Evolution & Learning in Games

Econ 243B

Jean-Paul Carvalho

Lecture 11.

Stochastic Stability

Analyzing Large-Dimensional Markov Processes

- We have seen that stationary distributions for reversible Markov processes can be computed.
- When the noise level ε is positive and the population size N is finite, individual behavior gives rise to an irreducible (and aperiodic) Markov process on a finite state space.
- We know there exists a unique stationary distribution $\mu^{\varepsilon,N}$ of such a process.
- However, this stationary distribution puts positive weight on every state.
- As we take the limit as $\varepsilon \to 0$ (the small noise limit) or as $N \to \infty$ (the large population limit) or both, the stationary distribution concentrates weight on a small set of states (often a single state).

Stochastic Stability

- There are different ways of defining **stochastic stability** depending on which limits are taken and in what order.
- Results can differ depending on the definition used.
- We shall focus on the small noise limit:

Definition: *Stochastic Stability* (Foster and Young 1990). A state x is stochastically stable if $\lim_{\varepsilon \to 0} \mu^{\varepsilon,N}(x) > 0$.

- In other words, for ε small, the process spends virtually all the time (as $t \to \infty$) in the stochastically stable set of states.

Risk dominance

- Consider the following general symmetric normal-form game:

$$A = \begin{pmatrix} a & c \\ b & d \end{pmatrix}.$$

- Let x equal the weight on strategy 1. The mixed equilibrium is $x^* = \frac{a-b}{(a-b)+(d-c)}$.
- $x = 1$ *risk dominates* $x = 0$ if the product of the deviation losses is at least as high for $x = 1$, and *vice versa* (Harsanyi and Selten 1988).
- In a symmetric game this means: $(d - c) \geq (a - b)$, which implies that $x^* \leq \frac{1}{2}$.
- If the inequalities are strict then $x = 1$ strictly risk dominates state $x = 0$.

Two-Strategy Coordination Games

- Consider the following symmetric normal-form *pure* coordination game:

$$A = \begin{pmatrix} 2 & 0 \\ 0 & 3 \end{pmatrix}.$$

- Note that:
  - Strategy 1 is a best response if and only if $x \geq x^* = \frac{2}{5}$, where x^* is the mixed Nash equilibrium.
  - Strategy 0 is a best response if and only if $x \leq \frac{2}{5}$.
  - As $x^* < \frac{1}{2}$, $x = 1$ is strictly risk dominant (here it is both the payoff dominant and risk dominant equilibrium).

Two-Strategy Coordination Games

Theorem 10.1. In two-strategy coordination games, for N sufficiently large, the $BRM(\varepsilon)$ revision protocol selects the state in which everyone plays the risk-dominant Nash equilibrium as the unique stochastically stable state.

Proof

- When individuals are randomly matched to play a symmetric coordination game, the only possible recurrence classes of the unperturbed $BRM(\varepsilon)$ dynamic are the two coordination equilibria (which are absorbing states).
- Let state 0 be the "All 0" ($x = 0$) coordination equilibrium and state 1 be the "All 1" ($x = 1$) coordination equilibrium.
- If $\frac{\mu_1^\varepsilon}{\mu_0^\varepsilon} \to 0$ as $\varepsilon \to 0$, then state 0 is stochastically stable.
- If $\frac{\mu_1^\varepsilon}{\mu_0^\varepsilon} \to \infty$ as $\varepsilon \to 0$, then state 1 is stochastically stable.

Proof

- Recall from Theorem 9.4 that:

$$\frac{\mu_1}{\mu_0} = \prod_{j=1}^{N} \frac{N-j+1}{j} \cdot \frac{\rho_{01}\left(F(\frac{j-1}{N}), \frac{j-1}{N}\right)}{\rho_{10}\left(F(\frac{j}{N}), \frac{j}{N}\right)}. \tag{1}$$

- Note that for $x > x^*$, strategy 1 is a best response and is chosen with prob. $1 - \frac{\varepsilon}{2}$. For $x < x^*$, strategy 0 is a best response.

- Substituting into (1), we get the expression for the $BRM(\varepsilon)$ dynamic:

$$\frac{\mu_1}{\mu_0} = \prod_{j=\lceil Nx^* \rceil}^{N} \frac{N-j+1}{j} \cdot \frac{1-\frac{\varepsilon}{2}}{\frac{\varepsilon}{2}} \prod_{j=1}^{\lfloor Nx^* \rfloor} \frac{N-j+1}{j} \cdot \frac{\frac{\varepsilon}{2}}{1-\frac{\varepsilon}{2}}. \tag{2}$$

Proof

- Rearranging:

$$\frac{\mu_1}{\mu_0} = \left(\frac{1-\frac{\varepsilon}{2}}{\frac{\varepsilon}{2}}\right)^{N-\lceil Nx^*\rceil+1} \left(\frac{\frac{\varepsilon}{2}}{1-\frac{\varepsilon}{2}}\right)^{\lfloor Nx^*\rfloor} \prod_{j=1}^{N} \frac{N-j+1}{j}. \quad (3)$$

Recognizing that $\lceil Nx^*\rceil = \lfloor Nx^*\rfloor + 1$, so that $N - \lceil Nx^*\rceil + 1 = N - \lfloor Nx^*\rfloor$, we can write:

$$\frac{\mu_1}{\mu_0} = \left(\frac{1-\frac{\varepsilon}{2}}{\frac{\varepsilon}{2}}\right)^{N-2\lfloor Nx^*\rfloor} \prod_{j=1}^{N} \frac{N-j+1}{j}, \quad (4)$$

which is proportional to:

$$\left(\frac{1-\frac{\varepsilon}{2}}{\frac{\varepsilon}{2}}\right)^{N-2\lfloor Nx^*\rfloor}. \quad (5)$$

Proof

- As $\varepsilon \to 0$:

$$\left(\frac{1-\frac{\varepsilon}{2}}{\frac{\varepsilon}{2}}\right) \quad \to \infty. \tag{6}$$

- Therefore $\frac{\mu_1}{\mu_0} \to \infty$ if $N > 2\lfloor Nx^* \rfloor$ and $\frac{\mu_1}{\mu_0} \to 0$ if $N < 2\lfloor Nx^* \rfloor$.

- For N sufficiently large, this is equivalent to:

  - $\frac{\mu_1}{\mu_0} \to \infty$ if $x^* < \frac{1}{2}$ —"All 1" is the unique stochastically stable state,

  - $\frac{\mu_1}{\mu_0} \to 0$ if $x^* > \frac{1}{2}$ —"All 0" is the unique stochastically stable state.

Pure Coordination

Corollary 10.2. Under the $BRM(\varepsilon)$ revision protocol, $\lim_{\varepsilon \to 0} \mu_1^\varepsilon = 1$ in the pure coordination game, i.e. "All 1" is the unique stochastically stable state.

Impure Coordination

- Now consider the following symmetric normal-form coordination game called **Stag Hunt**:

$$A = \begin{pmatrix} 2 & 2 \\ 0 & 3 \end{pmatrix}.$$

- Now $x^* = \frac{2}{3} > \frac{1}{2}$ is the mixed Nash equilibrium. Therefore, $x = 0$ is strictly risk-dominant, but Pareto <u>inefficient</u>.

Corollary 10.3. Under the $BRM(\varepsilon)$ revision protocol, $\lim_{\varepsilon \to 0} \mu_0^\varepsilon = 1$ in the stag hunt game, i.e. "All 0" (all hare) is the unique stochastically stable state.

Stochastic Stability

- Thus far, we have learned how to identify stochastically stable states for *reversible Markov processes*, for which a stationary distribution can be computed.

- Here we shall broaden the scope of our stochastic stability analysis and in doing so we shall require new concepts and techniques.

Regular Perturbed Markov Process

Definition. For each $\varepsilon \in [0, \varepsilon^*]$, let $\{X_t^\varepsilon\}$ be a Markov process. $\{X_t^\varepsilon\}$ is a *regular perturbed Markov process* if it is (a) irreducible for all $\varepsilon \in (0, \varepsilon^*]$ and (b) if for every $x, x' \in X$, $P_{xx'}^\varepsilon$ approaches $P_{xx'}^0$ at an exponential rate, i.e.

$$\lim_{\varepsilon \to 0} P_{xx'}^\varepsilon = P_{xx'}^0,$$

and

$$\text{if } P_{xx'}^\varepsilon > 0 \text{ for some } \varepsilon > 0\text{, then } 0 < \lim_{\varepsilon \to 0} P_{xx'}^\varepsilon / \varepsilon^{r(x,x')} < \infty$$

for some $r(x, x') \geq 0$.

Resistances

- The real number $r(x, x')$ is called the **resistance** (or cost) of the transition $x \to x'$.
- There cannot be two distinct exponents that satisfy the last condition, therefore $r(x, x')$ is uniquely identified.
- $P^0_{xx'} > 0$ if and only if $r(x, x') = 0$, i.e. transitions that occur under the unperturbed process have zero resistance.
- It is straightforward to show that the BRM(ε) and logit choice protocols generate regular Perturbed Markov processes.

Computing Stochastically Stable States

- Let $E_1, E_2, \ldots, E_K$ be recurrence classes of the unperturbed process ($\varepsilon = 0$).
- For each pair of distinct recurrence classes E_j and E_k a *jk*-path is a sequence of states $\zeta = (x_1, x_2, ..., x_q)$ that begins in E_j and ends in E_k.
- The resistance is the sum of the resistances of its edges, i.e. $r(\zeta) = r(x_1, x_2) + r(x_2, x_3) + ... + r(x_{q-1}, x_q)$.
- Let $r_{jk} = \min r(\zeta)$ be the least resistance over all *jk* paths.
- $r_{jk} > 0$ for all $j \neq k$ by the definition of a recurrence class.

Computing Stochastically Stable States

- Construct a complete directed graph of K vertices, one for each recurrence class.
- The weight on the directed edge $j \to k$ is r_{jk}.
- A tree rooted at vertex j (a *j-tree*) is a set of $K - 1$ directed edges such that *from every vertex different to j there is a unique directed path in the tree to j*.
- The resistance of a **rooted tree** is the sum of the resistance of its edges.
- The **stochastic potential** γ_j of E_j is the minimum resistance over all trees rooted at j.

Stochastic Potential & Stochastic Stability

Theorem 10.4. Let $\{X_t^\varepsilon\}$ be a regular perturbed Markov process and let μ^ε be the unique stationary distribution of $\{X_t^\varepsilon\}$ for each $\varepsilon > 0$.

Then $\lim_{\varepsilon \to 0} \mu^\varepsilon = \mu^0$ exists, and μ^0 is a stationary distribution of the unperturbed process $\{X_t^0\}$.

The stochastically stable states are precisely those states that are contained in the recurrence classes of $\{X_t^0\}$ having minimum stochastic potential.

Coordination Games Again

- Consider a 2×2 coordination game. There are two recurrence classes E_0 ('all 0') and E_1 ('all 1').
- With a certain sequence of nonbest responses, the perturbed process can transit between recurrence classes.
- Let $r(E_0, E_1)$ be the ***resistance*** of the transition $E_0 \to E_1$.
- Under the $BRM(\varepsilon)$ protocol, this is the minimum number of errors required for the process to transit from state E_0 to state E_1.

Coordination Games Again

- The corresponding trees are:

- Therefore, for N sufficiently large, we see that E_1 is stochastically stable if and only if it is risk dominant.

A Sufficient Condition for Stochastic Stability

- When there are many recurrence classes of the unperturbed Markov process $\{X_t^0\}$, the spanning tree method of finding a stochastically stable state can *sometimes* be difficult to implement.

- Consider the following alternative method introduced by (Ellison 1993) that provides a sufficient condition for stochastic stability (the spanning tree method provides both a necessary and sufficient condition).

- Let the unperturbed process have distinct recurrence classes $E_1, E_2, \ldots E_K$.

- Once it hits one of these recurrence classes, the unperturbed process locks in to that class.

The Radius-Coradius Theorem

- However, with a certain sequence of nonbest responses, the perturbed process can transit between recurrence classes.

- Once again, let $r(E, E_\ell)$ be the ***resistance*** (or cost) of the transition $E \to E_\ell$.

- Define the **radius** of recurrence class E by:

$$R(E) = min_\ell \; r(E, E_\ell)$$

  and the coradius of E by:

$$CR(E) = max_\ell \; r(E_\ell, E)$$

The Radius-Coradius Theorem

Theorem 10.5 If $R(E) > CR(E)$, then E is the unique stochastically stable class of $\{X_t^\varepsilon\}$.

The intuition for this is as follows:

- A lower bound on the probability of exiting recurrence class E is on the order of $\varepsilon^{R(E)}$.
- An upper bound on the probability of entering recurrence class E is on the order of $\varepsilon^{CR(E)}$.
- As $\varepsilon \to 0$, $\varepsilon^{R(E)-CR(E)} \to 0$ iff $R(E) > CR(E)$.
- In this case, the proportion of time spent in E is far, far greater than the proportion of time spent outside of E as the noise level becomes vanishingly small.