Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 17.
Fast Convergence

The Very Long Run?

Based on notes by H. Peyton Young

Kreindler & Young (2013 GEB, 2014 PNAS)

- In the long run, the stochastic dynamic spends almost all the time in the stochastically stable states.
- However, the expected waiting time to reach a stochastically stable state grows exponentially as the error rate ε becomes arbitrarily small.
- Is SS only relevant in the very long run?

Intermediate Error Rates

In this lecture, we see that for intermediate values of ε there is:

- Sharp selection.
- Fast convergence.

These results hold when agents respond to:

(a) the distribution of actions in the whole population,

(b) random samples from the population,

(c) their neighbors in a network.

Model

Population size N, large but finite.

2×2 symmetric pure coordination game:

	A	B
A	$1+\alpha$ $1+\alpha$	0 0
B	0 0	1 1

- (B, B) is the status quo.
- (A, A) is the innovation, with $\alpha > 0$ the payoff gain to the deviation.

General Coordination Game

Can generalize to any 2×2 symmetric coordination game:

	A	B
A	a, a	c, d
B	d, c	b, b

with α redefined as the normalized potential difference:

$$\alpha = \frac{(a-d)-(b-c)}{b-c}.$$

Strategy Revisions

- Time is discrete.
- Each period lasts $\tau = \frac{1}{N}$ units of time.
- Each period, one randomly chosen agent revises:

Step 1. Gathers information on current play,

Step. 2 Chooses a (myopic) noisy best response.

Logit Learning

Step 1. Agent forms estimate x of adoption rate of A.

(a) **Full information:** agent knows the current proportion of adopters in the population.

(b) **Partial information:** agent randomly samples d other players.

Step 2. Noisy best response given by logit function:

$$\Pr(\text{Choose } A|x) = \frac{e^{\beta(1+\alpha)x}}{e^{\beta(1+\alpha)x} + e^{\beta(1-x)}}.$$

Error rate (at zero adoption) is $\varepsilon = \frac{1}{1+e^{\beta}}$.

Logit Learning: $\alpha = 0.5, \beta = 3.5$

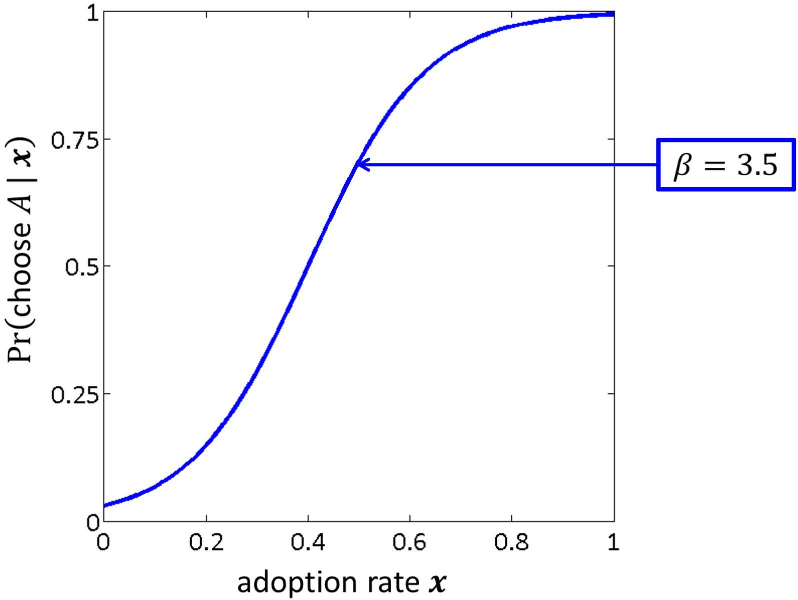

Logit Learning: $\alpha = 0.5$, various β

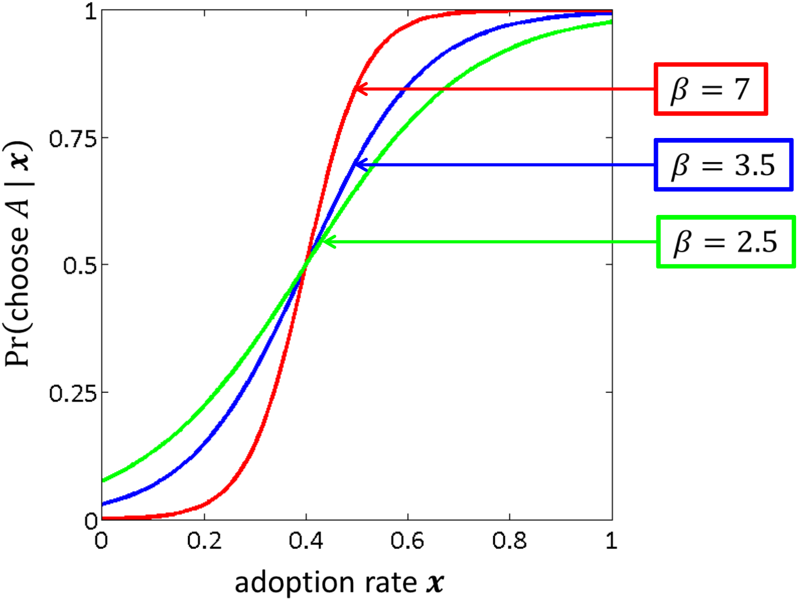

Convergence Times

- Study process $\Gamma_N(\alpha, \beta)$ starting from all- B state (status quo).
- State variable: adoption rate $x(t)$ of innovation A.
- Define **waiting time** to adoption level $p < 1$:

$$T_N(\alpha, \beta, p) = \mathsf{min}\{t : x(t) \geq p\}.$$

Sample Adoption Path

$\varepsilon = 5\%, \alpha = 100\%, p = 99\%, N = 1000$

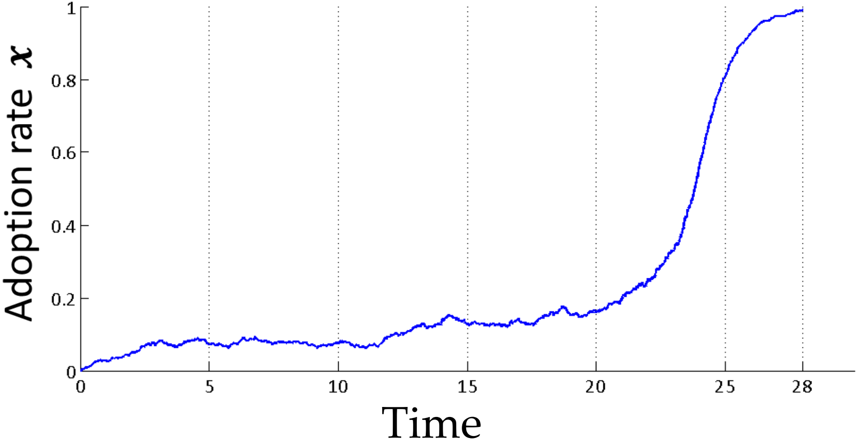

Fast Convergence - definitions

Definition 1. The family $\Gamma_N(\alpha, \beta)$ exhibits *fast convergence* if the expected waiting time until a majority of agents play A is bounded independently of N, or

$$E\,T_N\left(\alpha, \beta, \tfrac{1}{2}\right) < S(\alpha, \beta) \text{ for all } N.$$

Definition 2. The family $\Gamma_N(\alpha, \beta)$ exhibits *fast convergence* to p if the expected waiting time to adoption level p is bounded independently of N, or

$$E\,T_N\left(\alpha, \beta, p\right) < S(\alpha, \beta, p) \text{ for all } N.$$

Fast Convergence - result

Theorem 16.1. Let

$$
\begin{aligned}
h(\beta) &= \frac{e^{\beta-1}+4-e}{\beta}-2 \ \text{ for } \beta > 2 \\
h(\beta) &= 0 \ \text{ for } 0 < \beta \leq 2
\end{aligned}
$$

If $\alpha > h(\beta)$, then $\Gamma_N(\alpha, \beta)$ exhibits *fast convergence*.

Error Rate

Recall that the error rate ε is the probability of choosing A when you expect your opponent to choose B:

$$\varepsilon = \frac{e^0}{e^0 + e^\beta}.$$

- $\beta = 2$ means $\varepsilon \approx 12\%$.
- $\beta = 3$ means $\varepsilon \approx 5\%$.

Threshold for Fast Convergence

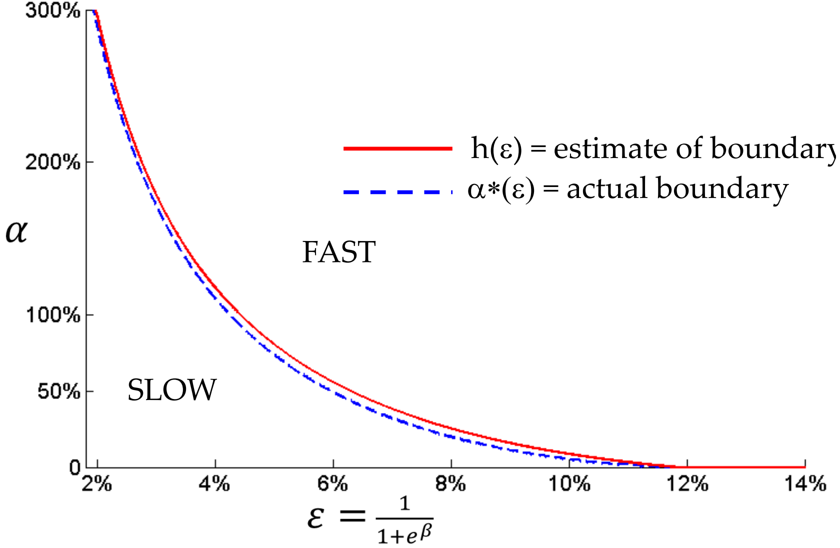

Proof Strategy

$$\Pr(\text{Choose } A|x) = f(x;\alpha,\beta) = \frac{e^{\beta(1+\alpha)x}}{e^{\beta(1+\alpha)x} + e^{\beta(1-x)}}.$$

Recall that the continuous-time mean logit dynamic is the differential equation:

$$\dot{x} = f(x;\alpha,\beta) - x \ \text{ with } x(0) = 0.$$

The logit equilibria are the fixed points:

$$f(x^*;\alpha,\beta) = x^*.$$

The key is to find combinations of α, β such that the lowest fixed point is greater than $1/2$.

Proof Strategy

The status quo equilibrium disappears when α is big enough.

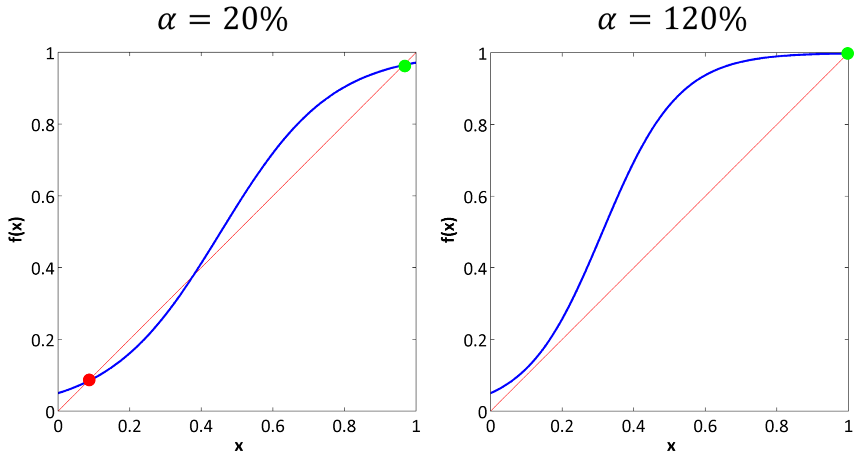

Proof Strategy

Let x_0 be the tangency point. Then:

$$f'(x_0; \alpha, \beta) = 1 \tag{1}$$

and

$$f(x_0; \alpha, \beta) = x_0. \tag{2}$$

Note that:

$$f'(x_0; \alpha, \beta) = \beta(2 + \alpha) f(x_0) \left(1 - f(x_0)\right) = 1. \tag{3}$$

Combining (1) and (3) yields

$$x_0(1 - x_0) = \frac{1}{\beta(2 + \alpha)}.$$

Proof Strategy

If x_0 is small, x_0^2 is very small.

Hence

$$\beta(2+\alpha)x_0 \approx 1 \tag{4}$$

$$\begin{aligned} f(x_0) &= \frac{e^{\beta(1+\alpha)x_0}}{e^{\beta(1+\alpha)x_0} + e^{\beta(1-x_0)}} \\ &= \frac{1}{1+e^{\beta-\beta(2+\alpha)x_0}} \approx \frac{1}{1+e^{\beta-1}} \\ &\approx x_0 \approx \frac{1}{\beta(2+\alpha)}. \end{aligned}$$

Therefore,

$$\beta(2+\alpha) \approx e^{\beta-1} + 1.$$

This defines the approximate combinations of α and β required to lift $f(x)$ off the 45-degree line.

In fact, we need $\beta(2+\alpha) > e^{\beta-1} + 4 - e$.

Average Waiting Times

$\varepsilon = 5\%, p = 99\%$

	$N = 100$	$N = 1000$	$N = 10,000$
$\alpha = 70\%$	33	101	$> 8,000$
$\alpha = 80\%$	25	36	35

$\varepsilon = 10\%, p = 50\%$ *(top row)*, $p = 90\%$ *(bottom row)*

	$N = 100$	$N = 1000$	$N = 10,000$
$\alpha = 4\%$	38	190	$> 8,000$
$\alpha = 25\%$	19	20	21

Partial Information

Analogous results hold when agents draw random *samples* from the population:

- Sample size $d < \infty$ fixed independently of N.
- Updating function becomes:

$$f_d(x;\alpha,\beta) = \sum_{k=0}^{d} \binom{d}{k} x^k (1-x)^{d-k} f\left(\tfrac{k}{d};\alpha,\beta\right).$$

Theorem 16.2. Assume $d \geq 3$. If $\alpha > \mathsf{min}\{h(\beta), d-2\}$, the process exhibits fast convergence.

Payoff Heterogeneity

Now suppose agents choose exact best responses, but payoffs are perturbed by small shocks:

Δ is the payoff gain from choosing A:

$$\Delta = (1+\alpha)x - (1-x).$$

Let ϵ_A and ϵ_B be i.i.d. (idiosyncratic) payoffs from playing A and B. Then:

$$\Pr(\text{Choose } A|x) = \Pr(\epsilon_A - \epsilon_B + \Delta > 0).$$

- If ϵ_A and ϵ_B are extreme value distributed, this is logit choice.
- For comparison, let us consider normally distributed payoff shocks, $\epsilon_A, \epsilon_B \sim N(0, 1/\beta)$.

Fast Convergence

Response function (left) and fast diffusion threshold (right).

Normally distributed payoff shocks (gray dots), and extreme value (black line) for $d = 15$.

Conclusion

Evolutionary selection occurs within realistic time frames for plausible levels of error and/or payoff heterogeneity.

For the role of networks in fast convergence, see:

- Morris (2000 ReStud),
- Kreindler & Young (2014 PNAS),
- Arieli & Young (2016 Ecta).