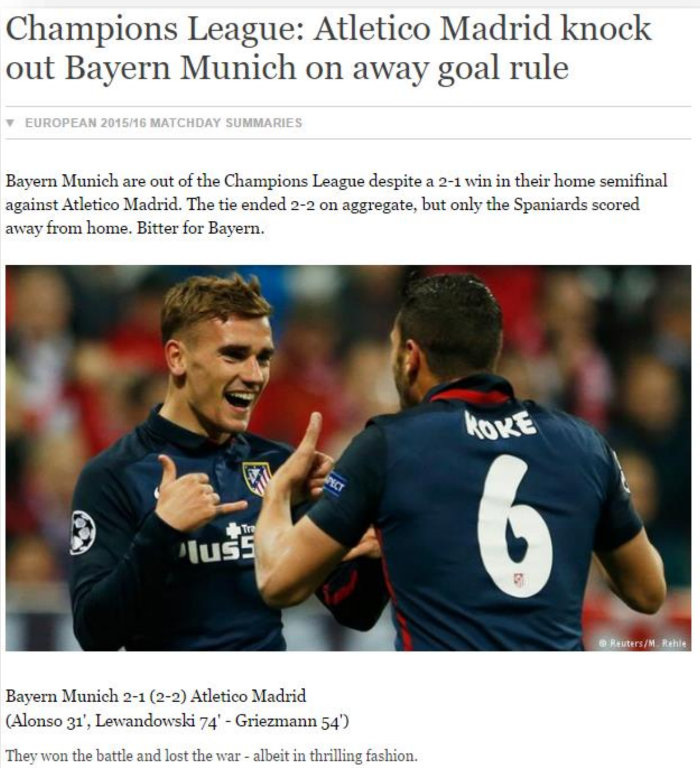
Champions League: Atletico Madrid knock out Bayern Munich on away goal rule

▼ EUROPEAN 2015/16 MATCHDAY SUMMARIES

Bayern Munich are out of the Champions League despite a 2-1 win in their home semifinal against Atletico Madrid. The tie ended 2-2 on aggregate, but only the Spaniards scored away from home. Bitter for Bayern.

© Reuters/M. Rehle

Bayern Munich 2-1 (2-2) Atletico Madrid
(Alonso 31', Lewandowski 74' - Griezmann 54')

They won the battle and lost the war - albeit in thrilling fashion.

- A description of all players' preferences over payoffs

Utility function for each player

Agent vs. Mechanism Design

- **Agent strategy design:** Game theory can be used to compute the expected utility for each decision, and use this to determine the best strategy (and its expected return) against a rational player

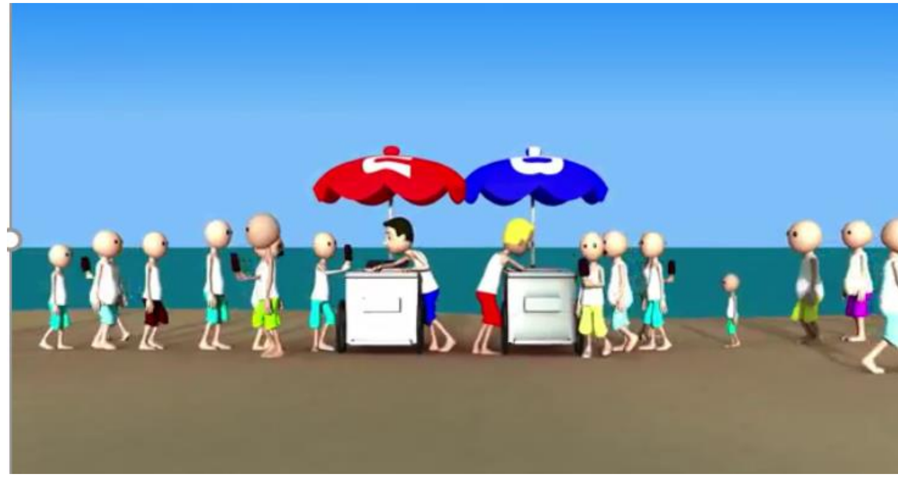

Strategy ≡ Policy

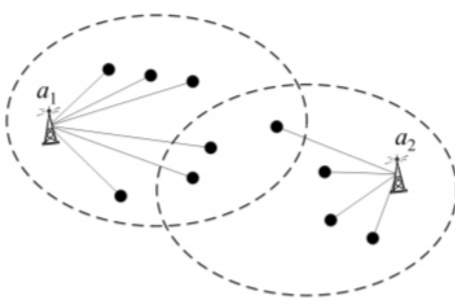

- **System-level mechanism design:** Define the rules of the game, such that the *collective utility* of the agents is maximized when each agent strategy is designed to maximize its *own utility* according to ASD

MAKING DECISIONS: BASIC DEFINITIONS

- Decision-making can involve choosing:
  - one *single action* or
  - a *sequence of actions*
- *Action outcomes can be certain or subject to uncertainty*
- A set A of alternative actions to choose from is given, it can be either discrete or continuous
- **Payoff (for a single agent):** function $\pi: A \rightarrow \mathbb{R}$ that associates a numerical values with every action in A
- **Optimal action** a^* (for a single agent scenario): $\pi(a^*) \geq \pi(a) \;\; \forall a \in A$
- **Payoff (for a multi-agent scenario):** The payoff of the action a for agent i depends on the actions of the other players! $\pi: A^n \rightarrow \mathbb{R}$
- **Strategy:** rule for choosing an action at every point a decision might have to be made (depending or not on the other agents)
- The strategy defines the ***behavior*** of an agent
- The observed behavior of an agent following a given strategy is the **outcome of the strategy**

Pure vs. Randomized Strategies

- **Pure strategy:** a strategy in which there is *no randomization*, one specific action is selected with certainty at each decision node
- All possible pure strategies define the <u>pure strategy set</u> S
- A ***decision tree*** can be used to represent a *sequence of decisions*

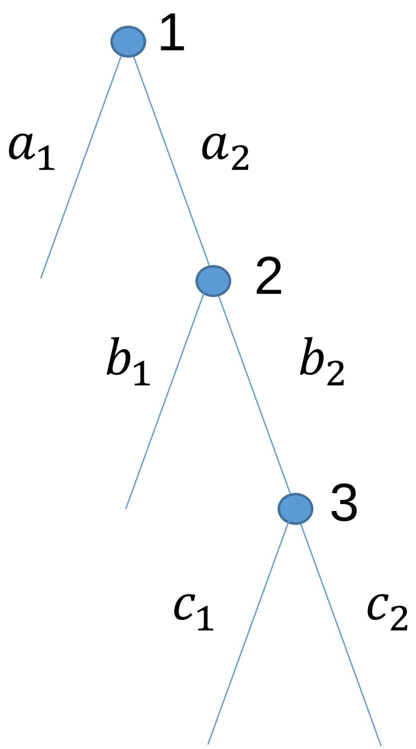

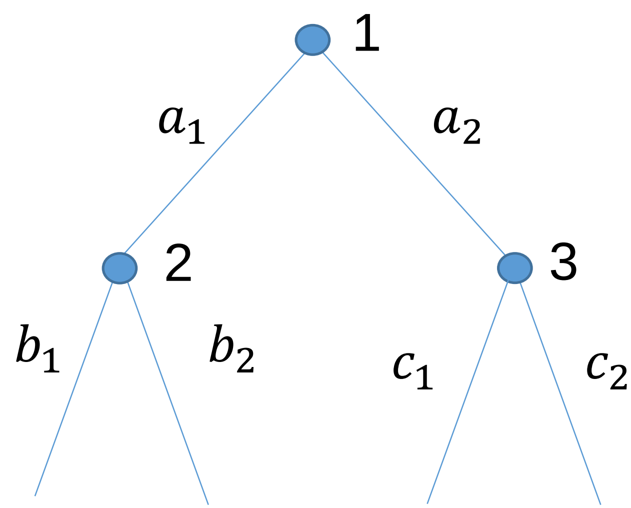

$$A_1 = \{a_1, a_2\}, \quad A_2 = \{b_1, b_2\}, \quad A_3 = \{c_1, c_2\}$$

- Three action sets (actions may the be same), that result in the pure strategy set: $S = \{a_1b_1c_1, a_1b_1c_2, a_1b_2c_1, a_1b_2c_2, a_2b_1c_1, a_2b_1c_2, a_2b_2c_1, a_2b_2c_2\}$

PURE VS. RANDOMIZED STRATEGIES

- In a game, we may <u>observe</u> only a subset of the possible *outcomes* of a strategy, depending on starting conditions and strategies from other agents

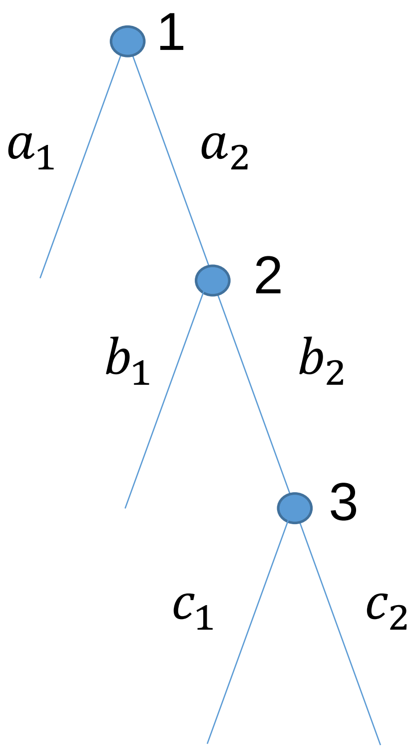

- Strategies that give the same outcome lead to the same payoff
- *Reduced strategy set:* the set formed by all pure strategies that lead to indistinguishable outcomes
- Let the pure strategy set be $\{a_1, a_2\}$, the behavior specifies using a_1 with probability p, and a_2 with probability $p-1$
- A **mixed strategy** β specifies the probability $p(s)$ with which each of the pure strategies $s \in S$ are used

- **Payoff** for using β (for a single agent): $\pi(\beta) = \sum_{a \in A} p(a)\pi(a)$
- Payoff in an *uncertain world*: $\pi(\beta|x) = \sum_{a \in A} p(a)\pi(a|x)$, x is the state

Strategies (Policies)

- Strategy: tells a player what to do for every possible situation throughout the game (complete algorithm for playing the game). It can be ***deterministic*** or ***stochastic***

- Strategy set: what strategies are available for the players to play. The set can be ***finite*** or ***infinite*** (e.g., beach war game)

- Strategy profile: a set of strategies for <u>all players</u> which fully specifies all actions in a game. **A strategy profile must include one and only one strategy for every player**

- Pure strategy: one specific element from the strategy set, a single strategy which is played 100% of the time (***deterministic***)

- Mixed strategy: assignment of a probability to each pure strategy. Pure strategy ≡ degenerate case of a mixed strategy (***stochastic***)

INFORMATION

- Complete information game: Utility functions, payoffs, strategies and "types" of players are *common knowledge*
- Incomplete information game: Players may not possess full information about their opponents (e.g., in auctions, each player knows its utility but not that of the other players). "*Parameters*" of the game are not fully known

- Perfect information game: Each player, when making any decision, is perfectly informed of all the events that have previously occurred (e.g., chess) **[Full observability]**
- Imperfect information game: Not all information is accessible to the player (e.g., poker, prisoner's dilemma) **[Partial observability]**

Turn-taking vs. Simultaneous moves

- **Static games**
- All players take actions “simultaneously”
- → Imperfect information games
- Complete information
- Single-move games

Morra

- **Dynamic games**
- Turn-taking games
- Fully observable ↔ Perfect Information Games
- Complete Information
- Repeated moves

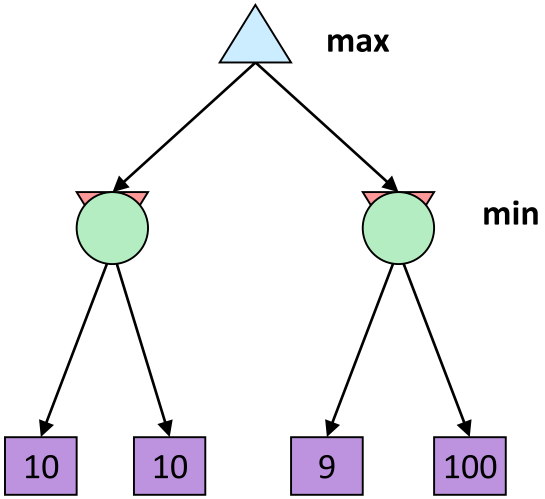

(Strategic-) Normal-Form Game

- Let's focus on <u>static games</u>
- There is a *strategic interaction* among players
- A game in normal form consists of:
  - Set of **players** $N = \{1, \dots, n\}$
  - **Strategy set** S
  - For each $i \in N$, a **utility function** u_i defined over the set of all possible *<u>strategy profiles</u>*, $u_i: S^n \to \mathbb{R}$
  - If each player $j \in N$ plays the strategy $s_j \in S$, the utility of player i is $u_i(s_1, \dots, s_n)$ that is the same as player i's payoff when strategy profile $(s_1, \dots, s_n)$ is chosen

Payoff matrix

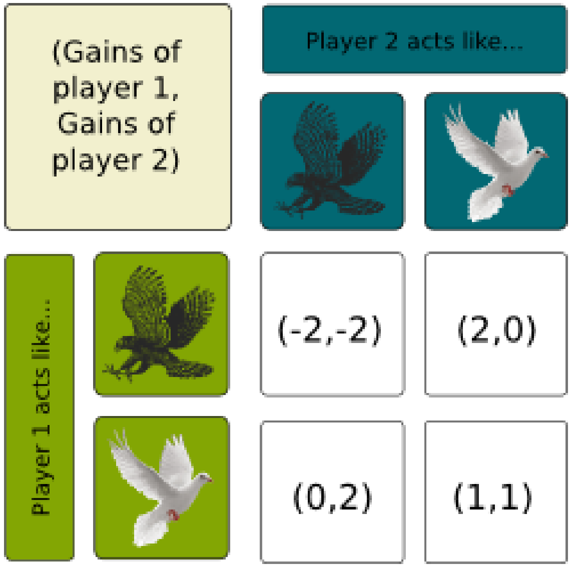

The Ice Cream Wars

- $N = \{1,2\}$
- $S = [0,1]$
- s_i is the fraction of beach
- .....

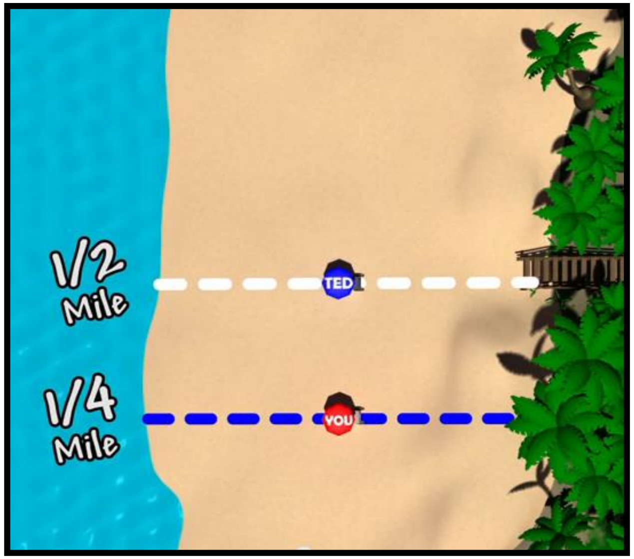

- $u_i(s_i, s_j) = \begin{cases} \frac{s_i+s_j}{2}, & s_i < s_j \\ 1 - \frac{s_i+s_j}{2}, & s_i > s_j \\ \frac{1}{2}, & s_i = s_j \end{cases}$

THE PRISONER'S DILEMMA (1962)

- Two men are charged with a crime
- They can't communicate with each other
- They are told that:
  - If one rats out and the other does not, the rat will be freed, other jailed for 9 years
  - If both rat out, both will be jailed for 6 years
- They also know that if neither rats out, both will be jailed for 1 year

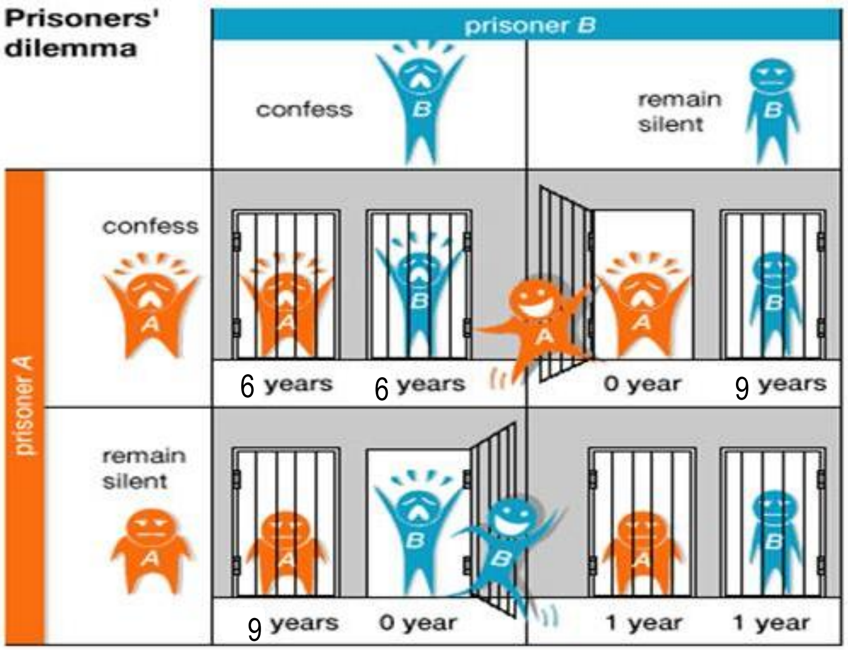

THE PRISONER'S DILEMMA (1962)

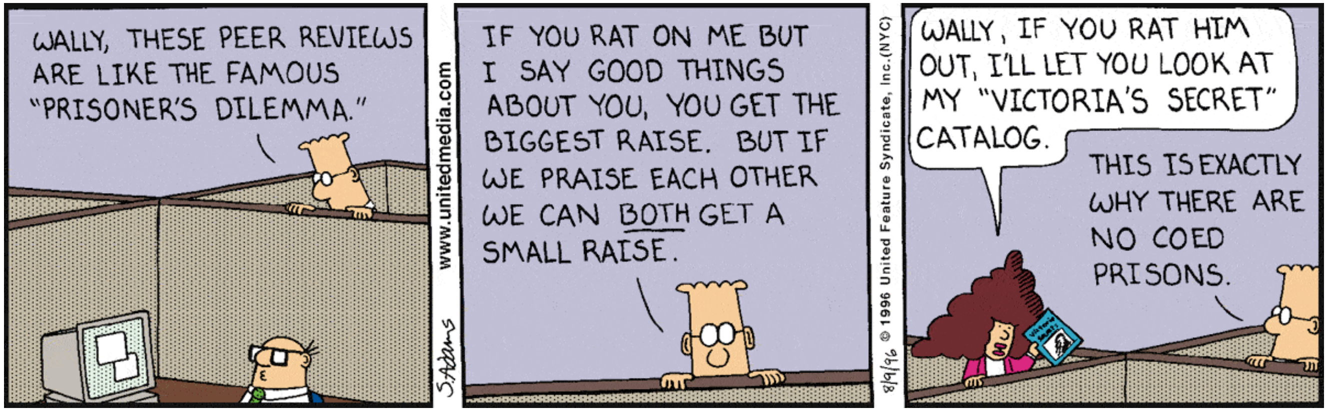

Prisoner's dilemma: Payoff matrix

Don't confess = Don't rat out
Cooperate with each other

Confess = Defect
Don't cooperate to each other, act selfishly!

A \ B	Don't Confess	Confess
Don't Confess	-1,-1	-9,0
Confess	0,-9	-6,-6

What would you do?

PRISONER'S DILEMMA: PAYOFF MATRIX

		B	
		Don't Confess	Confess
A	Don't Confess	-1,-1	-9,0
	Confess	0,-9	-6,-6

B Don't confess:

- If *A* don't confess, *B* gets -1
- If A confess, *B* gets -9

B Confess:

- If *A* don't confess, *B* gets 0
- If A confess, *B* gets -6

Rational agent B opts to *confess*

PRISONER'S DILEMMA

- Confess (Defection, Acting selfishly) is a dominant strategy for *B*: no matters what A plays, the best reply strategy is always to confess

- **(Strictly) dominant strategy**: yields a player strictly higher payoff,. *no matter which decision(s) the other player(s) choose*

- Weakly: ties in some cases

- Confess is a dominant strategy also for *A*

- *A* will reason as follows: *B*'s dominant strategy is to Confess, therefore, given that we are both rational agents, *B* will also Confess and we will both get 6 years.

Prisoner's dilemma

- But, is the dominant strategy (C,C) the **best** strategy?

		B: Don't Confess	B: Confess
A	Don't Confess	-1,-1	-9,0
A	Confess	0,-9	-6,-6

PARETO OPTIMALITY VS. EQUILIBRIA

- Being selfish is a dominant strategy, but the players can do much better by cooperating: (-1,-1), which is the *Pareto-optimal outcome*

- **Pareto optimality**: an outcome such that there is no other outcome that makes every player at least as well off, and at least one player strictly better off
  → Outcome (Don’t Confess, Don’t confess): (-1,-1)

- A strategy profile forms an **equilibrium** if no player can benefit by switching strategies, *given that every other player sticks with the same strategy*, which is the case of (Confess, Confess)

- An equilibrium is a **local optimum in** the space of the strategies

Understanding the Dilemma

- (Self-interested & Rational) agents would choose a strategy that does not bring the maximal reward
- The ***dilemma*** is that the equilibrium outcome is worse for both players than the outcome they would get if both refuse to confess
- Related to the tragedy of the commons

https://en.wikipedia.org/wiki/Tragedy_of_the_commons

On TV: Golden Balls

http://youtu.be/S0qjK3TWZE8

- If both choose Split, they each receive half the jackpot.
- If one chooses Steal and the other chooses Split, the Steal contestant wins the entire jackpot.
- If both choose Steal, neither contestant wins any money
- Watch the video!

The professor's dilemma

		Class: Listen	Class: Sleep
Professor	Make effort	$10^6, 10^6$	-10,0
	Slack off	0,-10	0,0

Dominant strategies?

Nope, if Class listen, and Professor slacks off, Sleep provides a higher payoff!
No *dominant* strategy: best strategy it doesn't matter what other player's strategy

NASH EQUILIBRIUM (1951)

- Can we find an **equilibrium** also in absence of a dominant strategy?
- At equilibrium, each player's strategy is a best response to strategies of others
- Formally, a Nash equilibrium is *strategy profile* $s = (s_1 \dots, s_n) \in S^n$ such that:

$$\forall i \in N, \forall s_i' \in S, u_i(s) \geq u_i(s_i', s_{-i})$$

John F. Nash,
Nobel Prize in Economics, 1994

(NOT) NASH EQUILIBRIUM

A beautiful mind, the movie about (?) John Nash

http://youtu.be/CemLiSI5ox8

Russel Crowe was wrong

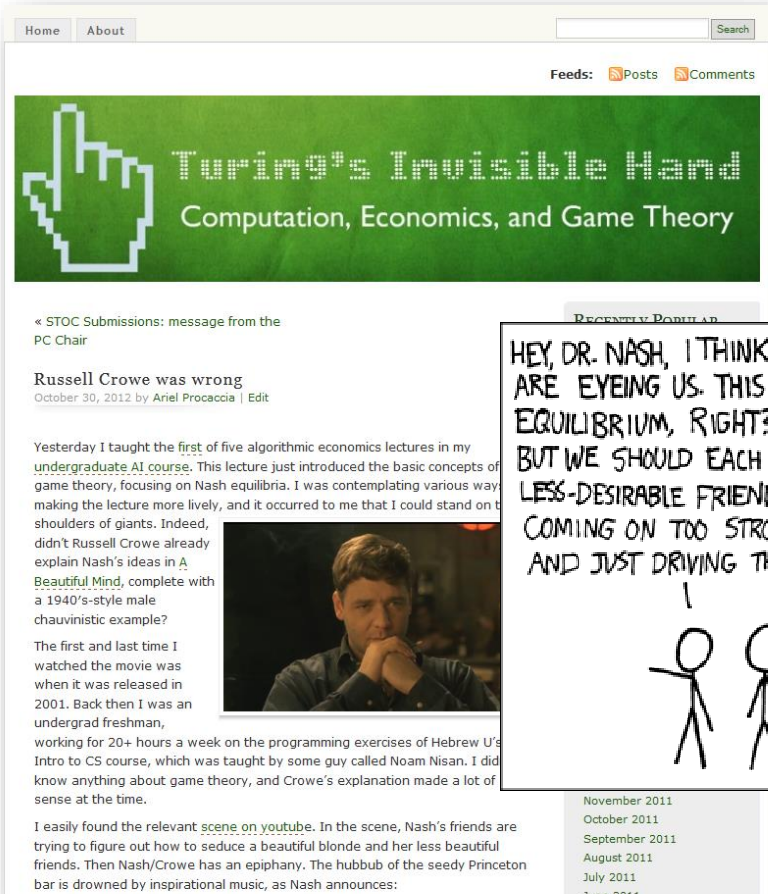

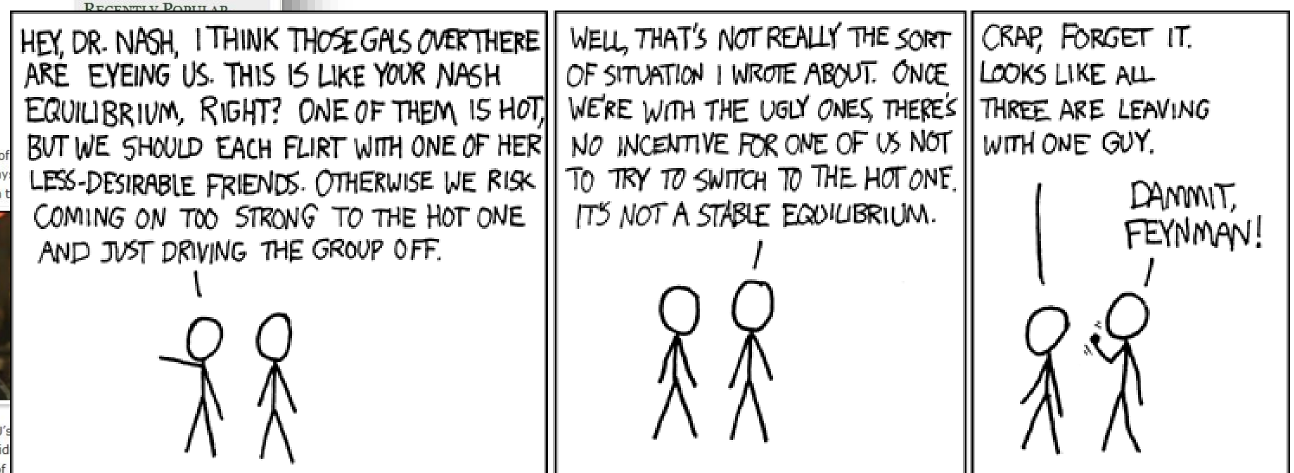

End of the Ice Cream Wars

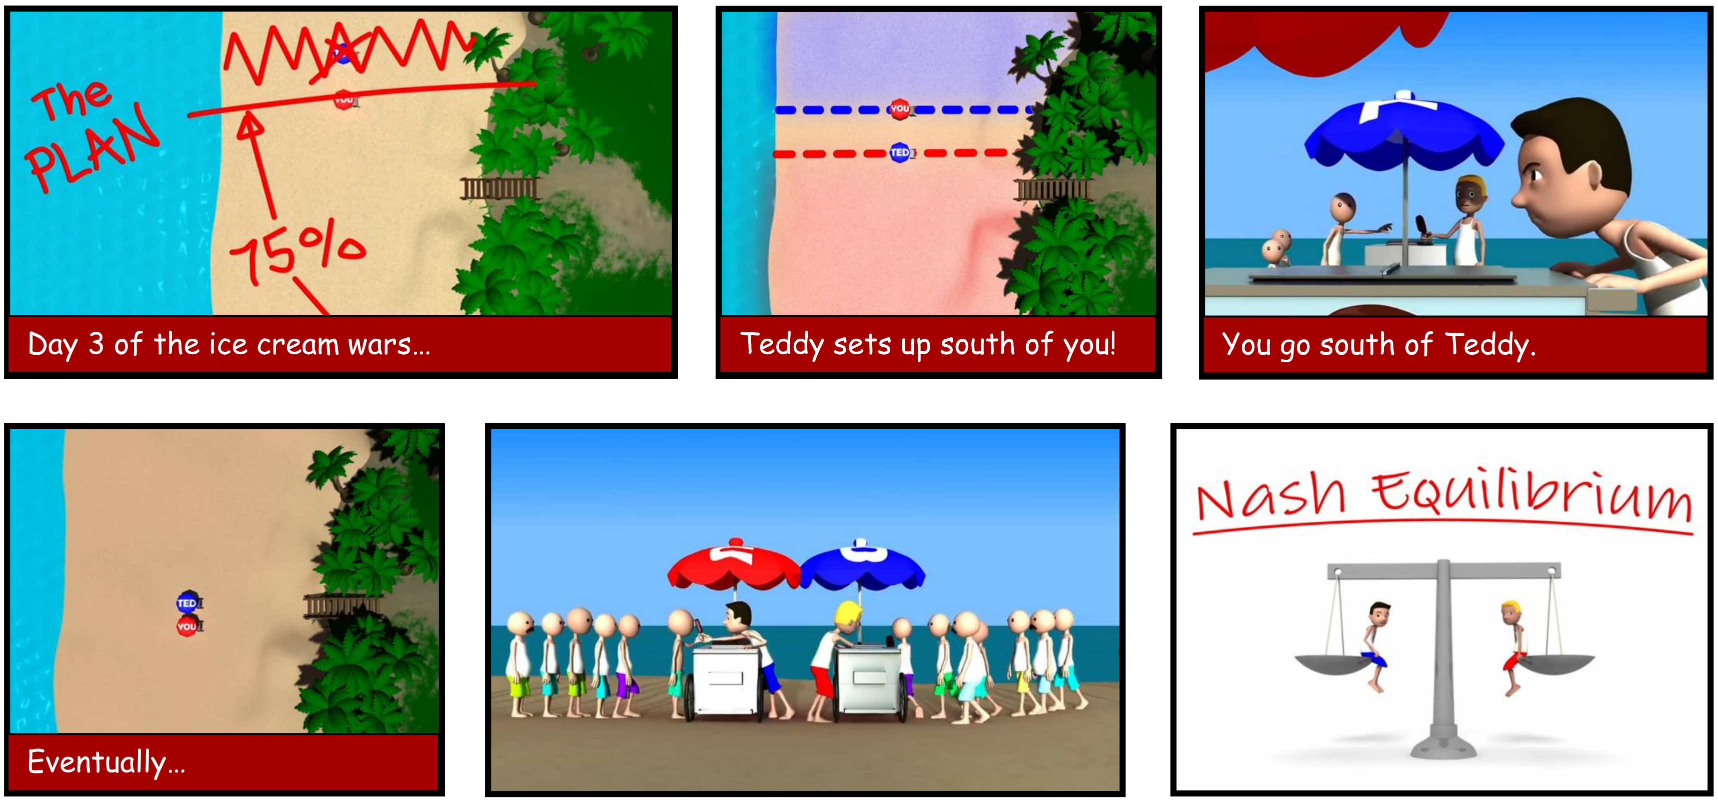

Shops’ logistics …