Game Theory

Lecture 14

Introduction to Dynamic Computation of Game-Theoretic Solution Concepts

&

Some Fundamental Limits on Adaptive Heuristics

A RATIONAL CLASSIFICATION OF DYNAMICS

- We consider dynamic models where the same game is played repeatedly over time.
- One can roughly classify dynamic models in game theory and economic theory into three classes:
  - ***learning dynamics***,
  - ***evolutionary dynamics***, and
  - ***adaptive heuristics***.

Learning Dynamics

- In a **(*Bayesian*)** ***learning dynamic***, each player starts with a ***prior belief*** on the relevant data (the "state of the world"), which usually includes the game being played and the other players' types and (repeated game) strategies.
- Every period, after observing the actions taken (or, more generally, some information about these actions), each player updates his beliefs (using ***Bayes' rule***).
- He then plays optimally given his updated beliefs.
- Roughly speaking, conditions like "***the priors contain a grain of truth***" guarantee that in the long run play is close to the Nash equilibria of the game.

	C	*D*
C	–1, –1	–4, 0
D	0,–4	–3,–3

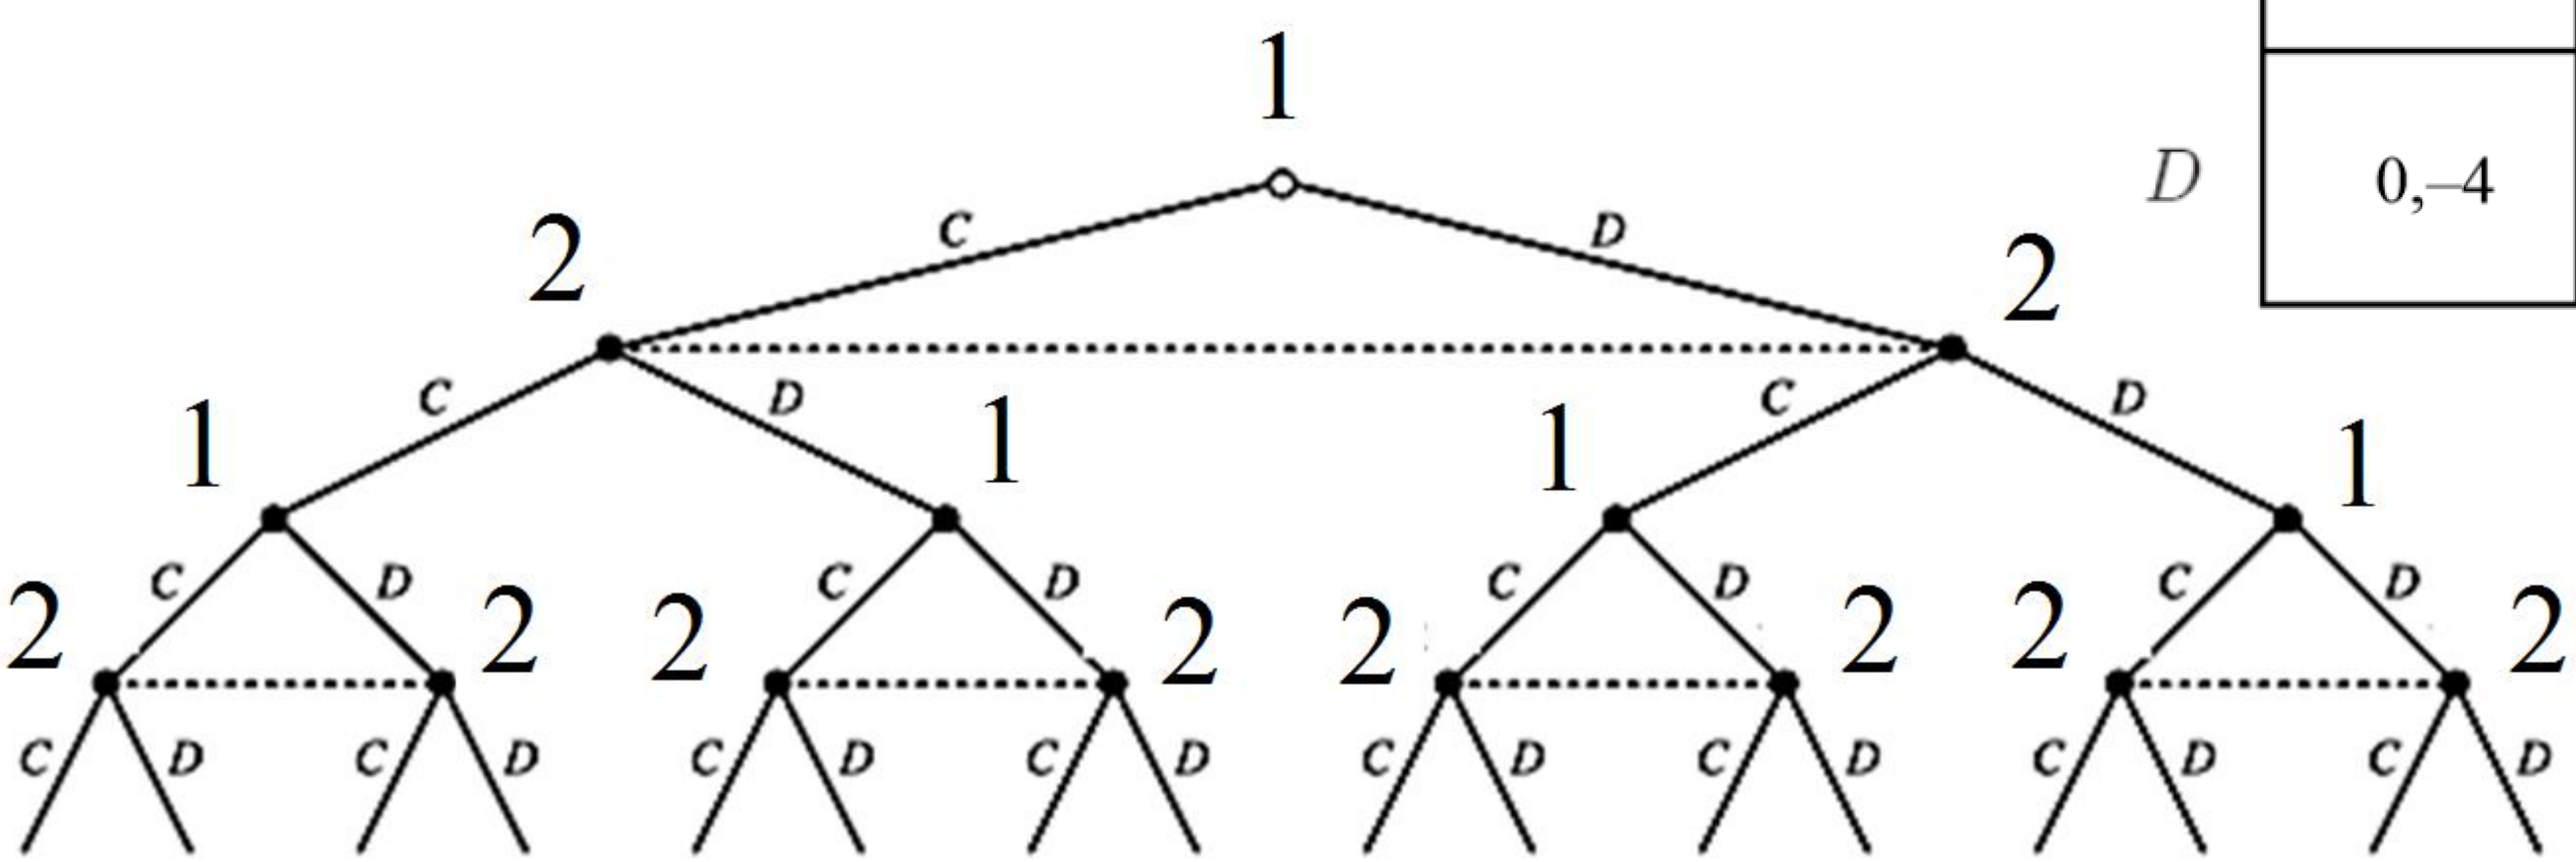

A small portion of repeated prisoner's dilemma

Evolutionary Dynamics:
Adapting game theory to evolutionary games

- Unlike in classical game theory, in evolutionary game theory (EGT), players do not choose their strategy and cannot change it: they are born with a strategy and their offspring inherit that same strategy; in other words,
  - Each such individual always plays the same one-shot action (this fixed action is his "genotype").
- **EGT does not require players to act rationally!**
  - In fact, in the context of the process of evolution, every organism acts as if it were a rational creature, by which we mean a creature whose behavior is directed toward one goal: *to maximize the expected number of its reproducing descendants.*
  - We say that it acts "as if" it were rational in order to stress that the individual organism is not a strategically planning creature.
    - If an organism's inherited properties are not adapted to the struggle for survival, however, it will simply not have descendants.

Evolutionary Dynamics (cont'd)

- If we relate to an organism's number of offspring as a payoff, we have described a process that is propelled by the maximization of payoffs.

- Since the concept of equilibrium in a game is also predicated on the idea that only strategies that maximize expected payoffs (against the strategies used by the other players) will be chosen, we have a motivation for using ideas from game theory in order to explain evolutionary phenomena.

- The focus is on the dynamic process that develops under conditions of many random encounters between individuals in the population, along with the appearance of **random mutations**.

  - A mutation is an individual in the population characterized by a particular behavior:
    - E.g., it may be of type dove, or type hawk. More generally, a mutation can be of type x ($0 \le x \le 1$); that is, the individual will behave as a dove with probability x, and as a hawk with probability $1 - x$.

		Population: Dove (y)	Population: Hawk ($1 - y$)
Mutation	Dove	4, 4	2, 8
	Hawk	8, 2	1, 1

Adaptive Heuristics (Natural Dynamics)

- We use the term ***heuristics*** for rules of behavior that are simple, unsophisticated, simplistic, and myopic (unlike the (Bayesian) "learning" models).
  - These are "rules of thumb" that the players use to make their decisions.
  - We call them ***adaptive*** if they induce behavior that reacts to what happens in the play of the game, in directions that, loosely speaking, seem "better."
    - Thus, always making the same fixed choice, and always randomizing uniformly over all possible choices, are both heuristics.
      - But these heuristics are not adaptive, since they are not at all responsive to the situation (i.e., to the game being played and the behavior of the other participants).
    - In contrast, ***fictitious play*** is a prime example of an adaptive heuristic: at each stage one plays an action that is optimal against the frequency distribution of the past actions of the other players.

Degrees of Rationality

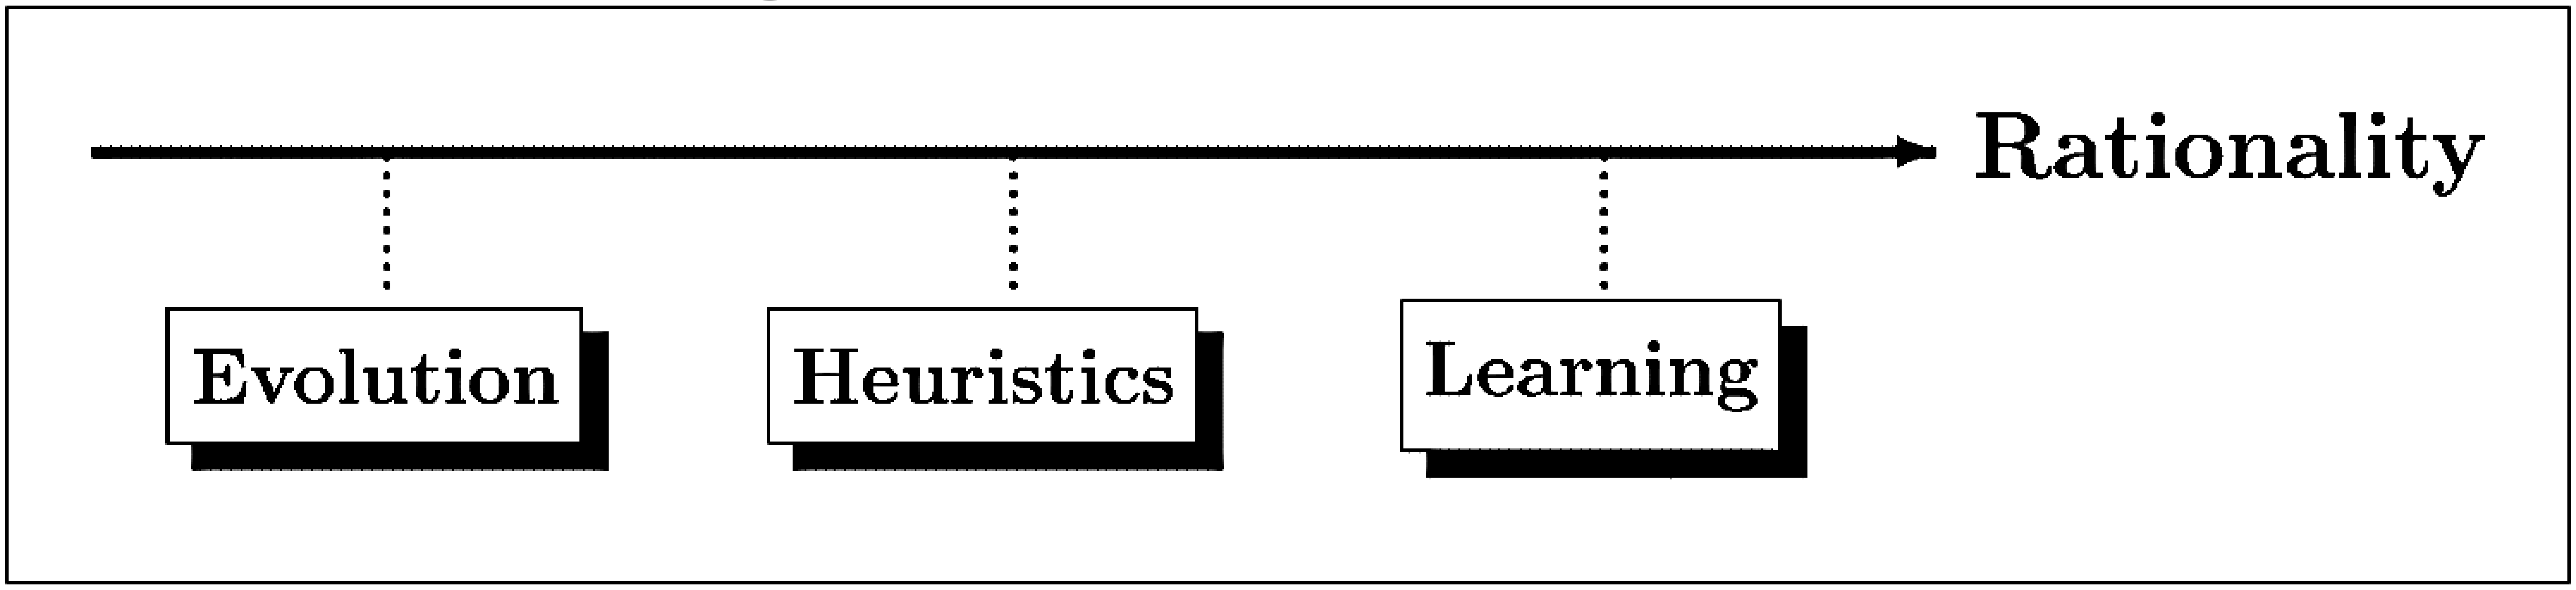

- ***Learning dynamics*** require high levels of rationality. Indeed, repeated-game strategies are complex objects; even more so are beliefs (i.e., probability distributions) over such objects; moreover, in every period it is necessary to update these beliefs, and, finally, to compute best replies to them.

- At the other extreme are ***evolutionary dynamics***. Here the individuals in each population do not exhibit any degree of rationality; their behavior ("phenotype") is completely mechanistic, dictated by their "genotype." They do not compute anything—they just "are there" and play their fixed actions. What may be viewed as somewhat rational is the aggregate dynamic of the population (particularly the selection component), which affects the relative proportions of the various actions.

- ***Adaptive heuristics (a.k.a., natural dynamics)*** lie in between:
  - ➢ on the one hand, the players do perform certain usually simple computations given the environment, and so the behavior is not fixed as in evolutionary dynamics;
  - ➢ on the other hand, these computations are far removed from the full rationality and optimization that is carried out in learning models.

Some Fundamental Limits on Adaptive Heuristics

Fundamental limit on dynamics

FACT

There are no general, natural dynamics leading to Nash equilibrium

- *"general"* : **in all games**
  rather than: in specific classes of games:
  - two-person zero-sum games
  - potential games
  - supermodular games
  - . . .

- *"leading to Nash equilibrium"* :
  at a Nash equilibrium (or close to it) from some time on

- *"natural"* :
  - **adaptive** (reacting, improving, ...)
  - **simple and efficient**:
    - **computation** (performed at each step)
    - **time** (how long to reach equilibrium)
    - **information** (of each player)

  bounded rationality

Some Examples of Unnatural Dynamics

Dynamics that are NOT ***"natural"*** :

- **exhaustive search**
  (deterministic or stochastic)
- using a **mediator**
- **broadcasting** the private information and then performing **joint** computation
- **fully rational learning**
  (prior beliefs on the strategies of the opponents, Bayesian updating, optimization)

Natural Dynamics: Information

UNCOUPLED DYNAMICS :

Each player knows *only* his own payoff (utility) function

(does *not* know the payoff functions of the other players)

(privacy-preserving, decentralized, distributed ...)

Review of Notation

N-person game in strategic (normal) form:

- **Players**

$$i = 1, 2, \ldots, N$$

- For each player i: **Actions**

$$a^i \text{ in } A^i$$

- For each player i: **Payoffs (utilities)**

$$u^i(a) \equiv u^i(a^1, a^2, \ldots, a^N)$$

- **Time**

$$t = 1, 2, \ldots$$

- At period t each player i chooses an **action**

$$a_t^i \text{ in } A^i$$

according to a probability distribution

$$\sigma_t^i \text{ in } \Delta(A^i)$$

Uncoupled Dynamics: Generic Definition

Fix the set of players $1, 2, \ldots, N$ and their action spaces $A^1, A^2, \ldots, A^N$

- A general dynamic:

$$\sigma_t^i \equiv \sigma_t^i(\text{HISTORY}\,;\ \text{GAME})$$

$$\equiv \sigma_t^i(\text{HISTORY}\,;\ u^1, \ldots, u^i, \ldots, u^N)$$

- An **UNCOUPLED** dynamic:

$$\sigma_t^i \equiv \sigma_t^i(\text{HISTORY}\,;\ u^i)$$

- Simplest **uncoupled dynamics**:

$$\sigma_t^i \equiv f^i(a_{t-1}; u^i)$$

where $a_{t-1} = (a_{t-1}^1, a_{t-1}^2, \ldots, a_{t-1}^N) \in A$ are the actions of all the players in the previous period

- Only last period matters (*"1-recall"*)
- Time t does not matter (*"stationary"*)

Impossibility Result

Theorem. *There are* **NO** *uncoupled dynamics with 1-recall*

$$\sigma_t^i \equiv f^i(a_{t-1}; u^i)$$

that yield almost sure convergence of play to pure Nash equilibria of the stage game in all games where such equilibria exist.

Consider the following two-person game, which has a unique pure Nash equilibrium **(R3,C3)**

	C1	C2	**C3**
R1	1,0	0,1	1,0
R2	0,1	1,0	1,0
R3	0,1	0,1	1,1

Assume ***by way of contradiction*** that we are given an uncoupled, 1-recall, stationary dynamic that yields almost sure convergence to pure Nash equilibria when these exist

Impossibility Result (Cont'd)

- Suppose the play at time $t-1$ is **(R1,C1)**
- **ROWENA is best replying at (R1,C1)**
  - by *1-recall* , *stationarity* and *uncoupledness*,
    - **ROWENA is best replying at $t-1$**
    - **$\Rightarrow$ ROWENA will play the same action at t**

      $\Rightarrow$ ROWENA will play R1 also at t
- Similarly for COLIN:

  A player who is best replying cannot switch

	C1	C2	C3
R1	1,0 ↔	0,1 ↕	1,0 ↔
R2	0,1 ↕	1,0 ↔	1,0 ↔
R3	0,1 ↕	0,1 ↕	1,1

$\Rightarrow$ (R3,C3) **cannot be reached**
(unless we start there)

2-Recall: Possibility

Theorem. **THERE EXIST** *uncoupled dynamics with* **2-RECALL**

$$\sigma_t^i \equiv f^i(a_{t-2}, a_{t-1}; u^i)$$

that yield almost sure convergence of play to pure Nash equilibria of the stage game in every game where such equilibria exist.

Define the strategy of each player i as follows:

IF:

- Everyone played the same in the previous two periods: $a_{t-2} = a_{t-1} = a$; and
- Player i best replied: $a^i \in \mathrm{BR}^i(a^{-i}; u^i)$

THEN: At t player i *plays a^i again*: $a_t^i = a^i$

ELSE: At t player i *randomizes uniformly over A^i*

"Good":

- simple

"Bad":

- exhaustive search
- all players must use it
- takes a long time

How long to equilibrium?

- An uncoupled dynamic leading to Nash equilibria is **TIME-EFFICIENT** if the **TIME IT TAKES** is **POLYNOMIAL** in the number of players (rather than: *exponential*)

Theorem. *There are* **NO TIME-EFFICIENT** *uncoupled dynamics that reach a pure Nash equilibrium in all games where such equilibria exist.*

In fact: exponential, like exhaustive search

- **Perhaps we are asking too much?**
- For instance, the size of the data (the payoff functions) is *exponential* rather than polynomial in the number of players

- The exponential bounds for Nash equilibrium procedures are not due just to the complexity of the input, i.e., to the payoff functions being of exponential size, but rather to the intrinsic complexity of reaching Nash equilibria:
  - **NASH EQUILIBRIUM**: a ***fixed-point*** of a non-linear map
  - **CORRELATED EQUILIBRIUM**: a solution of finitely many ***linear inequalities***

So, what about correlated equilibria?

THERE EXIST ***general, natural dynamics***
leading to CORRELATED EQUILIBRIA

- ***"general"***: in all games
- ***"natural"***:
  - **adaptive** (also: close to "behavioral")
  - **simple and efficient**:
    computation, time, information
- ***"leading to correlated equilibria"***:
  statistics of play become close to
  CORRELATED EQUILIBRIA

- It has been proved that the number of steps need to reach an approximate correlated $\varepsilon -$ equilibria is polynomial rather than exponential in the number of players.

- ***Regret Matching***
- General regret-based dynamics
- **"REGRET"**: the increase in past payoff, if any,
  if a different action would have been used

- **"MATCHING"**: switching to a different action
  with a probability that is proportional to the
  regret for that action

Summary

Nash Equilibrium

"DYNAMICALLY DIFFICULT"

Correlated Equilibrium

"DYNAMICALLY EASY"

References

- S. Hart, "Adaptive Heuristics," Econometrica, Vol. 73, No. 5, pp. 1401–1430, September, 2005.

- S. Hart and Y. Mansour, "How long to equilibrium? The communication complexity of uncoupled equilibrium procedures," *Games and Economic Behavior*, Vol. 69, pp. 107–126, 2010.

- S. Hart and A. Mas-Colell, "Stochastic uncoupled dynamics and Nash equilibrium," Games and Economic Behavior, Vol. 57, pp. 286–303, 2006.