

Second School on the
MATHEMATICS OF ECONOMICS
(21 August - 1 September 2000)

Learning in Games

D. Levine
(UCLA)

Dynamics in Games

- Review of statics
- Equilibrium dynamics
- “disequilibrium dynamics:” learning and evolution
- convergence to equilibrium

Nash Equilibrium Statics

Static Simultaneous Move Game

an N player game $i = 1 \dots N$, $P(S)$ 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 $s_{-i} \in S_{-i} \equiv \times_{j \neq i} S_j$

$$\sigma_{-i} \in \Sigma_{-i} \equiv \times_{j \neq i} \Sigma_j$$

$u_i(s)$ 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 $i \in 1, \dots, N$

$$u_i(\sigma) = \max_{\sigma'_i} u_i(\sigma'_i, \sigma_{-i})$$

Theorem: a Nash equilibrium exists in a finite game

this is more or less why Kakutani's fixed point theorem was invented

$B_i(\sigma)$ is the set of best responses of i to σ_{-i} , and is UHC convex valued

this is a very passive notion of equilibrium, which often causes confusion

Examples

Prisoner's Dilemma Game

	cooperate	defect
cooperate	2,2	0,3
defect	3,0	1,1

a unique dominant strategy equilibrium (D,L)

Matching Pennies

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

A unique mixed strategy equilibrium at 50-50

Multiple equilibria

Pure Coordination Game: Driving on the Right or the Left

	R	L
R	1,1	0,0
L	0,0	1,1

three equilibria (R,R) (L,L) ((.5R,.5L),(.5R,.5L))


too many equilibria??

Equilibrium Cycles


Drive on the right on Sundays; on the left during the rest of the week

Dynamics in Economics


Cycles versus Stochastic Fluctuations



Shiller



<http://www-personal.buseco.monash.edu.au/~hyndman/TSDL/physics/star.dat>



chaos theory has done poorly in economics

economists learn from their mistakes – clever newcomers are very good at repeating our old errors

Disequilibrium Nash Adjustment

Expectational Error Adjustment

If we are not at a Nash equilibrium, someone has erroneous beliefs

Dynamics studied by economists are driven by error correction:
erroneous beliefs should be changed

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...

Population versus individual models

Individual model: specify the beliefs of each individual and how they adjust beliefs and behavior

“learning model”

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

population model: specify fraction of the population changing to a “better strategy” based on some measure of population performance

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)

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

This has advantages in the population setting, because the state can be measure by what people are doing in the current period.


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

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



Multi-dimensional cycling: Shapley game, Jordan game

Shapley Game

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

Jordan Game

Three player matching pennies; player 1 wants to match player 2, who wants to match player 3 who wants to do the opposite of player 1

“Classical” Case of Fictitious Play

- keep track of frequencies of opponents' play
- begin with an initial or prior sample
- play a best-response to historical frequencies
- not well defined if there are ties, but for generic payoff/prior there will be no ties

- optimal procedure against i.i.d. opponents
- how well does fictitious play do if the i.i.d. assumption is wrong?

How well can fictitious play do in the long-run?

- notice that fictitious play only keeps track of frequencies: can fictitious play do as well in the long-run as if those frequencies (but not the order of the sample) was known in advance?
- alternatively: suppose that a player is constrained to play the same action in every period, so that he does not care about the order of observations

Universal Consistency

let u_t^i be actual utility at time t

let ϕ_t^{-i} be frequency of opponents' play (joint distribution over S^{-i})

suppose that for *all* (note that this does not say “for almost all”) sequences of opponent play

$$\liminf_{T \rightarrow \infty} (1/T) \sum_{t=1}^T u_t^i - \max_{s^i} u^i(s^i, \phi_T^{-i}) \geq 0$$

then the learning procedure is universally consistent

*Is fictitious play universally consistent? Fudenberg and Kreps example
– the anti-driving game*

0,0	1,1
1,1	0,0

this coordination game is played by two identical players

suppose they use *identical deterministic* learning procedures

then they play UL or DR and get 0 in every period

this is not individually rational, let alone universally consistent

Theorem [Monderer, Samet, Sela; Fudenberg, Levine]: fictitious play is consistent provided the frequency with which the player switches strategies goes to zero

Smooth Fictitious Play

instead of maximizing $u^i(s^i, \phi_{t-1}^i)$ maximize

$$u^i(\sigma^i, \phi_{t-1}^i) + \lambda v^i(\sigma^i)$$

where v^i is smooth, concave and has derivatives that are unbounded at the boundary of the unit simplex

example: the *entropy*

$$v^i(\sigma^i) = -\sum_{s^i} \sigma^i(s^i) \log \sigma^i(s^i)$$

as $\lambda \rightarrow 0$ this results in an approximate optimum to the original problem

however the solution to $u^i(\sigma^i, \phi_{t-1}^i) + \lambda v^i(\sigma^i)$ is smooth and interior
(always puts positive weight on all pure strategies)

- *Theorem [Blackwell, Hannan, Fudenberg and Levine and others]:*
smooth fictitious play is ε universally consistent with $\varepsilon \rightarrow 0$ as $\lambda \rightarrow 0$

Notice that these procedures is random

It must be: otherwise it can be defeated by a clever opponent

Think: matching pennies

But: Universal consistency is a long way from Nash Equilibrium...

Complex Structure of Equilibria and Computation of Equilibria

- Many isolated points or segments possibly separated by cycles
- Is it possible to define algorithms guaranteed to converge to a Nash equilibrium (move directly to the closest equilibrium!)
- Algorithms for finding equilibria (Scarf, Global Newton) are not decentralized
- Cycles aren't believable
- Can the system wander around forever without a pattern so that people are always confused? Note that chaotic systems are very predictable in the short run.
- A global convergence theorem would be more plausible if the target was geometrically simpler (for example, a convex set)

Correlated Equilibrium

Definition

r a probability distribution over S

i.e. may allow correlations

$u_i(r)$ expected utility

$r_{-i}(\cdot | s_i)$ conditional distribution of everyone else

$r_i(s_i)$ marginal over player i

r is a *correlated equilibrium* if $r_i(s_i) > 0$ implies

$$u_i(s_i, r_{-i}(\cdot | s_i)) \geq u_i(s_i', r_{-i}(\cdot | s_i))$$

Interpretation

- Nash equilibrium with a correlating device
- Correlating device “gives each player a private recommendation”
- Given the recommendation, it is optimal to follow it
- Any equilibrium in a game with (possibly partially private) correlating device isomorphic to a correlated equilibrium
- Since correlating devices exist in nature, it is hard to accept the logic of Nash equilibrium without accepting the logic of correlated equilibrium
- Important class of correlated equilibria: public randomizations over Nash equilibrium – for example flip a coin to see if we drive on the right or the left

Example

Chicken

6,6	2,7
7,2	0,0

three Nash equilibria (2,7), (7,2) and mixed equilibrium w/ probabilities (2/3,1/3) and payoffs (4 2/3, 4 2/3)

1/3	1/3
1/3	0

is a correlated equilibrium giving utility (5,5)

Structure

The set of correlated equilibrium strategies is convex

If $r, \tilde{r}$ are correlated equilibria so is $(1 - l)r + l\tilde{r}$ for $0 \leq l \leq 1$

Because equilibrium r is defined by a system of linear inequalities

$$\sum_{s_{-i}} u_i(s_i, s_{-i}) r(s_i, s_{-i}) \geq \sum_{s_{-i}} u_i(s_i', s_{-i}) r(s_i', s_{-i})$$

must hold for every i, s_i, s_i'

note that for the s_i 's with zero probability this just says: $0 \geq 0$.

Calibrated adjustment procedures

- Return to the ideas of universal consistency
- Universal consistency deficient because it does not take account of conditional probabilities
- Goal: if “enough” conditional probabilities are taken account of we get global convergence to the set of correlated equilibria

Conditional Probability Models: Experts

allow time dependent games

$$\liminf_{T \rightarrow \infty} (1/T) \sum_{t=1}^T u_t^i - (1/T) \sum_{t=1}^T \max_{s^i} u_t^i(s^i, s_t^{-i}) \geq 0$$

universal consistency theorem for smooth fictitious play holds, without change in proof

a “model” makes conditional probability forecasts

an “expert” makes recommendations about how to play

$$s_t^i = e^i(h_{t-1}^i)$$

$$\text{set } v_t^i(e^i, s_t^{-i}) = u^i(e^i(h_{t-1}^i), s_t^{-i})$$

conclusion: can do as well as if you knew who the best expert was in advance

Conditional Probability Models: Direct

classify observations into subsamples

countable collection of categories Ψ

classification rule $\psi^i: H \times S \rightarrow \Psi$

$$\psi^i(h_{t-1}^i, s_t^i)$$

$\phi_t^{-i}(\psi)$ empirical distribution of opponent's play conditional on the category ψ ; $n_t(\psi)$ is number of time category has occurred

effective categories: minimal finite subset $\Psi_t \in \Psi$ with all observations through time t

m_t denotes the number of effective categories

Assumption 1: $\lim_{t \rightarrow \infty} m_t / t = 0$

This is essentially the method of sieves

Universal Conditional Consistency

total utility actually received in the subsample ψ is $u_t^i(\psi)$

$$c_t^i(\psi) = \begin{cases} n_t(\psi) \max_{s^i} u^i(s^i, \phi_t^{-i}) - u_t^i(\psi) & n_t(\psi) > 0 \\ 0 & n_t(\psi) = 0 \end{cases}$$

universal conditional consistency

$$\limsup (1/T) \sum_{\psi \in \Psi_t} c_T^i(\psi) \leq 0$$

Non Calibrated Case

categorization rule depends only on history, not on own plans

- 1) given h_{t-1}^i , $\psi(h_{t-1}^i)$ chooses the category
- 2) play a smooth fictitious play against the sample in the chosen category $\phi_{t-1}^{-i}(\psi)$
- 3) add the new observation s_t^{-i} to the category $\psi(h_{t-1}^i)$

Works like smooth fictitious play within each category, so universally conditionally consistent

Calibrated Case

try to use a rule $\psi(h_{t-1}^i, s_t^i)$

focus on special case $\psi(s_t^i), \Psi = S$

each category ψ has a corresponding smooth fictitious play $\sigma^i(\phi_{t-1}^{-i}(\psi))$

suppose we choose category ψ with probability $\lambda(\psi)$, then overall play is

$$pr(s^i) = \sum_{\psi} \lambda(\psi) \sigma^i(\phi_{t-1}^{-i}(\psi))[s^i]$$

but categories correspond to own strategies: fixed point property:

$$\lambda(s^i) = pr(s^i)$$

$$\lambda(s^i) = \sum_{\psi} \lambda(\psi) \sigma^i(\phi_{t-1}^{-i}(\psi))[s^i]$$

unique fixed point, solvable by linear algebra

Interpretation of Calibration


weather forecasting example: calibrated beliefs, versus calibrated actions

consequence of universal calibration: global convergence to the set of correlated equilibria


Shapley Example

	A	M	B
A	0,0	0,1	1,0
M	1,0	0,0	0,1
B	0,1	1,0	0,0

smooth fictitious play (time in logs)



condition on opponents last period play (time in logs)



Discounted Learning

But what about the short-run?

In the short-run we are guessing...

A learning procedure $\hat{\rho}$ is ε -as good as a procedure ρ if for all sequences of discount factors $\{\beta_t\}$ and all histories h_t^i

$$\sum_{t=1}^{\infty} \beta_t u(\rho(h_{t-1}), s_t^{-i}) \leq \sum_{t=1}^{\infty} \beta_t u(\hat{\rho}(h_{t-1}), s_t^{-i}) + \varepsilon$$

Proposition: For any learning procedure ρ and any ε there exists a categorical smooth fictitious play $\hat{\rho}$ that is ε -as good as ρ

exploits the fact that the time average result must be true for at every
time