Evolution & Learning in Games
Econ 243B

Jean-Paul Carvalho

Lecture 12.
Stochastic Evolution and Stationary Distributions

Outline

- So far, we have focussed exclusively on the behavior of the mean dynamic, which we have treated as a (deterministic) approximation to the underlying stochastic process.
- The informal justification was that when the population is large, idiosyncratic noise will be averaged away.
- In this lecture, we shall formalize this notion for a large but finite population, in the finite-horizon case.

Outline

- When we want to analyze evolution over infinite horizons, we need a different set of concepts and tools, which we shall begin to introduce in the second part of the lecture.
- Eventually, we shall see that we can make sharp predictions about the long-run behavior of the stochastic dynamic, while characterizing its (approximate) medium-run behavior by the deterministic mean dynamic.
- Our stochastic analysis will focus (primarily) on the single-population case.

The Stochastic Evolutionary Process

- Let the population be *large but finite*, with N members.
- The set of feasible social states is then a discrete grid embedded in X:

$$\mathcal{X}^N = X \bigcap \tfrac{1}{N}\mathbb{Z}^n = \{x \in X : Nx \in \mathbb{Z}^n\}.$$

The Stochastic Evolutionary Process

- Once again:
  - A population game is denoted by $F : X \to \mathbb{R}$.
  - The strategy space is $S = \{1, \ldots n\}$,
  - Choices are described by the revision protocol $\rho = \mathbb{R}^n \times X \to \mathbb{R}_+^{n \times n}$.
- Each agent is also equipped with a rate 1 Poisson clock, where a ring signals the arrival of a revision opportunity for the clock's owner.

Markov Property

The following independence assumptions are made:

- Different clocks ring independently of each other,
- Strategy choices are made independently of the clocks' rings,
- Choices are based exclusively on the current state of the process x and current payoff vector π; beyond that the history of play does not matter.

Together these assumptions mean that the current state summarizes all the information on the history of play required to characterize the future motion of the process.

That is, the stochastic evolutionary process $\{X_t^N\}$ is a continuous-time **Markov process** on the finite state space $\mathcal{X}^N$.

Transition Probabilities

- Explicitly, the process is fully characterized by its transition probabilities $\{P^N_{xy}\}_{x,y \in \mathcal{X}^N}$.
- When an agent playing $i \in S$ receives a revision opportunity, he switches to $j \neq i$ with probability ρ_{ij}.
- The probability that the next revision involves a switch from strategy i to j is then $x_i \rho_{ij}$.
- The transition involves one less agent playing i and one more agent playing j, i.e. the state is shifted by $\frac{1}{N}(e_j - e_i)$.
- $P^N_{xy}(t)$ is the probability of transiting from state x to y in exactly one revision.
- If $P^N_{xy}(t)$ is independent of t, then the Markov process is *time homogenous*, and we write P^N_{xy}. This is the case we will be dealing with.

Transition Matrix

- P^N is a $|\mathcal{X}^N| \times |\mathcal{X}^N|$ matrix, called the *transition matrix*, whose elements are the transition probabilities.
- For all x, $\sum_{y \in \mathcal{X}^N} P^N_{xy} = 1$.
- This means that P^N is a *row stochastic matrix*, i.e. its row elements sum to one.
- For example, consider the following transition matrix for a two-state Markov process:

$$P^N = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{4} & \frac{3}{4} \end{pmatrix}.$$

Representation

Another way to describe this process is via a graph:

Markov Process

Observation 11.1. A population game F, a revision protocol ρ, a revision opportunity arrival rate of one, and a population size N define a Markov process $\{X_t^N\}$ on the state space $\mathcal{X}^N$. This process is described by some initial state $X_0^N = x_0^N$, the jump rates $\lambda_x^N = N$ and the transition probabilities:

$$P_{x,x+z}^N = \begin{cases} x_i \rho_{ij}\big(F(x), x\big) & \text{if } z = \frac{1}{N}(e_j - e_i),\, i, j \in S,\, i \neq j \\ 1 - \sum_{i \in S} \sum_{j \neq i} x_i \rho_{ij}\big(F(x), x\big) & \text{if } z = \mathbf{0} \\ 0 & \text{otherwise} \end{cases}$$

Deterministic Approximation

Theorem 11.1. Let $\{\{X_t^N\}\}_{N=N_0}^{\infty}$ be a sequence of continuous-time Markov processes on the state space $\mathcal{X}^N$.

Suppose that V is Lipschitz continuous. Let the initial conditions $X_0^N = x_0^N$ converge to state $x_0 \in X$ as $N \to \infty$, and let $\{x_t\}_{t\geq 0}$ be the solution to the mean dynamic starting from x_0.

Then for all $T < \infty$ and $\varepsilon > 0$,

$$\lim_{N\to\infty} \mathbb{P}\left(\sup_{t\in[0,T]} |X_t^N - x_t| < \varepsilon \right) = 1.$$

Conditions

- Lipschitz continuity guarantees unique solutions to the mean dynamic, but does not apply to the best response dynamic.
- Still, deterministic approximation results for the best response dynamic are available.
- More importantly, this is a finite-horizon result.

Conditions

- To see why the result does not extend to an infinite horizon, consider the logit dynamic in a potential game.
  - We have seen that this mean dynamic converges to a perturbed equilibrium from any initial state.
  - Over a finite horizon, the underlying Markov process closely tracks the mean dynamic, when the population is large.
  - When the population is large but finite, however, there is always a positive probability path from each state to every other state, so that every state in $\mathcal{X}^N$ is visited infinitely often with probability one.

Long-Run Analysis

- Over finite-horizons, the focus is on the mean dynamic.
- In the long run, however, the object of interest is the **stationary distribution** μ of the process $\{X_t\}$ (we are now dropping the N superscript).
- The stationary distribution tells us the frequency distribution of visits to each state as $t \to \infty$, i.e. almost surely the process spends proportion $\mu(x)$ of the time in state x.
- Before analyzing stationary distributions, let us introduce some definitions.

Communication

- State y is **accessible** from x if there exists a *positive probability path* from x to y, i.e. a sequence of states beginning in x and ending in y in which each one step transition between states has positive probability under P.
- States x and y **communicate** if they are each accessible from the other.
- A set of states E is **closed** if the process cannot leave it, i.e. for all $x \in E$ and $y \notin E$, $P_{xy} = 0$.
- An **absorbing state** is a singleton closed set.
- Every state in $\mathcal{X}$ is either transient or recurrent; and a state is **recurrent** if and only if it is a member of a *closed communication class*.

Recurrence

Theorem 11.2. Let $\{X_t\}$ be a Markov Process on a finite set $\mathcal{X}$.

Then starting from any state x_0, the frequency distribution of visits to states converges to a stationary distribution μ, which solves $\mu P = \mu$. In such a vector, $\mu(x) = 0$ for all transient states.

- This is why closed communication classes are commonly called recurrence classes.

Stationary Distributions

- FACT: every non-negative row stochastic matrix has at least one left eigenvector with eigenvalue one, i.e. there exists a μ such that $\mu P = \mu$.

- In other words there exists at least one stationary distribution.

Theorem 11.3.

(i) If a stationary distribution μ is unique, it is the long run frequency distribution independent of the initial state.

(ii) If there are multiple stationary distributions, then the long run frequency distribution is among these, and can be any one of them.

Stationary Distributions

- In our previous example, $\mu = (\frac{1}{3}, \frac{2}{3})$ (check by showing that μ solves $\mu P = \mu$).

- Consider the following Markov process:

- The solutions to the stationarity equation are:

$$\mu_1 = (\tfrac{1}{3}, \tfrac{2}{3}, 0, 0), \mu_2 = (0, 0, 0, 1),$$

  or any convex combination of μ_1 and μ_2.

Irreduciblity

- A Markov Process is **irreducible** if there is a positive probability path from each state to every other, i.e. if all states in $\mathcal{X}$ communicate, or equivalently if $\mathcal{X}$ forms a single recurrent class.

Theorem 11.4. If the Markov process $\{X_t\}$ is irreducible then it has a unique stationary distribution, and this stationary distribution is independent of the initial state.

In this case, we say that the process $\{X_t\}$ is **ergodic**, its long-run behavior does not depend on initial conditions.

k-Step Ahead Probabilities

- We may not only want to know the proportion of time the process spends in each state (given by μ), but also the probability of being in a given state at some future point in time.
- Recall that $\mathbb{P}(X_1 = y | X_0 = x) = P_{xy}$ is a one step transition probability.
- Two step transition probabilities are computed by multiplying P by itself:

$$\begin{aligned}
\mathbb{P}(X_2 = y | X_0 = x) &= \sum_{z \in \mathcal{X}} \mathbb{P}(X_2 = y, X_1 = z | X_0 = x) \\
&= \sum_{z \in \mathcal{X}} \mathbb{P}(X_1 = z | X_0 = x) \mathbb{P}(X_2 = y | X_1 = z, X_0 = x) \\
&= \sum_{z \in \mathcal{X}} P_{xz} P_{zy} \\
&= \left(P^2\right)_{xy}.
\end{aligned}$$

Aperiodicity

- By induction, the t-step transition probabilities are given by the entries of the tth power of the transition matrix:

$$\mathbb{P}(X_t = y | X_0 = x) = \left(P^t\right)_{xy}.$$

- Let $\mathcal{T}_x$ be the set of all positive integers T such that there is a positive probability of moving from x to x in exactly T periods.

- The process is **aperiodic** if for every $x \in \mathcal{X}$, the greatest common denominator of $\mathcal{T}_x$ is 1.

This holds whenever the probability of remaining in each state is positive, i.e. $P_{xx} > 0$ for all $x \in \mathcal{X}$.

Aperiodicity

- If $\{X_t\}$ is *irreducible* and *aperiodic*, then with probability one:

$$\text{for all } x_0, x \in \mathcal{X} \qquad \lim_{t \to \infty} \left(P^t\right)_{x_0 x} = \mu(x).$$

- Therefore, from any initial state x_0 both the proportion of time the process spends in each state up through time t and the probability of being in each state at time t converge to the stationary distribution μ.

- Hence the stationary distribution provides a lot of information about the long run behavior of the process.

Full Support Revision Protocols

- Irreducibility and aperiodicity are desirable properties of a Markov process.
- They are both generated by full support revision protocols, i.e. revision protocols in which all strategies are chosen with positive probability.
- Let us consider two examples which are extensions of best response protocols.

Full Support Revision Protocols

- **Best response with mutations:**
  - A revising agent switches to his current best response with probability $1 - \varepsilon$, and chooses a strategy uniformly (mutates) with probability $\varepsilon > 0$.
  - Let us refer to this protocol as $BRM(\varepsilon)$.
  - In case of best response ties, it is often assumed that a non-mutating agent sticks with her current strategy if it is a best response; otherwise she chooses at random among from the set of best responses.

Full Support Revision Protocols

- **Logit Choice:**

$$\rho_{ij}(\pi) = \frac{exp(\eta^{-1}\pi_j)}{\sum_{k\in S} exp(\eta^{-1}\pi_k)}.$$

- We can define $\varepsilon^{-1} = \exp(\eta^{-1})$, where ε is an increasing function of η. As $\eta \to 0$, $\varepsilon \to 0$.

- In this way, we can rewrite the revision protocol as follows:

$$\rho_{ij}(\pi) = \frac{\varepsilon^{-\pi_j}}{\sum_{k\in S} \varepsilon^{-\pi_k}}.$$

Stationary Distributions

- There are two problems:
  - It may not be possible to compute the stationary distribution explicitly,
  - Even if it is possible to do so, the stationary distribution may spread weight widely over the state space.

Analyzing Large-Dimensional Markov Processes

- In this lecture:
  - We shall study a class of **reversible** Markov processes whose stationary distributions are easy to compute;
- In the next lecture, we shall:
  - Introduce the concept of **stochastic stability**, which can drastically reduce the number of states which attract positive weight in the stationary distribution.

Reversible Markov Processes

- Reversible Markov processes permit easy computation even if the state space $\mathcal{X}$, and hence the $|\mathcal{X}| \times |\mathcal{X}|$ transition matrix P, is large.
- A process $\{X_t\}$ is **reversible** if it admits a *reversible distribution*, i.e. a probability distribution μ on $\mathcal{X}$ that satisfies the following *detailed balance conditions*:

$$\mu_x P_{xy} = \mu_y P_{yx} \quad \text{for all } x, y \in X. \tag{1}$$

- Such a process is reversible in the sense that it looks the same whether time is run forward or backward.

Reversible Markov Processes

- Recall that a stationary distribution μ satisfies:

$$\sum_{x \in \mathcal{X}} \mu_x P_{xy} = \mu_y \quad \text{for all } y \in X. \tag{2}$$

- Summing (1) over x we get:

$$\begin{aligned} \sum_{x \in \mathcal{X}} \mu_x P_{xy} &= \sum_{x \in \mathcal{X}} \mu_y P_{yx} \\ &= \mu_y \sum_{x \in \mathcal{X}} P_{yx} \\ &= \mu_y. \end{aligned} \tag{3}$$

- Therefore, a reversible distribution is also a stationary distribution.

Reversible Markov Processes

- There are two contexts in which the stochastic evolutionary process $\{X_t\}$ is known to be reversible:

  1. two-strategy games (under arbitrary revision protocols),
  2. potential games under exponential protocols.

- We shall now study the first and leave the second to later.

Two-Strategy Games

- Let $F : X \to \mathbb{R}^2$ be a two strategy game, with strategy set $\{0, 1\}$, full support revision protocol $\rho : \mathbb{R}^2 \times X \to \mathbb{R}^{2\times 2}$, and finite population size N.
- This defines an irreducible and aperiodic Markov Process $\{X_t\}$ on the state space $\mathcal{X}^N$.
- For this class of games, let $x \equiv x_1$. The state of the process is fully described by x.
- Therefore, the state space is $\mathcal{X}^N = \{0, \frac{1}{N}, \ldots, 1\}$, a uniformly spaced grid embedded in the unit interval.

Birth and Death Processes

- Because revision opportunities arrive independently in continuous time, agents switch strategies sequentially.
- This means that transitions are always between adjacent states.
- If in addition the state space is linearly ordered (which it is in a two-strategy game), then we refer to the Markov process as a birth and death process.
- We shall now show that a stationary distribution for such processes can be calculated in a straightforward way.

Birth and Death Processes

- In a birth and death process, there are vectors $p, q \in \mathbb{R}^{|\mathcal{X}|}$ with $p_1 = q_0 = 0$ such that the transition matrix takes the following form:

$$P_{xy} \equiv \begin{cases} p_x & \text{if } y = x + \frac{1}{N}, \\ q_x & \text{if } y = x - \frac{1}{N}, \\ 1 - p_x - q_x & \text{if } y = x, \\ 0 & \text{otherwise} \end{cases}$$

- The process is *irreducible* if $p_x > 0$ for all $x < 1$ and $q_x > 0$ for all $x > 0$, which is what we shall assume.

Birth and Death Processes

- Because of the "local" structure of transitions, the reversibility condition reduces to:

$$\mu_x q_x = \mu_{x-1/N} p_{x-1/N}$$

  for all $x > 0$.
- Applying the formula inductively, we have:

$$\mu_x q_x q_{x-1/N} \cdots q_{1/N} = \mu_0 p_0 p_{1/N} p_{2/N} \cdots p_{x-1/N}.$$

- That is, the process running 'down' from state x to zero should look like the process running 'up' from zero to state x.

Stationary Distribution

- Rearranging, we see that the stationary distribution satisfies:

$$\frac{\mu_x}{\mu_0} = \prod_{j=1}^{Nx} \frac{p_{(j-1)/N}}{q_{j/N}} \qquad \text{for all } x \in \{\tfrac{1}{N}, \dots, 1\}. \tag{4}$$

- For a full support revision protocol ρ, the upward and downward probabilities are given by:

$$\begin{aligned} p_x &= (1-x)\rho_{01}\big(F(x), x\big) \\ q_x &= x\rho_{10}\big(F(x), x\big). \end{aligned} \tag{5}$$

Stationary Distribution

- Substituting the expressions in (5) into (4) yields:

$$\frac{\mu_x}{\mu_0} = \prod_{j=1}^{Nx} \frac{p_{(j-1)/N}}{q_{j/N}} = \prod_{j=1}^{Nx} \frac{1 - \frac{j-1}{N}}{\frac{j}{N}} . \frac{\rho_{01}\left(F\left(\frac{j-1}{N}\right), \frac{j-1}{N}\right)}{\rho_{10}\left(F\left(\frac{j}{N}\right), \frac{j}{N}\right)} \tag{6}$$

for all $x \in \{\frac{1}{N}, \ldots, 1\}$.

Stationary Distribution

Simplifying, we have the following result:

Theorem 11.5. Suppose that a population of N agents plays the two-strategy game F using the full support revision protocol ρ. Then the stationary distribution for the evolutionary process $\{X_t^N\}$ on $\mathcal{X}^N$ is given by:

$$\frac{\mu_x}{\mu_0} = \prod_{j=1}^{Nx} \frac{N-j+1}{j} \cdot \frac{\rho_{01}\left(F(\frac{j-1}{N}), \frac{j-1}{N}\right)}{\rho_{10}\left(F(\frac{j}{N}), \frac{j}{N}\right)} \qquad \text{for } x \in \{\tfrac{1}{N}, \ldots, 1\}, \quad (7)$$

with μ_0 determined by the requirement that $\sum_{x \in \mathcal{X}} \mu_x = 1$.