Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Adaptive Play by Idiosyncratic Agents (2004 GEB)
David P. Myatt and Christopher Wallace

Motivation

- In the standard approach covered in the first part of the course, players within each population had identical payoff functions and strategy sets.
- This paper examines the consequences of introducing heterogeneity in payoffs.
- Random (idiosyncratic) payoffs introduce noise into the revision process which plays a role similar to that of errors in the Young (1998) framework.

Pairwise Interactions

- Agents are matched to play a game of the following form:

	1	2
1	a, a	0, 0
2	0, 0	d, d

- Assume $a > d > 0$, so that there are two pure-strategy Nash equilibria, the coordination equilibria $(1, 1)$ and $(2, 2)$, with the first being risk dominant.

- In the mixed Nash equilibrium $(x^*, 1 - x^*)$, $x^* = \frac{d}{a+d}$ which is less than $\frac{1}{2}$ since $a > d$ by assumption.

Payoff Heterogeneity

- Instead of being matched to play the exact game above, each player's payoffs are perturbed:

$$\begin{array}{l} \tilde{a} = a + \sigma\varepsilon_a \\ \tilde{d} = d + \sigma\varepsilon_d \end{array} \quad \text{where}$$

$$\begin{bmatrix} \varepsilon_a \\ \varepsilon_d \end{bmatrix} \sim N\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \xi_a^2 & \rho\xi_a\xi_d \\ \rho\xi_a\xi_d & \xi_d^2 \end{bmatrix} \right).$$

- This is a specific kind of random utility model: the normal distribution is convenient especially due to its unbounded support which allows either strategy to be dominant with positive probability.

- The common scaling factor σ will be allowed to vanish for the limiting results.

Revision Protocol

- There is a finite population of n players.
- Each period, every individual plays a fixed pure strategy against a randomly selected opponent from the $n-1$ remaining players.
- The state of play $z \in Z = \{0, \ldots, n\}$ is the number of players using strategy 1.

Revision Protocol

- *Birth-death process*: At the end of each period a randomly selected player is replaced and the entrant equipped with new (perturbed payoffs) $\tilde{a}$ and $\tilde{d}$.
- The entrant observes the strategy distribution among the remaining $n-1$ players and selects a (myopic) best response.
- This defines a (time-homogeneous) Markov chain on the state space Z.

Best Responses

- Beginning in state z, and following the exit of an incumbent, either $i = z$ or $i = z - 1$ of the remaining incumbents will be using strategy 1.

- Let $x = i/(n-1)$ denote the fraction of agents using strategy 1.

- The entrant chooses strategy 1 whenever $x(a + \sigma\varepsilon_a) > (1-x)(d + \sigma\varepsilon_d)$.

- Rearranging:

$$(1-x)\varepsilon_d + x\varepsilon_a < [xa - (1-x)d]/\sigma.$$

Best Responses

- The LHS is a normally distributed r.v. with mean 0 and variance $x^2\xi_a^2 + (1-x)^2\xi_d^2 - 2x(1-x)\rho\xi_a\xi_d$.

- Hence the entrant chooses pure strategy 1 with probability:

$$Pr[1|x] = \Phi\left(\frac{xa-(1-x)d}{\sigma\sqrt{x^2\xi_a^2 + (1-x)^2\xi_d^2 - 2x(1-x)\rho\xi_a\xi_d}}\right).$$

where Φ denotes the cdf of the standard normal distribution.

Basins of Attraction

- The **basins of attraction** of strategies 1 and 2 are:

$$Z_1 = \{\lceil (n-1)x^* + 1\rceil, \ldots, n\} \quad \textit{and} \quad Z_2 = \{0, \ldots, \lfloor (n-1)x^* \rfloor\}.$$

- The **basin depth** faced by an entrant is $\kappa(x)^2$:

$$\kappa(x) = \frac{xa - (1-x)d}{\sigma\sqrt{x^2\xi_a^2 + (1-x)^2\xi_d^2 - 2x(1-x)\rho\xi_a\xi_d}}.$$

- Define $\kappa_i = \kappa\left(\frac{i}{n-1}\right)$.

Basins of Attraction

- If $z \in Z_1$, the "flow of play" is toward strategy 1: $Pr[1|x] > 1/2$. If $z \in Z_2$, the "flow of play" is toward strategy 2.

- $z = \lceil (n-1)x^* \rceil$ belongs to neither basin of attraction. The most likely choice of the entrant depends on the identity of the exiting player.

Basin Depths

- While the basins of attraction describe the flow of play, the basin depth measures the difficulty of moving against the flow.

- To see this, consider a state $z \in Z_1$. An entrant is most likely to play strategy 2 when a player using strategy 1 exits.

- In this case, $Pr[2|x] = 1 - \Phi(\kappa_i/\sigma)$ where $i = z - 1$. The larger is κ_i^2, the lower is $Pr[2|x]$ (recall the mean is zero).

Transition Probabilities

- Due to step-by-step revisions, all transitions are local: $p_{z,z'} = 0$ for $|z - z'| > 1$.

- For states $z < n$, the probability of a step up is:

$$p_{z,z+1} = \frac{n - z}{n} \times \Phi\left(\frac{\kappa_z}{\sigma}\right).$$

- The other transitions are:

$$p_{z,z-1} = \frac{z}{n} \times \left[1 - \Phi\left(\frac{\kappa_{z-1}}{\sigma}\right)\right].$$

$$p_{z,z} = \frac{n - z}{n} \times \left[1 - \Phi\left(\frac{\kappa_z}{\sigma}\right)\right] + \frac{z}{n} \times \Phi\left(\frac{\kappa_{z-1}}{\sigma}\right).$$

Asymptotic Behavior

- The Markov chain is irreducible and aperiodic, hence there is a unique stationary distribution (or ergodic distribution) which describes the long-run behavior of the process independently of initial conditions.

- We know that two-strategy games under arbitrary revision protocols generate reversible Markov processes, which permit easy computation of the stationary distribution.

- Indeed, we can employ Theorem 11.5 (from Lecture 11) to derive the stationary distribution, which Myatt and Wallace denote by π.

- The local maxima of the stationary distribution coincide with the Bayesian Nash equilibria of the underlying game.

Stationary Distribution & BNE

- To see why, suppose $x < \Phi(\kappa(x)/\sigma)$ of incumbent players are using strategy 1.
- A strategy 1 player is less likely to exit [with prob. x] than enter [with approx. prob. $\Phi(\kappa(x)/\sigma)$]. In expectation, the number of strategy 1 players is growing.
- The opposite occurs when $x > \Phi(\kappa(x)/\sigma)$.
- So the process moves toward stable fixed points of $\Phi(\kappa(x)/\sigma)$, which are Bayesian Nash equilibria.

Stationary Distribution & BNE

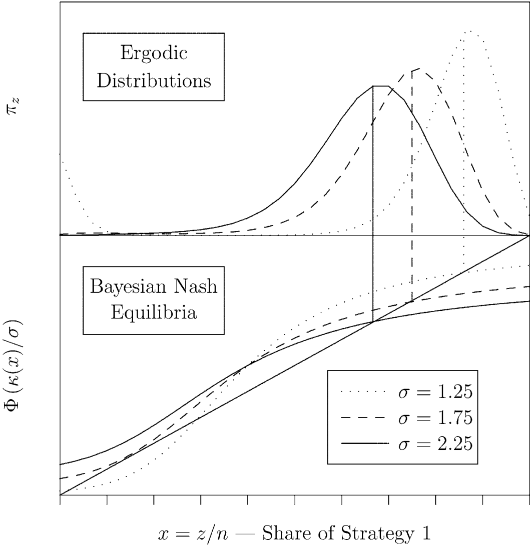

Fig. 1. Parameters are $a = 3$, $d = 2$, $\xi_a = \xi_d = 1$, $\rho = 0$ and $n = 30$.

Equilibrium Selection: large n limit

Proposition 1. As $n \to \infty$ all the weight in the stationary distribution focusses on a single Bayesian Nash equilibrium.

Equilibrium Selection: small σ limit

- Bergin and Lipman (1996, Ecta) demonstrate that for any recurrence class E one can choose state-dependent mutations in such a way that E is stochastically stable.

- Blume (1999, SFI working paper) shows that the risk dominant equilibrium is selected if a mild "skew symmetry" condition holds: the probability of a mutation depends only on the absolute difference between the payoffs to the two strategies.

- When $\xi_a = \xi_d$, the noise process in the current paper is skew symmetric:

  Hence as $\sigma \to 0$, the risk-dominant equilibrium is selected in this case.

Equilibrium Selection: small σ limit

- For $\xi_a \neq \xi_d$, we need to define:

  The **basin volumes** are $B_1 = \sum_{z \in Z_1} \kappa_{z-1}^2$ and $B_2 = \sum_{z \in Z_2} \kappa_z^2$.

  Proposition 2. If $B_1 > B_2$ then $\lim_{\sigma \to 0} \pi_n = 1$ and if $B_2 > B_1$ then $\lim_{\sigma \to 0} \pi_0 = 1$.

  Corollary 1. For $n = 2$, strategy 1 is selected ($\lim_{\sigma \to 0} \pi_n = 1$) whenever $a/\xi_a > d/\xi_d$.