Basics of Evolutionary Game Theory

Last modified: September 20, 2012

Static Simultaneous Move Game

an N player game $i = 1 \ldots N$, $P(S)$ are probability measure on S

finite strategy spaces, $\sigma^i \in \Sigma^i \equiv P(S^i)$ are mixed strategies

$s \in S \equiv \times_{i=1}^{N} S^i$ are the strategy profiles

$\sigma \in \Sigma \equiv \times_{i=1}^{N} \Sigma^i$

other useful notation $s^{-i} \in S^{-i} \equiv \times_{j \neq i} S^j$

$\sigma^{-i} \in \Sigma^{-i} \equiv \times_{j \neq i} \Sigma^j$

$u^i(s) = u^i(s^i, s^{-i})$ payoff or utility

$u^i(\sigma) \equiv \sum_{s \in S} u^i(s) \prod_{j=1}^{N} \sigma^j(s^j)$ is expected utility

Nash Equilibrium

players correctly anticipate on another's strategies

σ is a *Nash equilibrium* profile if for each $i \in 1, \ldots N$
$u^i(\sigma) = \max_{\tilde{\sigma}^i} u^i(\tilde{\sigma}^i, \sigma^{-i})$

Theorem: a Nash equilibrium exists in a finite game

Disequilibrium Adjustment

- not a Nash equilibrium: someone has erroneous beliefs
- dynamics driven by error correction: erroneous beliefs should be changed

Individual "Learning" Models

beliefs modified through experience – who do you play?

- playing repeatedly against a fixed opponent with or without myopia
- pick a players at random from a large population everyone sees play
- players randomly matched, results of all matches revealed anonymously
- players matched randomly see results only of own match (this is how experiments are usually conducted)

specify the beliefs of each individual and how they adjust beliefs and behavior

example: best-response dynamic – everyone plays best response to previous periods play

Best Response Dynamic

expectations: $\overline{\sigma}_{t+1}^{-i} = \sigma_t^{-i}$

$s_{t+1}^i = B^i(\sigma_t^{-i})$

Matching Pennies

	H	T
H	1,-1	-1,1
T	-1,1	1,-1

- error driven cycles
- “cob-web”
- not many people would play this way…

Partial Adjustment

Best-response is too abrupt – consider the cob-web cycle

- Partial best-response: adjust in direction of improving payoff based on previous period play

expectations: $\sigma_{t+1}^i = \alpha B^i(\sigma_t^{-i}) + (1-\alpha)\sigma_t^i$

continuous time:

$$\sigma_{t+1}^i - \sigma_t^i = \alpha\left(B^i(\sigma_t^{-i}) - \sigma_t^i \right)$$

$$\alpha = a\Delta$$

$$\dot{\sigma}_t^i \approx \frac{\sigma_{t+1}^i - \sigma_t^i}{\Delta} = a\left(B^i(\sigma_t^{-i}) - \sigma_t^i \right)$$

- Fictitious play: play best response to a long-term average

Shapley example

0,0	1,2	2,1
2,1	0,0	1,2
1,2	2,1	0,0

note that (0,0) is never hit, but always in Nash equilibrium

Population Models

- evolution: better strategies do better/ random mutation
- population model: specify fraction of the population changing to a “better strategy” based on some measure of population performance
- partial best response can be a population model
- another example: replicator – strategies that are doing better than average grow
- population and individual approaches are generally compatible: every individual model gives rise to a population model, and most population models are compatible with sensible individual behavior
- it is possible to specify population models that don't make sense at the individual level (genetic algorithms)

Replicator

- strategies that are doing better than average grow

$$\frac{\dot{\sigma}_t^i(s^i)}{\sigma_t^i(s^i)} = \alpha(u^i(s^i,\sigma_t^{-i}) - u^i(\sigma_t))$$

$$\dot{\sigma}_t^i(s^i) = \alpha\sigma_t^i(s^i)(u^i(s^i,\sigma_t^{-i}) - u^i(\sigma_t))$$

- steady states at "relative best response"
- relative = relative to those strategies actually used
- as a stimulus-response model
- probability matching issues
- as a model of social learning

One-dimensional case

Two player, two action symmetric game

There is only one sensible dynamic: move in the direction of increasing individual payoffs

Driving game: 1 if agree, 0 if disagree

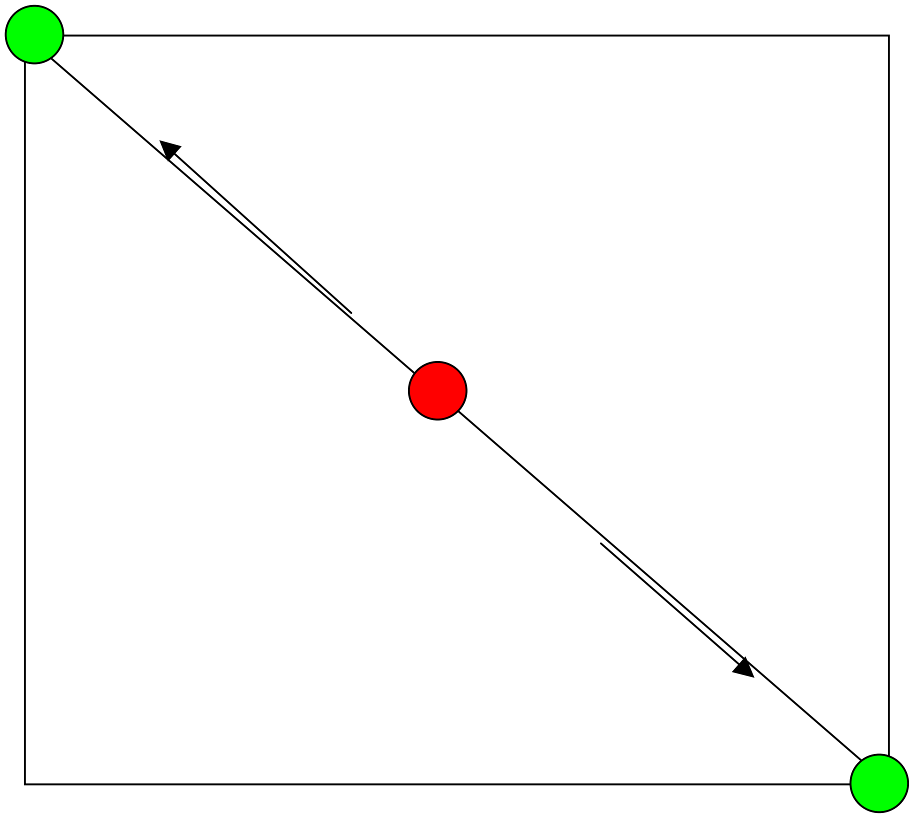

Anti-driving game: get 1 if disagree, 0 if agree

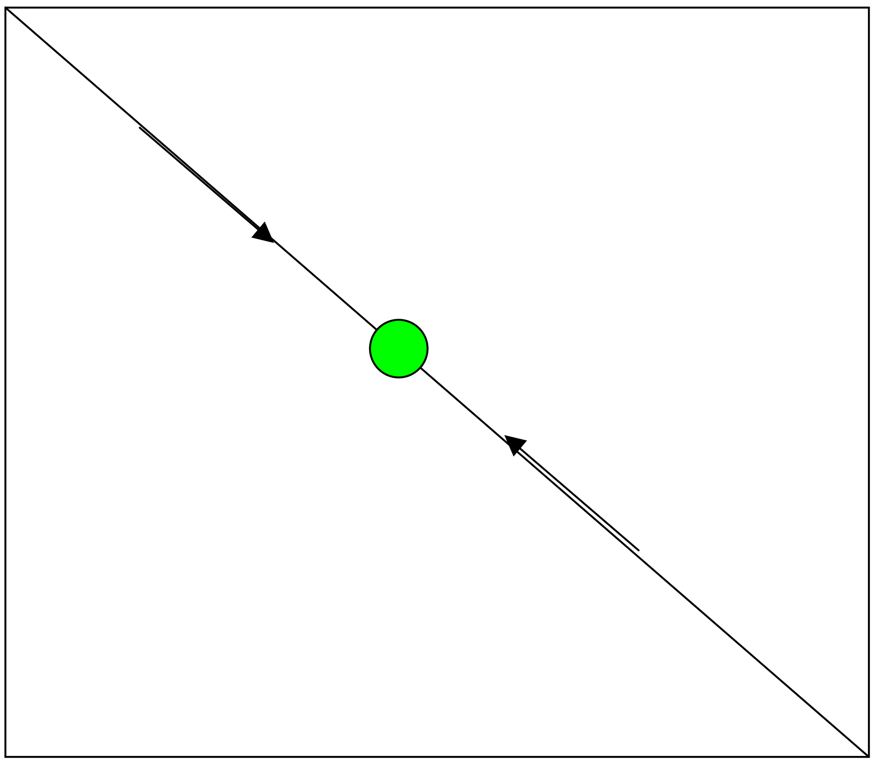

Stochastic Evolutionary Model

Kandori, Mailath, Rob and Young

finite population of N players

state variable σ_t

- deterministic dynamic – discrete time replicator or partial best response
- mutations: with probability ε one player is randomly chosen to "mutate" on to randomly chosen strategy
- everyone else follows deterministic dynamic
- induces a Markov process $M(\varepsilon)$ on the state space Σ

The Markov Process

- For $\varepsilon > 0$ the process $M(\varepsilon)$ is aperiodic and irreducible and hence has a unique invariant distribution $\mu(\varepsilon)$
- When $\varepsilon = 0$ all steady states (Nash equilibria usually) and asymptotic cycles of the deterministic dynamic are ergodic classes - we denote them by $\Sigma(0)$
- Resistance and regularity:

a scalar valued function $Q(\varepsilon)$ is *regular* if $r[Q] \equiv \lim_{\varepsilon \to 0} \log Q(\varepsilon) / \log \varepsilon$ exists and $r[Q] = 0$ implies $\lim_{\varepsilon \to 0} Q(\varepsilon) > 0$

- $\log \varepsilon^r / \log \varepsilon = r$

Theorem: $M(\varepsilon)$ is regular

Theorem (Young): $\mu(0) \equiv \lim_{\varepsilon \to 0} \mu(\varepsilon)$ exists and puts weight only on $\Sigma(0)$

Stochastically Stable Sets

- *Stochastically stable sets* are those points in $\Sigma(0)$ that get positive weight according to $\mu(0)$
- The point being of course that in general not all of $\Sigma(0)$ is stochastically stable
- Description of what the Markov process looks like for ε small

The Resistance of Trees

- T is a tree whose nodes are in the set $\Sigma(0)$ with any set of edges
- $D(\sigma)$ is the unique node from σ in the direction of the root
- a σ -tree is a tree whose root is σ , denoted $T(\sigma)$
- for any two points σ_0, σ_t in $\Sigma(0)$a path from σ_0 to σ_t is a sequence of points $\sigma_0, \sigma_1, \ldots \sigma_{t-1}, \sigma_t$ where the transition from σ_τ to $\sigma_{\tau+1}$ has positive probability for $\varepsilon > 0$
- the resistance of a path is the sum of resistances between points in the path $\sum_{\tau=0}^{t-1} r(\tau, \tau+1)$
- the resistance $r(\sigma_0, \sigma_t)$ is the least resistance of any path between
- the resistance $r(T(\sigma))$ of the σ -tree $T(\sigma)$ is the sum over non-root nodes σ_τ of $r(\sigma_\tau, D(\sigma_\tau))$

Least Cost Trees

$$r(T(\sigma)) = \sum_{\sigma_\tau \in \Sigma(0) \backslash \sigma} r(\sigma_\tau, D(\sigma_\tau))$$

the resistance $r(\sigma)$ is the least resistance of all σ -trees

Theorem (Young): σ is a stochastically stable if and only if $\sigma \in \Sigma(0)$ and $r(\sigma) = \min_{\sigma_\tau \in \Sigma(0)} r(\sigma_\tau)$

Example with Three Nodes

Resistances row to column

	A	B	C
A		6	5
B	4		3
C	2	1	

Side Trees

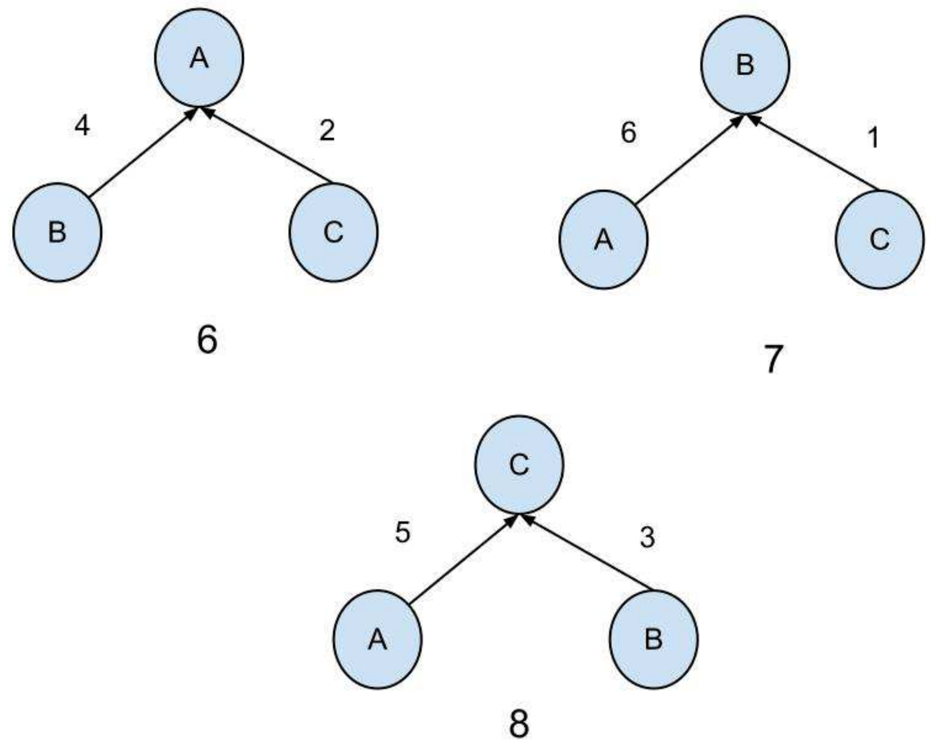

Cross Trees

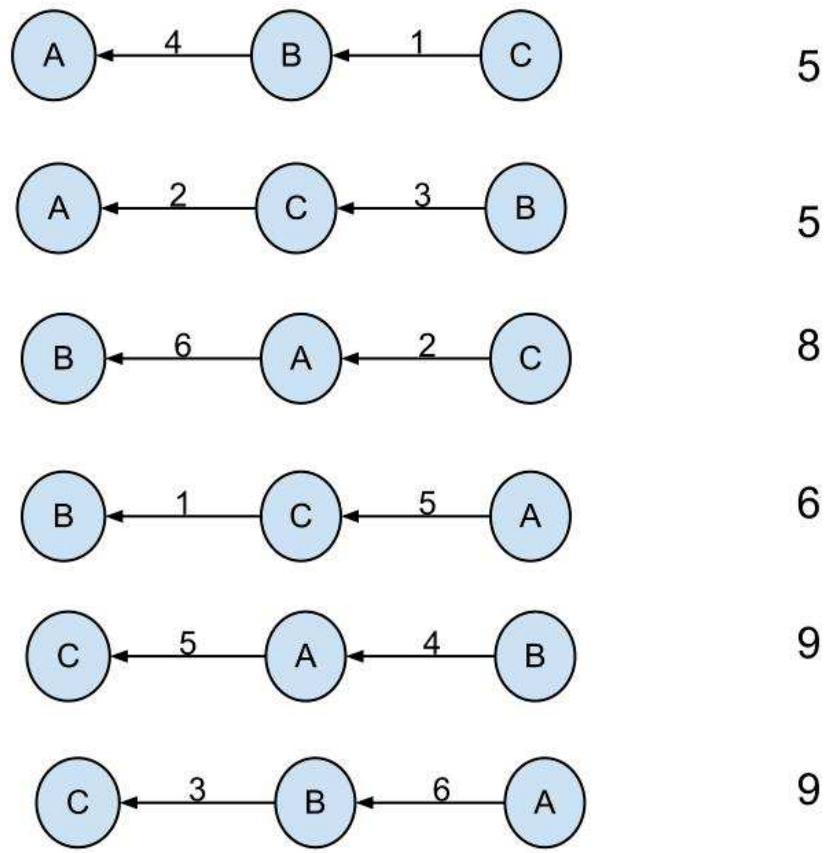

Half Dominance

In a symmetric game a pure strategy Nash equilibrium the symmetric strategy $s=(s^i, s^i, \ldots, s^i)$ is ½-dominant if it is a strict best response for everyone to play s^i when the probability of all other players snce and proofimultaneously playing s^i is at least ½

A half dominant Nash equilibrium is stochastically stable

Proof by tree trimming

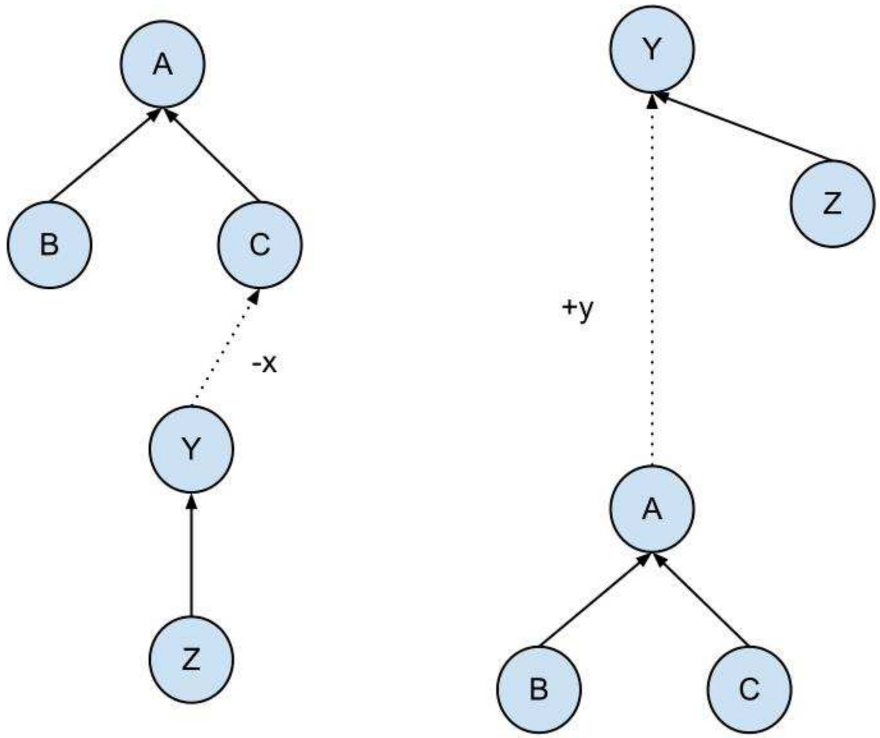

Need $y < x$

Stag-Hunt and 2x2 Coordination Games

$1 > x > y > 0$

x, x	$y, 0$
$0, y$	1,1

indifference $p = pr(x) = (1 - y)/(x - y + 1)$

1 is pareto efficient

x is half dominant if and only if $p < 1/2$

i.e. $1 < x + y$

for example $x = 3/4, y = 1/2$

Relative Waiting Times

- $N(1-p)$ mutations to go from $x \rightarrow 1$
- $\varepsilon^{N(1-p)}$, waiting time inversely proportional to this
- Np mutations to go from $1 \rightarrow x$
- ε^{Np}, waiting time inversely proportional to this

Radius and Co-radius

Radius: least number of mutations to get out to a different point in $\Sigma(0)$

Coradius: least number of mutations to get back from any point in $\Sigma(0)$

Radius > coradius implies stochastic stability

Comments

Nachbar: it can take a long time to learn to eliminate dominated strategies (deterministic dynamic)

Ellison: the very long run can be very long, but much shorter with local interaction